The Gender Pay Gap: Micro Sources and Macro Consequences*

Iacopo Morchio[†] Christian Moser[‡]

June 12, 2025

Abstract

Using linked employer-employee data from Brazil, we document a large gender pay gap due to women working at lower-paying employers. To interpret this fact, we develop an equilibrium search model with endogenous firm pay, amenities, and hiring. We provide a constructive proof of identification of all model parameters. The estimated model suggests that amenities are important for both men and women and that compensating differentials explain half of the gender pay gap. Equal-treatment policies partly close gender gaps but are not output- or welfare-improving.

Keywords: Wage Inequality, Amenities, Equilibrium Search Model, Linked Employer-Employee Data, Compensating Differentials, Taste-Based Discrimination, Monopsony Power

JEL Classification: E24, J16, J31, J32

*We thank the editor, an anonymous coeditor, and three anonymous referees for constructive feedback. We are grateful for discussions by Rajashri Chakrabarti, Rasmus Lentz, Andreas Mueller, Laura Pilossoph, and Rachel Schuh. We have benefited from comments by Mark Aguiar, Jim Albrecht, Jesper Bagger, Sydnee Caldwell, Carlos Carrillo-Tudela, Loukas Karabarbounis, Pat Kline, Kevin Lang, Guido Menzio, Andreas Mueller, Richard Rogerson, Yona Rubinstein, Ayşegül Şahin, Todd Schoellman, Isaac Sorkin, Gabriel Ulyssea, Chris Taber, Rune Vejlin, and Gianluca Violante. We also thank attendants at the 2018 and 2024 GEA Conferences, the 2019 AEA Annual Meeting, the 2019 Search and Matching Network Conference, the 2019 Bank of Italy/CEPR/IZA Conference, the 2019 and 2024 SED Annual Meetings, the 2019 and 2020 NBER Summer Institutes, the 2019, 2020, and 2024 Stanford SITE Workshops, the 2019 Dale T. Mortensen Centre Conference, the 2019 Boston Macro Juniors Conference, the 2019, 2020, and 2023 Columbia Junior Micro Macro Labor Conferences, the 2020 EALE/SOLE/AASLE World Conference in Berlin, the 2020 Annual USC Marshall Macro Day, the 2021 Federal Reserve System's Applied Microeconomics Conference, the 2021 LACEA-LAMES Meeting, the 2023 Annual Conference of AIEL, the 2024 NLS-E Conference at Princeton, the 2024 NBER Conference on Wage Dynamics in the 21st Century, the 2024 FRB of Atlanta's Annual Employment Conference, as well as seminar participants at Columbia, BU, CREST, Iowa State, Oslo, Bristol, Edinburgh, Vienna, Linz, the FRB of Philadelphia, Temple, the FRB of Minneapolis, the VMACS Junior Conference, Western Ontario, ASU, Emory, the FRB of Atlanta, the FRB of St. Louis, the Essex/RHUL/Bristol SaM Webinar, Banco de Portugal, CEMFI, NYU, UC Berkeley, Bonn, UCL, UW-Madison, UBC, FGV EESP, UT Austin, Houston, the FRB of Richmond, UCL, Georgetown, Cardiff, NHH, US Census, and Michigan for helpful comments. Ian Ho, Rachel Williams, and Giacomo Ricciardi provided outstanding research assistance. Moser gratefully acknowledges financial support from the Ewing Marion Kauffman Foundation, the Sanford C. Bernstein & Co. Center for Leadership and Ethics at Columbia Business School, INET, and the Hewlett Foundation. Moser also thanks the FRB of Minneapolis and the Heller-Hurwicz Economics Institute at the University of Minnesota for their generous hospitality during a significant share of the period of work on this project. Any errors are our own.

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

[‡]Columbia University, NBER, CEPR, and IZA. Email: c.moser@columbia.edu.

1 Introduction

The fact that women, compared to men, often work at lower-paying firms appears worrisome not only from a distributional perspective but also in terms of society's allocation of talent (Lentz and Mortensen, 2010; Hsieh et al., 2019). However, a holistic assessment of gender gaps requires looking beyond monetary pay and output. In this spirit, the empirical literature has convincingly established that nonpay job attributes are important in general (Hall and Mueller, 2018) and that women have a higher willingness to pay for amenities such as work flexibility in particular (Goldin, 2014). Thus, it seems natural to revisit the observed gender differences in sorting through the lens of a richer framework that takes into account firm heterogeneity in both pay and nonpay attributes.

This paper's goal is to identify the microeconomic sources of the gender pay gap in order to assess their macroeconomic consequences. Our analysis proceeds in four steps. First, we establish empirical facts on gender differences in sorting across firm pay. Second, we interpret these facts by developing an equilibrium search model with endogenous firm pay, amenities, and employment. Third, we provide a constructive proof of identification of all model parameters based on linked employer-employee data. Finally, we use the estimated model to shed light on the structure of pay and amenities across firms, to quantify the role of compensating differentials, and to simulate counterfactual equal-pay and equal-hiring policies. Altogether, our results highlight an important role of nonpay job attributes in explaining gender differences in pay and employment across firms.

In the first step, we leverage rich linked employer-employee data from Brazil, which has a significant gender pay gap of 13.3 log points, making this an interesting context to study gender inequality. An advantage of studying Brazil is that it offers remarkably detailed information on gender-relevant labor market variables such as workers' education, occupation, tenure, work hours, and employment histories with information on parental leaves. To estimate firm pay differences for identical workers by gender, we follow Card et al.'s (2016) extension of the seminal two-way fixed effects (FEs) framework by Abowd, Kramarz, and Margolis (1999, henceforth AKM). Controlling for worker heterogeneity, we find a gender firm pay gap of 11.3 log points (i.e., 85.0% of the overall gender pay gap), mostly due to women sorting to lower-paying firms relative to men.

In the second step, one of our key contributions is to develop an equilibrium search model based on the seminal framework by Burdett and Mortensen (1998), with endogenous gender differences in firm pay, amenities, and hiring. Workers differ in their gender and ability and search for jobs both in and out of employment. Firms differ in their productivity, gender preferences, and amenity costs and face a convex increasing vacancy cost schedule. Our model accommodates several competing explanations for the gender pay gap, including gender-specific *compensating differentials* (Rosen, 1986), *taste-based discrimination* (Becker, 1971), and *labor market frictions* (Manning, 2003). The model features pay and utility dispersion both within and between worker types. The existence of pay differences, however, is neither necessary nor sufficient for the existence of utility differences. Job-to-job transitions may entail pay declines. Firms can discriminate based on their pay offers, amenities, or hiring decisions. Men and women climb different firm ladders. Finally, even nondiscriminatory firms without regard for gender may treat women differently as a best response to the labor market environment. To sum up, the gender-specific distributions of firm pay, amenities, and hiring are all jointly determined in equilibrium, making it a formidable task to isolate each of the model's features.

In the third step, we make a methodological contribution by separately identifying all model parameters, including labor market objects, gender-specific firm types, and economy-wide elasticities of the vacancy and amenity cost functions. We demonstrate that our equilibrium model admits a log-additive wage equation with separable worker and firm components, akin to Card et al.'s (2016) extension of AKM, which is helpful for two reasons. First, it enables us to interpret the gender-specific employer FEs from our empirical analysis through the lens of the structural model. Second, it allows us to control for gender-specific selection based on ability (Mulligan and Rubinstein, 2008). To identify labor market objects, we follow a revealed-preferences argument based on worker flows. Next, we identify gender-specific firm types. From a bird's eye view, we leverage the intuition that firms' unobserved surplus maps into hiring decisions. We show how to invert this model mapping to recover firm-level utility offers. By comparing those to firm pay, we back out the gender-specific amenity values at each firm. We then estimate employer preferences over gender by comparing equilibrium outcomes between men and women at the same firm. Finally, we show how to pin down the economy-wide elasticity parameters guiding vacancy and amenity costs based on aggregate moments. A noteworthy aspect of our approach is that we do not rely on any distributional assumptions in identifying a large number of model parameters in an environment with firm heterogeneity in both observed pay and unobserved amenities (cf. Bontemps et al., 1999, 2000).

In the fourth step, our substantive contribution is to use the estimated model to revisit the observed patterns of gender-specific sorting across firms. By identifying rich gender-firm heterogeneity, our model provides several novel insights regarding the structure of firm pay and amenities. We find that amenities play an important role for both genders, with a mean amenity share of 48.8% for men and 52.2% for women. However, higher-ranked employers for men mostly offer higher pay, while for women they offer higher amenities. Compensating differentials explain the lion's share of overall firm pay dispersion, with utility dispersion only accounting for 4.4% of pay dispersion for men and 3.6% of that for women. This is all the more striking given that we find significant labor market frictions in Brazil. Taking into account gender differences in amenities, the gender gap in total compensation becomes 4.6 log points (i.e., 40.7 percent of the gender pay gap). Altogether, these results suggest that compensating differentials are central to understanding gender-specific sorting across firms.

Given the importance of firm-level amenities, we return to our motivation regarding the micro sources and macro consequences of the gender pay gap. We leverage the equilibrium nature of our model to decompose gender gaps in pay, amenities, and utility. As a result of shutting down firm heterogeneity in amenities, the gender pay gap closes by 47.6 percent, largely due to a relocation of women toward formerly male-dominated firms. Our equilibrium framework naturally lends itself to an analysis of counterfactual equal-pay and equal-hiring policies. The bottom line is that both policies close part of the gender pay gap but lower worker welfare due to their adverse incentive effects on firms' pay, amenity, and hiring decisions. Thus, our results underline the importance of studying such policies in general equilibrium.

A cautionary note regarding the interpretation of our results is in order. We have quantified the contribution of gender-specific preferences over workplace amenities toward gender differences in sorting and pay across firms. We think of workers' preferences over amenities, such as those relating to work-life balance, as capturing the induced demand for certain workplace or employer characteristics due to household arrangements and societal norms that differ across men and women (Goldin, 2014, 2023; Cubas et al., 2021, 2023). In reality, women may seek out high-amenity jobs not because of preferences but because of constraints imposed by nonmarket factors. Without such accommodations, some women may not be able to work at all, given their responsibilities. Our estimates of amenity valuations and associated compensating differentials should be interpreted in this broader sense.

Related Literature. A burgeoning literature highlights the role of employer heterogeneity in explaining empirical pay dispersion (Card et al., 2013, 2018; Alvarez et al., 2018; Song et al., 2018). Our work builds on Card et al.'s (2016) extension of the seminal framework by AKM with gender-specific employer pay components.[1] Relative to their work, we make three contributions. First, we offer a microfoundation for the wage equation in Card et al. (2016) and empirical patterns of worker sorting

[1]Complementary work on gender gaps in firm pay includes Sorkin (2017), Coudin et al. (2018), Bruns (2019), Barth et al. (2021), Casarico and Lattanzio (2022), Cruz and Rau (2022), Palladino et al. (2022), Lentz et al. (2023), and Vattuone (2023).

across firms by gender based on a tractable *equilibrium* model. Second, whereas Card et al. (2016) rationalize gender gaps in employer pay through differences in exogenous bargaining parameters across the sexes, we identify more than one reason behind them: compensating differentials (Rosen, 1986), taste-based discrimination (Becker, 1971), and labor market frictions (Manning, 2003). Third, we simulate a series of counterfactual experiments, including equal-treatment policies, for which the equilibrium nature of our model is crucial as firms adjust their pay, amenities, and hiring.

Our equilibrium search model builds on the influential framework by Burdett and Mortensen (1998), which has been developed in different directions by Bontemps et al. (1999, 2000), Moscarini and Postel-Vinay (2013, 2018), Meghir et al. (2015), Lise and Robin (2017), Bagger and Lentz (2019), Bilal et al. (2022), and Engbom and Moser (2022). In these models, firms are heterogeneous only in productivity. Consequently, workers agree on a firm ranking, and wage gains go hand in hand with efficiency gains. Departing from this tradition, we develop a model with richer firm heterogeneity, which we identify based on linked employer-employee data. Specifically, we allow firms to differ not just in productivity but also in preferences over gender and amenity costs. As a result, firm pay, amenities, and hiring are jointly determined in equilibrium. In spite of this added complexity, we provide a constructive proof of identification of all model parameters, including labor market objects, gender-specific firm types, and economy-wide elasticities of the vacancy and amenity cost functions. By allowing for this richness, several novel insights emerge regarding gender-specific compensation structures across firms. Compared to the aforementioned work, our model has radically different implications for the interpretation of empirical pay dispersion in relation to misallocation and welfare.[2]

There is ample empirical evidence that job amenities matter for labor market outcomes (Hamermesh, 1999; Pierce, 2001; Hall and Mueller, 2018; Sockin, 2022; Maestas et al., 2023), especially for women (Goldin, 2014, 2023; Juhn and McCue, 2017; Mas and Pallais, 2017, 2019; Wiswall and Zafar, 2017; Chen et al., 2020; Cubas et al., 2021, 2023; Kim et al., 2024). However, quantifying the role of firm-level amenities is complicated by both data limitations and theoretical challenges. Regarding data, many firm-level amenities are unobserved to the analyst, so their valuations must be inferred under additional assumptions.[3] Regarding theory, valuations of firm-level amenities are not easily

[2]For example, Bagger et al. (2014) find *"large marginal output and wage gains associated with labor reallocation"* (p. 5) and conclude that *"the finding that misallocation persists under these circumstances suggests that labor market frictions are important barriers to growth"* (p. 1). Furthermore, Lentz and Mortensen (2010) conclude that *"the reallocation of employment from less to more productive firms will yield efficiency gains"* (p. 577) and that *"workers will find it in their interest to seek out higher-paying employers"* (p. 577). Neither of these conclusions necessarily follows in our model with heterogeneity in employer amenities.

[3]Several works have estimated the values of *specific* job amenities like employer health insurance (Dey and Flinn, 2005), job security (Jarosch, 2023), fatality risk (Lavetti and Schmutte, 2018), commuting costs (Flemming, 2020), location (Heise and Porzio, 2023), family friendliness (Hotz et al., 2018; Xiao, 2023), and working conditions (Bonhomme and Jolivet, 2009).

obtained in existing models, especially in the presence of frictions.[4]

Three related papers that also use linked employer-employee data stand out in this context. First, Taber and Vejlin (2020) develop a model with comparative advantage, search frictions, and compensating differentials. Despite its rich features, some of which we abstract from, they nonparametrically identify nearly all model parameters. A notable exception is workers' bargaining power, which is not identified since their revealed-preferences approach pins down *ordinal* but not *cardinal* utility.

Second, Sorkin (2018) embeds discrete choice within a search model, which can identify what he terms the *"Rosen motive"* of compensating differentials, capturing amenities dispersion conditional on a firm's value. In contrast, his model cannot identify what he terms the *"Mortensen motive,"* capturing amenities dispersion correlated with firm values. This is because, in his model, the variance of utility is not pinned down due to an arbitrary but necessary normalization of the scale parameter of the type-I extreme value distribution guiding idiosyncratic utility draws.[5] By allowing firms to choose hiring subject to a commonly used vacancy cost function, the parameters of which we identify, our general-equilibrium model recovers the entire distribution of firm values and amenities, which would not have been possible in the partial-equilibrium models of Taber and Vejlin (2020) or Sorkin (2018).

Third, Lamadon et al. (2022) develop an equilibrium labor market model featuring compensating differentials and rent sharing. Their model is frictionless, which implies that observed wages reflect productive attributes and idiosyncratic preferences following a nested logit structure. We complement their work by developing a distribution-free model that features search frictions, yet remains point-identified.[6] Interestingly, we find a similarly important role for amenities and compensating differentials as they do, even under large search frictions. In addition, our model has several unique implications for gender inequality and equal-treatment policies not previously considered.

Outline. The paper is structured as follows. Section 2 describes the data. Section 3 presents motivating empirical facts. Section 4 develops the equilibrium model. Section 5 provides a constructive identification proof. Section 6 shows estimation results. Section 7 analyzes gender-specific compensation structures across employers. Section 8 simulates counterfactuals. Finally, Section 9 concludes.

Theoretical work by Hwang et al. (1998), Lang and Majumdar (2004), and Albrecht et al. (2018) develop models with *nonspecific* job amenities. Relatedly, Sullivan and To (2014), Hall and Mueller (2018), Jung and Kuhn (2019), Luo and Mongey (2019), and Hsieh et al. (2019) estimate nonspecific amenity values using individual-level (i.e., not linked employer-employee) data.

[4] In this sense, our model parallels Boerma and Karabarbounis (2021) who identify unobserved home production values.

[5] See Appendix C.8 for a detailed comparison between our model and that in Sorkin (2018).

[6] Lamadon et al. (2022) note that *"while incorporating [search frictions] would be interesting, it would also present severe challenges to identification, especially if one allows for two-sided heterogeneity"* (p. 210). We formally identify our model, though it is worth noting that our notion of worker heterogeneity and our production structure are substantially simpler than theirs.

2 Data Description

Dataset. Our main data source is the Brazilian linked employer-employee register *Relação Anual de Informações Sociais* (RAIS), which covers all workers at tax-registered employers. Starting in 2007, there is information on worker absences, including parental leaves. In 2015, the country entered a severe recession. Therefore, we restrict attention to the eight-year period from 2007 to 2014.

Variables. The data contain unique identifiers for workers and establishments (henceforth "employers" or "firms"). For each job spell, we observe the start and end dates, mean monthly earnings (henceforth "wage" or "pay"), as well as the worker's gender, educational attainment in nine categories, worker age in years, tenure in years, contractual work hours per week, five-digit sector codes with 672 categories, municipality codes with 5,565 categories, and six-digit occupation codes with 2,383 categories. We exploit the full panel of the data going back to 1985 together with the tenure variable to impute actual—not just potential—formal-sector work experience in years.[7]

Sample Selection. We select workers between the ages of 18 and 54 earning at least the federal minimum wage. For each worker-year combination, we keep the highest paid among all longest employment spells. We then iteratively drop singleton observations defined by gender-employer combinations and worker identifiers. We also impose a minimum employer size threshold of 10 non-singleton workers—i.e., workers who are observed at least one more time at a future date.[8] Finally, we require that employers appear in our sample in at least four out of the eight years. Together, these selection criteria leave us with a set of reasonably large and stable employers for which pay policies can be credibly estimated with minimal limited-mobility bias.[9] In order to separately identify worker and employer pay components as well as employer ranks based on worker flows, we focus on observations in the largest strongly connected set, which requires flows into and out of all employers in the set. Our selection criteria do not substantially alter the raw gender pay gap.

Summary Statistics. Table 1 presents summary statistics on our final sample.[10] The pooled sample comprises 267.3 million worker-years, including 56.3 million unique workers and 607.0 thousand

[7] The distinction between actual and potential experience is important given Brazil's sizable informal sector, as shown in Figure A.1 in Appendix A.1, though gender differences are small and thus explain little of the empirical gender pay gap.

[8] While the RAIS data cover only Brazil's formal sector, the employer size restriction implies that the vast majority of informal employers would in any case be excluded from our analysis (Ulyssea, 2018; Dix-Carneiro et al., 2021).

[9] See also Andrews et al. (2008, 2012), Kline et al. (2020), Borovičková and Shimer (2020), and Bonhomme et al. (2023).

[10] Appendix A.2 compares summary statistics based on the raw data (Appendix Table A.1), the selected sample (Appendix Table A.2), the connected set (Appendix Table A.3), and comparisons between them (Appendix Tables A.4–A.5).

unique employers. Around 38.2% of these observations are for women who are more likely to be White, more educated, older, work at significantly larger employers, work shorter hours, and have longer tenure. Importantly, the raw gender pay gap in our sample is 13.3 log points.

Table 1. Summary statistics, 2007–2014

	Overall	Men	Women
Mean log real monthly earnings (std. dev.)	7.211 (0.693)	7.262 (0.697)	7.129 (0.679)
Mean years of education (std. dev.)	11.1 (3.3)	10.4 (3.3)	12.1 (2.9)
Mean years of age (std. dev.)	33.6 (9.4)	33.5 (9.4)	33.8 (9.4)
Mean employer size (std. dev.)	2,815 (16,418)	1,774 (11,509)	4,497 (22,059)
Mean contractual work hours (std. dev.)	41.7 (5.1)	42.6 (3.9)	40.3 (6.4)
Mean years of tenure (std. dev.)	3.9 (5.6)	3.6 (5.2)	4.5 (6.1)
Share Nonwhite	0.378	0.409	0.327
Share female	0.382		
Mean log gender earnings gap	0.133		
Number of worker-years	267,318,328	165,149,632	102,168,696
Number of unique workers	56,297,308	33,761,656	22,535,652
Number of unique employers	607,029	403,585	203,444

Note: This table reports summary statistics for workers in the final sample, separately for the overall population, for men only, and for women only. Since information on race is missing for a significant number of observations, conditional means are reported for the share of Nonwhite workers. *Source:* RAIS, 2007–2014.

3 Empirical Gender Pay Gaps and Employer Heterogeneity

The goal of this section is to highlight the roles of employer pay heterogeneity, worker sorting across employers, and workplace amenities in relation to the gender pay gap.

3.1 Measuring Gender-Specific Employer Pay

We start by estimating a variant of Card et al.'s (2016) extension of the seminal two-way FEs framework due to AKM, which allows for gender-specific employer pay components. Formally, we model log earnings of individual i in year t working at employer $j = J(i,t)$, denoted by $\ln w_{ijt}$, as

$$\ln w_{ijt} = \alpha_i + \psi_{G(i)j} + X_{it}\beta_{G(i)} + \varepsilon_{ijt}, \tag{1}$$

where α_i is a person FE; $\psi_{G(i)j}$ is a gender-specific employer FE for workers of gender $G(i) \in \{M, F\}$; X_{it} is a vector of time-varying worker characteristics including a set of restricted education-age dummies as well as dummies for hours, occupation, tenure, actual experience, and education-year combinations, all subject to gender-specific returns $\beta_{G(i)}$; and ε_{ijt} is a residual.[11] By including person FEs, we control for selection of men and women across employers based on time-invariant worker pay characteristics. Our focus lies in the distribution of gender-specific employer FEs ψ_{Mj} and ψ_{Fj}.

Any two-way FEs model requires a normalization of the level of employer FEs relative to person FEs. In our case, the inclusion of gender-specific employer FEs in equation (1) requires two normalizations—one for each gender.[12] In previous work, Card et al. (2016) and Gerard et al. (2021) normalized the FEs of employers in each connected set to have mean zero in the restaurant and fast-food sector, which they argued has a low surplus. However, those papers were concerned only with employer heterogeneity in pay. Our setting with gender-specific amenities and compensating differentials calls for a different normalization of employer FEs, which we derive in Appendix D.4 based on the theoretical framework in Section 4. Guided by our theory, we would like to normalize employer pay for a subset of employers that (i) rank near the bottom in the utility ladder for both genders; (ii) treat men and women as near-perfect substitutes in production; and (iii) provide a similar (e.g., close to zero) amenity value to men and women who they employ. We impose condition (i) based on our empirical estimates of revealed-preference ranks derived in Section 5.3. We further assume that workers in the restaurant and fast-food sector satisfy conditions (ii) and (iii), which we think of as a reasonable assumption given the nature of jobs in this sector.

We now turn to our object of interest in equation (1)—namely, the gender-specific employer FEs.[13] Panel A of Figure 1 plots the distribution of employer FEs by gender. The distribution for women has a 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 employers. The distribution is relatively dispersed compared to its mean of 2.4 log points. Altogether, this evidence suggests that the sorting of women into lower-paying employers is a significant source of gender pay differences.

To dissect the structure of pay and the relative contribution of employer heterogeneity, Table 2 presents a variance decomposition of log earnings based on both the conventional plug-in estimator and the leave-out estimator by Kline, Saggio, and Sølvsten (2020, henceforth KSS), which uses a jack-

[11]To simultaneously identify age, time, and worker FEs, we restrict the age-pay profile be flat around ages 45–49. We view this as an attractive alternative to the approach advocated by Card et al. (2018), since it allows us to verify that our restriction leads to smooth education-age FEs around this age window. See Appendix Figure B.6 for details.

[12]To see this, note that we could transform $\alpha_i \mapsto \alpha_i + k$ and $\psi_{G(i)j} \mapsto \psi_{G(i)j} - k$ for all workers i of a given gender $g = G(i)$, for any constant $k \in \mathbb{R}$, without changing the sum of the components in equation (1).

[13]Appendix B.1 shows auxiliary results relating to the AKM equation, including estimated gender-specific hours FEs (Figure B.1), occupation FEs (Figure B.2), actual-experience FEs (Figure B.3), tenure FEs (Figure B.4), education-year FEs (Figure B.5), and education-age FEs (Figure B.6). Further tests of the log-additivity and exogenous mobility assumptions of similar specifications and data are presented in Alvarez et al. (2018), Engbom and Moser (2022), and Gerard et al. (2021).

Figure 1. Predicted AKM employer FEs for women and men

A. Gender-specific employer FE distributions B. Distribution of within-employer FE differences

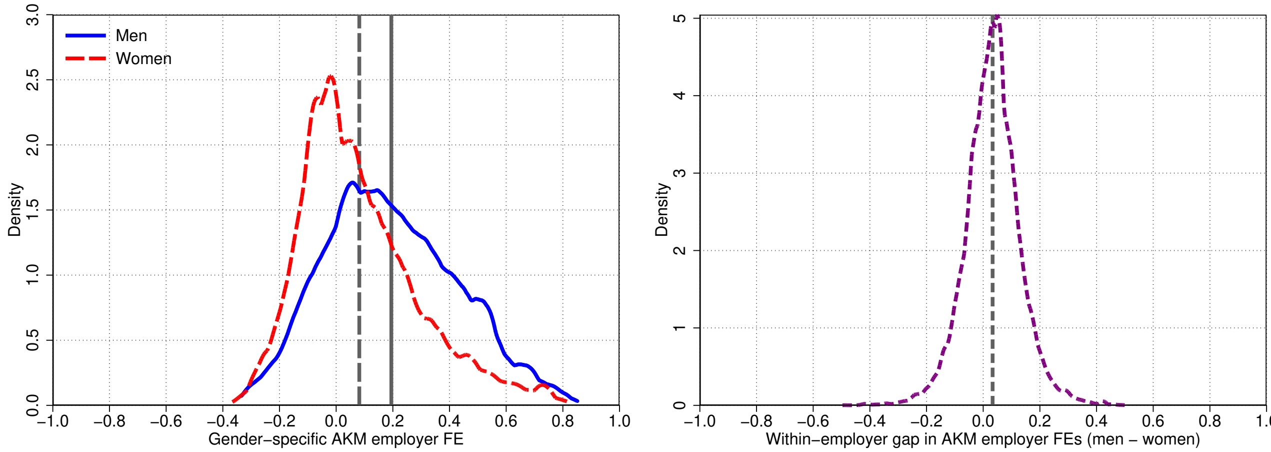

Note: This figure shows kernel density plots of estimated gender-specific employer FEs based on estimating earnings equation (1). Panel A shows the distributions of gender-specific employer FEs ψ_{gj} separately by gender. Panel B shows the distribution of within-employer FE differences $\psi_{Mj} - \psi_{Fj}$ weighted by total employment. Vertical patterned lines show the means of the respective distributions. *Source:* RAIS, 2007–2014.

knife correction for limited-mobility bias.[14] Men have a slightly higher raw (conditional) variance of earnings, with 49.7 (25.8) log points, compared to 48.2 (25.0) log points for women. For both genders, the largest plug-in (leave-out) variance component is due to estimated worker FEs, which account for 23.3% (37.7%) for men and 24.3% (39.6%) for women. Importantly for us, the plug-in (leave-out) estimates of employer FEs account for 13.0% (25.0%) of the variance of log earnings for men and 11.5% (22.2%) of that for women. The close similarity of the plug-in versus leave-out estimates of the variance of employer FEs suggests that the latter are precisely estimated, alleviating concerns about limited-mobility bias in our sample. Overall, these estimates suggest that employer heterogeneity explains a substantial share of earnings dispersion for both genders. The plug-in (leave-out) estimate of the correlation between person and employer FEs is around 21.2% (24.5%) for men and 25.5% (29.7%) for women. Finally, the plug-in (leave-out) estimate of the coefficient of determination or R^2 is upward of 92.1% (77.7%).

3.2 Gender Differences in Sorting Across Employer Pay and Amenities

Our starting point is a total gender gap in employer FEs of 11.3 log points—see the last row of Table 2. It is worth noting that this constitutes 85.0% of the overall gender pay gap of 13.3 log points from Table

[14] Engbom and Moser (2022) show that the leave-out estimator by KSS delivers substantially similar results for a sample of men across time periods from 1994–2018 in the same RAIS data from Brazil. For a discussion of limited-mobility bias and alternative ways to address it, see Bonhomme et al. (2019), Borovičková and Shimer (2020), and Bonhomme et al. (2023).

Table 2. Variance decompositions based on plug-in and leave-out estimates

	Men		Women	
	Plug-in	Leave-out	Plug-in	Leave-out
Variance of log earnings	0.497	0.258	0.482	0.250
Variance components:				
Employer FEs (%)	0.065 (13.0%)	0.064 (25.0%)	0.056 (11.5%)	0.055 (22.2%)
Person FEs (%)	0.116 (23.3%)	0.097 (37.7%)	0.117 (24.3%)	0.099 (39.6%)
Correlation	0.212	0.245	0.255	0.297
R^2	0.921	0.777	0.929	0.793
Mean employer FE	0.197	0.197	0.081	0.081

Note: This table shows the variance components of log earnings based on equation (1). The variance components correspond to the variance decomposition $Var(\ln w_{ijt}) = Var(\alpha_i) + Var(\psi_{G(i)j}) + Var(X_{it}\beta_{G(i)}) + 2\sum Cov(\cdot) + Var(\varepsilon_{ijt})$. The "plug-in" columns refer to conventional plug-in estimates, while the "leave-out" columns refer to estimates that correct for limited-mobility bias, following KSS. The variances shown in the "plug-in" columns are the variances of raw log earnings, while those in the "leave-out" columns are the variances of residualized log earnings conditional on the same controls as those in equation (1). *Source:* RAIS, 2007–2014.

1 above. Such an important role for firms in explaining the gender pay gap is striking in comparison to previous work by Card et al. (2016, Table III) who report that around 21% of the gender wage gap in Portugal is due to the contribution of firm components. That firms play a larger role in our context can be rationalized by the baseline gender pay gap increasing significantly after the inclusion of Mincerian controls for education as well as the overall importance of firms in explaining pay dispersion in Brazil (Alvarez et al., 2018; Firpo and Portella, 2019; Engbom and Moser, 2022).

We closely follow Card et al. (2016) in dissecting the employer pay gap into parts between versus within employers.A Kitagawa-Oaxaca-Blinder decomposition allows us to write the total gender gap in employer FEs as

$$\underbrace{\left(\mathbb{E}\left[\psi_{MJ(i,t)}\middle| G(i)=M\right]-\mathbb{E}\left[\psi_{MJ(i,t)}\middle| G(i)=F\right]\right)}_{\text{between-employer gap}}+\underbrace{\mathbb{E}\left[\psi_{MJ(i,t)}-\psi_{FJ(i,t)}\middle| G(i)=F\right]}_{\text{within-employer gap}}, \tag{2}$$

where $\mathbb{E}[\cdot]$ is the mean operator across individuals i and years t. Equation (2) decomposes the gender gap in employer FEs into two terms. The *between-employer pay gap* is the gender-weighted difference in mean male-employer FEs. It reflects differences in pay between men and women that are due to their different allocations across employers. The *within-employer pay gap* is the mean difference in gender-specific employer FEs weighted by the distribution of women. It reflects differences in pay between women and men at the same employer.[15]

[15] To see that the between-employer gap is invariant to the normalization of gender-employer FEs, note that for any $k \in \mathbb{R}$, $\mathbb{E}[\psi_{gJ(i,t)} + k|G(i) = M] - \mathbb{E}[\psi_{gJ(i,t)} + k|G(i) = F] = \mathbb{E}[\psi_{gJ(i,t)}|G(i) = M] - \mathbb{E}[\psi_{gJ(i,t)}|G(i) = F]$ for $g \in \{M, F\}$. The

Results from the decomposition in equation (2) are shown in Table 3. Out of the total gender gap in employer FEs of 11.3 log points, a majority share of 78.7% is attributed to the between-employer pay gap. This suggests that women, compared to men, systematically work at lower-paying employers and that gender-specific sorting accounts for most of the gender pay gap.[16] These findings accord well with other contexts: Combining the within-employer pay gap of 2.4 log points in Table 3 with the mean of men's employer FEs of 19.7 log points in Table 2, we conclude that women have a relative rent-sharing disadvantage of $2.4/19.7 = 12.3\%$, compared to the 10% rent-sharing disadvantage reported by Card et al. (2016, p. 618) for Portugal.

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

	Between-employer gap		Within-employer gap	
Gender gap in employer FEs	Level	Share (%)	Level	Share (%)
0.113	0.089	78.7	0.024	21.3

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gap in employer FEs into a between-employer gap and a within-employer gap. The decomposition corresponds to equation (2) and uses men's employer FEs for computing the between-employer component. An alternative decomposition using women's employer FEs for computing the between-employer component is presented in Table B.1 in Appendix B.2. *Source:* RAIS, 2007–2014.

While there are many causes behind gendered sorting across employers, we are particularly interested in quantifying the extent to which both pay and nonpay employer attributes contribute toward the gender pay gap. To this end, we develop an empirical equilibrium model of firm heterogeneity in pay, amenities, and employment.[17]

4 Equilibrium Model of Employer Pay, Amenities, and Size

In this section, we develop an equilibrium search model based on the seminal framework by Burdett and Mortensen (1998) that features endogenous firm pay, amenities, and hiring that may differ across workers of different genders. Our main contribution is to use such a framework to shed new light on the fundamental drivers behind gender gaps in pay, amenities, and employment across firms.

within-employer gap, on the other hand, depends on the normalization of gender-employer FEs, as discussed in the text.

[16] An alternative decomposition using women's employer FEs for computing the between-employer component is presented in Appendix Table B.1. Appendix Figure B.7 illustrates the estimates underlying the two decompositions.

[17] Beyond labor-demand-side factors related to firms' pay and nonpay attributes, labor-supply-side factors could be important. In our structural analysis, we allow for gender differences in labor market attachment and job mobility. Our framework is flexible enough to allow for the separate analysis of more granular population subgroups—for example, different education groups or worker groups split by age group or parental status.

4.1 Workers

Workers are infinitely lived, risk-neutral, and discount the future at rate ρ. They differ in their *gender* $g \in \{M, F\}$ and *ability* $z > 0$, with associated measure $\mu_{gz} \geq 0$ satisfying $\sum_g \int_z \mu_{gz}\, dz = 1$.

Job Search. Workers find themselves either employed or nonemployed.[18] From both states, they engage in random job search within labor markets segmented 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, a worker of type (g, z) at firm j receives flow utility from total compensation $x = w + \beta_g(j)a$, where $w \geq 0$ is the monetary wage, $a \geq 0$ is a workplace amenity, and $\beta_g(j) > 0$ is the relative preference weight on amenities for workers of gender g at firm j.[19] While nonemployed, a worker of type (g, z) receives flow utility from home production $x = b_g z$ for some $b_g \in \mathbb{R}$.

Workers receive voluntary job offers at rate λ^U_{gz} from nonemployment and at rate λ^E_{gz} from employment. While those job offers admit free disposal, workers also receive involuntary job offers at rate λ^G_{gz}, which we think of as capturing idiosyncratic reasons for switching jobs. We write $\lambda^E_{gz} = s^E_g \lambda^U_{gz}$ and $\lambda^G_{gz} = s^G_g \lambda^U_{gz}$, where $s^E_g > 0$ and $s^G_g \geq 0$ are the relative hazards of voluntary and involuntary on-the-job offers, which may differ by gender g.

A job offer is an opportunity to work at a firm with wage w and amenity level a. Since workers of type (g, z) rank a firm j according to its flow utility $x = w + \beta_g(j)a$, their decisions depend only on the flow-utility offer distribution $F_{gz}(x)$ and not additionally on the distribution of (w, a).

Jobs are endogenously terminated when a worker of type (g, z) with flow utility x accepts a higher-utility job at rate $\lambda^E_{gz}(1 - F_{gz}(x))$. Jobs are exogenously terminated either when the worker moves into nonemployment at rate δ_g or when the worker relocates to a randomly drawn job at rate λ^G_{gz}.[20]

[18] Throughout the paper, we think of "nonemployed" workers in the model as capturing real-world 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. When 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.

[19] With "preferences," we have in mind the induced demand for certain workplace amenities and other employer characteristics due to household arrangements and societal norms that differ across men and women (Goldin, 2014, 2023).

[20] In Appendix C.9, we present a version of our model with heterogeneity in firm-level separation rates, as in Jarosch (2023), which we assume to be nonincreasing in firm ranks. Later, we discuss the identification of this more general model and its quantitative implications for our estimates.

Value Functions. The Hamilton-Jacobi-Bellman (HJB) equation for an employed worker of type (g, z) in a job with flow utility x is

$$\rho S_{gz}(x) = x + \lambda^E_{gz} \int_{x' \geq x} \left[S_{gz}(x') - S_{gz}(x) \right] dF_{gz}(x') + \lambda^G_{gz} \int_{x'} \left[S_{gz}(x') - S_{gz}(x) \right] dF_{gz}(x') + \delta_g \left[W_{gz} - S_{gz}(x) \right], \tag{3}$$

where $S_{gz}(x)$ is the value of employment of a worker of type (g, z) in a job offering flow utility x and W_{gz} is the value of nonemployment of a worker of type (g, z).

Analogously, the HJB equation for a nonemployed worker of type (g, z) is

$$\rho W_{gz} = b_g z + (\lambda^U_{gz} + \lambda^G_{gz}) \int_{x'} \max \left\{ S_{gz}(x') - W_{gz}, 0 \right\} dF_{gz}(x'). \tag{4}$$

Policy Functions. Strict monotonicity of $S_{gz}(x)$ in x implies that optimal job acceptance of the nonemployed follows a threshold rule with reservation flow utility $\phi_{gz} \in \mathbb{R}$ such that a nonemployed worker accepts an offer if $x \geq \phi_{gz}$ and rejects it otherwise. Clearly, ϕ_{gz} equals the sum of the flow utility in nonemployment plus the forgone option value of receiving job offers while nonemployed:

$$\phi_{gz} = b_g z + (\lambda^U_{gz} - \lambda^E_{gz}) \int_{x' \geq \phi_{gz}} \frac{1 - F_{gz}(x')}{\rho + \delta_g + \lambda^G_{gz} + \lambda^E_{gz} \left[1 - F_{gz}(x') \right]} \, dx'. \tag{5}$$

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

Nonemployment. The steady-state nonemployment rate for each worker type is

$$u_{gz} = \frac{\delta_g}{\delta_g + \lambda^U_{gz} + \lambda^G_{gz}}. \tag{6}$$

Utility Distribution. The cross-sectional distribution of flow utilities for each worker type is

$$G_{gz}(x) = \frac{F_{gz}(x)}{1 + \kappa^E_{gz} \left[1 - F_{gz}(x) \right]}, \tag{7}$$

where $\kappa^E_{gz} \equiv \lambda^E_{gz} / (\delta_g + \lambda^G_{gz})$ governs the effective speed of climbing the firm ladder in utility space.

4.2 Firms

Firms differ in four dimensions: *productivity* $p > 0$, as in Burdett and Mortensen (1998), a set of *gender wedges* $\tau_g \in \mathbb{R}$ representing the firm's disutility from employing workers of gender g, as in Becker (1971), a set of *preference weights on amenities* $\beta_g > 0$ for workers of each gender g, as in Goldin (2014), and a set of *amenity cost shifters* $c_g^{a,0} > 0$ for each gender g, as in Hwang et al. (1998). Thus, a firm's type is $(p, \{\tau_g\}_g, \{\beta_g\}_g, \{c_g^{a,0}\}_g)$, which we assume is continuously distributed with CDF $\Gamma(\cdot)$.

Wages, Amenities, and Vacancies. Firms deliver value through wages w and amenities a.[21] The flow cost of providing amenity level $a \geq 0$ to workers of type (g, z) is paid per worker, as in Hwang et al. (1998), and given by

$$c_{gz}^a(a) = c_g^{a,0} \frac{(a/z)^{\eta^a}}{\eta^a} z, \tag{8}$$

where $\eta^a > 1$ is the economy-wide amenity cost elasticity. This formulation is consistent with amenities provided as piece rates relative to worker ability z, as for parental leaves and paid time off.

To hire workers, a firm posting $v \geq 0$ job vacancies for workers of type (g, z) pays a flow cost

$$c_{gz}^v(v) = c_g^{v,0} \frac{v^{\eta^v}}{\eta^v} z, \tag{9}$$

where $c_g^{v,0} > 0$ is a gender-specific vacancy cost shifter and $\eta^v > 1$ is the economy-wide vacancy cost elasticity. This formulation is consistent with recruiting costs being denominated either in terms of new hires' onboarding time or in terms of identical incumbent workers' time spent on recruiting.

Production. A firm with productivity p employing $\{l_{gz}\}_{gz}$ workers of each type produces output

$$y(p, \{l_{gz}\}_{gz}) = p \sum_g \int_a z l_{gz} \, dz. \tag{10}$$

Gender Wedges. We allow employers to have preferences over workers' gender captured by employer-specific gender wedges $\{\tau_g\}_g$. Two popular theories that map into this gender wedge include taste-based discrimination, as in Becker (1971), and firm-level comparative advantages in production, as in Goldin (1992). Without loss of generality, we normalize $\tau_g = \tau \mathbf{1}[g = F]$, where $\tau \in \mathbb{R}$ captures the rel-

[21] Appendix C.5 presents an alternative model, in which firms produce an amenity vector with a gender-specific vector of preference weights on amenities, and establishes conditions for *observational equivalence* and *counterfactual equivalence* between the two models.

ative disadvantage of women relative to men at the same employer, all else (i.e., estimates of worker ability, firm productivity, preference weights on amenities, and the firm's amenity cost shifters) equal. Also without loss of generality, we write gender wedges as an implicit tax on firm revenues, as in the literature on factor-input misallocation (e.g., Restuccia and Rogerson, 2008).[22]

Value Function. In summary, firms post wages, amenities, and vacancies in each market to maximize their steady-state flow payoff. The value $\Pi(\cdot)$ of a firm of type $(p, \{\tau_g\}_g, \{\beta_g\}_g, \{c_g^{a,0}\}_g)$ satisfies

$$\rho\Pi(\cdot) = \max_{\{w_{gz}, a_{gz}, v_{gz}\}_{gz}} \left\{ \sum_g \int_z \left[(1-\tau_g)pz - w_{gz} - c^a_{gz}(a_{gz}) \right] l_{gz}(x_{gz}, v_{gz}) - c^v_{gz}(v_{gz})\, dz \right\}$$

$$\text{s.t.} \quad w_{gz} + \beta_g a_{gz} = x_{gz}, \quad \forall (g,z), \tag{11}$$

where $l_{gz}(x_{gz}, v_{gz})$ is the equilibrium mapping from flow utility—i.e., the wages plus gender- and firm-specific amenity valuations—and vacancies into employment at a given firm.

4.3 Matching

The effective mass of job searchers and total mass of vacancies in market (g, z) are given by

$$U_{gz} = \mu_{gz} \left[u_{gz} + s_g^E (1-u_{gz}) + s_g^G \right], \quad V_{gz} = \int_j v_{gz}(j)\, d\Gamma(j). \tag{12}$$

A Cobb-Douglas matching function with constant returns to scale combines the mass of job searchers with the mass of vacancies to produce $m_{gz} = \chi_g V_{gz}^\alpha U_{gz}^{1-\alpha}$ matches between workers and firms, where $\chi_g > 0$ is the matching efficiency and $\alpha \in (0,1)$ is the matching elasticity. Labor market tightness is

$$\theta_{gz} \equiv \frac{V_{gz}}{U_{gz}}. \tag{13}$$

Workers' job-finding rates among the nonemployed, λ^U_{gz}, the voluntary job offer rates among the employed, λ^E_{gz}, the involuntary job offer rates, λ^G_{gz}, and firms' job-filling rates, q_{gz}, are given by

$$\lambda^U_{gz} = \chi_g \theta^\alpha_{gz}, \quad \lambda^E_{gz} = s^E_g \lambda^U_{gz}, \quad \lambda^G_{gz} = s^G_g \lambda^U_{gz}, \quad \text{and} \quad q_{gz} = \chi_g \theta^{\alpha-1}_{gz}. \tag{14}$$

[22] As will become clear soon, the exact functional form by which τ_g enters firms' payoff function (e.g., multiplicatively or additively) is immaterial for the model's identification and estimation.

4.4 Firm Size Distribution

The following Kolmogorov forward equation describes employment's law of motion given a firm's flow-utility and vacancy policies (x, v), the offer distribution $F_{gz}(x)$, and tightness θ_{gz} in market (g, z):

$$\dot{l}_{gz}(x,v) = \left[-\delta_g - \lambda_{gz}^G - \lambda_{gz}^E \left[1 - F_{gz}(x)\right]\right] l_{gz}(x,v) + \left[\frac{u_{gz} + (1-u_{gz})s_g^E G_{gz}(x) + s_g^G}{u_{gz} + (1-u_{gz})s_g^E + s_g^G}\right] v q_{gz}, \tag{15}$$

Solving for the stationary employment distribution in each market (g, z), firm sizes are given by

$$l_{gz}(x,v) = \left(\frac{1}{\delta_g + \lambda_{gz}^G + \lambda_{gz}^E\left[1 - F_{gz}(x)\right]}\right)^2 \frac{v}{V_{gz}} \mu_{gz}(u_{gz} + s_g^G)\lambda_{gz}^U(\delta_g + \lambda_{gz}^G + \lambda_{gz}^E). \tag{16}$$

4.5 Equilibrium Characterization

Appendix C.1 defines a *stationary equilibrium* of the economy. Our assumptions that markets are segmented and that technology is additively separable across worker types keep the analysis tractable by allowing us to divide the firm's problem into separate subproblems across markets. In particular, a firm with preference weight on amenities β_g offering wage w_{gz} and amenity level a_{gz} finds itself ranked according to its flow utility $x_{gz} = w_{gz} + \beta_g a_{gz}$ along a firm ladder in market (g, z). Next, we provide comparative statics results for the optimal policy functions of a firm of type $(p, \{\tau_g\}_g, \{\beta_g\}_g, \{c_g^{a,0}\}_g)$. To this end, we first define the *preference-adjusted amenity cost shifter* $\tilde{c}_g^{a,0} \equiv c_g^{a,0}/\beta_g$ as the ratio of the firm's amenity cost shifter to the marginal utility from amenities at that firm for workers of gender g.

Lemma 1 (Optimal Amenities). *A firm's optimal amenity policy function $a_{gz}^*(\cdot)$ can be written as $a_z^*(\tilde{c}_g^{a,0}) = (\tilde{c}_g^{a,0})^{1/(1-\eta^a)}z$. The optimal amenity cost function $c_{gz}^{a,*}(\cdot)$ can be written as $c_z^{a,*}(\tilde{c}_g^{a,0}) = (\tilde{c}_g^{a,0})^{1/(1-\eta^a)}z/\eta^a$.*

Proof. See Appendix C.2. □

Due to their convex per-worker cost, amenities are offered up to the point where the ratio of a firm's marginal cost of amenities to workers' marginal utility from amenities at that firm equals the same ratio for wages, which is unity. Lemma 1 states that the preference-adjusted amenity cost shifter, $\tilde{c}_g^{a,0}$, is sufficient to determine the optimal amenity level and the optimal amenity cost for a given worker ability z. Intuitively, a firm optimally produces higher amenity levels a^* either if it has a lower cost of providing the amenity, $c_g^{a,0}$, or if workers have a greater preference weight on amenities, β_g. An important implication of this is that we cannot distinguish between the optimal amenity level

a^*_{gz} and its preference weight β_g. We can only infer their product—i.e., the *amenity valuation* $\beta_g a^*_{gz}$.[23]

Next, we define a firm's *composite productivity* in market (g, z),

$$\tilde{p}_{gz} \equiv (1 - \tau_g) pz + \beta_g a^*_z(\tilde{c}^{a,0}_g) - c^{a,*}_z(\tilde{c}^{a,0}_g), \tag{17}$$

as revenues net of the gender wedge, $(1 - \tau_g)pz$, plus the optimized amenity valuation, $\beta_g a^*_z(\tilde{c}^{a,0}_g)$, minus optimized amenity costs $c^{a,*}_z(\tilde{c}^{a,0}_g)$. While amenities are endogenously produced in equilibrium, Lemma 1 allows us to treat the optimal amenity valuation $\beta_g a^*_{gz} = a^*_z(\tilde{c}^{a,0}_g)$ and optimal amenity cost $c^{a,*}_{gz} = c^{a,*}_z(\tilde{c}^{a,0}_g)$, and thus composite productivity $\tilde{p}_{gz}$, for workers of type (g, z) as an exogenous firm characteristic. Consequently, we can rewrite the firm's subproblem in market (g, z) as

$$\rho \Pi_{gz}(\tilde{p}_{gz}) = \max_{x,v} \left\{ \left[\tilde{p}_{gz} - x \right] l_{gz}(x, v) - c^v_{gz}(v) \right\}. \tag{18}$$

Equation (18) implies that composite productivity $\tilde{p}_{gz}$ is a sufficient statistic for a firm's type when solving for equilibrium in market (g, z). Thus, our model is isomorphic to one without amenities or gender wedges but with productivity p replaced by composite productivity $\tilde{p}$ and wages w replaced by flow utility x.[24] Firms with the same composite productivity $\tilde{p}$ offer the same flow utility from total compensation $x^*(\tilde{p})$ and post the same number of vacancies $v^*(\tilde{p})$ in equilibrium.

Lemma 2 (Optimal Vacancies)**.** *Keeping fixed all other parameters, a firm's optimal vacancy policy $v^*_{gz}(\cdot)$ is strictly increasing in composite productivity $\tilde{p}_{gz}$ and, thus, strictly increasing in productivity p, strictly decreasing in the gender wedge τ for women, and strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$.*

Proof. See Appendix C.3. □

The intuition behind Lemma 2 is that more profitable firms benefit more from each worker contact.

Lemma 3 (Optimal Flow Utility and Wages)**.** *Keeping fixed all other parameters, a firm's optimal flow utility offer $x^*_{gz}(\cdot)$ is strictly increasing in composite productivity $\tilde{p}_{gz}$ and, thus, strictly increasing in productivity p for all worker types, strictly decreasing in the gender wedge τ for women, and strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$. A firm's optimal wage offer $w^*_{gz}(\cdot)$ is strictly increasing in productivity p for all worker types and strictly decreasing in the gender wedge τ for women.*

[23] Equivalently, we could formulate a firm's type as consisting of a gender-specific wedge on the gender-neutral amenity cost shifter, which in our formulation above is absorbed in the preference weight on amenities.

[24] See Mortensen (2003) and Engbom and Moser (2022) for examples of such a model.

Proof. See Appendix C.4. □

Lemma 3 extends well-known comparative statics results for wages (e.g., Bontemps et al., 1999, 2000) to a richer environment with both wages and amenities. Intuitively, firms with a greater payoff from employment optimally offer higher flow utility in order to attract and retain more workers.

4.6 Equilibrium Properties

Our model has several notable equilibrium properties. First, search frictions give rise to gender-specific monopsony power (Robinson, 1933) across firms, which results in utility dispersion both within and across genders.[25] Search frictions imply that low- and high-utility firms coexist and relocation towards higher-utility firms is sluggish. As a result, gender differences both within and between firms depend on the severity of search frictions.

Second, observed wage differences are neither necessary nor sufficient for the existence of utility differences. Two firms can offer the same wage $w = w'$ but different utility $x \neq x'$ if $a \neq a'$. Conversely, two firms can offer different wages $w \neq w'$ but the same utility $x = x'$ if $a' = a + w - w'$. The competitive environment determines the extent to which amenities are priced into wages, giving rise to compensating differentials (Rosen, 1986). Consequently, gender pay gaps in the data may either understate or overstate inequality in utility, the measurement of which requires a structural model.

Third, the model rationalizes job-to-job transitions with wage declines through two channels. A worker may voluntarily transition from a job with wage-amenity combination (w, a) to one with (w', a') and $w' < w$ if $x' > x$. In addition, a worker may involuntarily transition from a job offering (w, a) to a randomly drawn job (w', a') with $w' < w$, regardless of $x' \lesseqgtr x$.

Fourth, firms have multiple levers to discriminate between genders—wages w, amenities a, and vacancies v. Firms choose these levers to maximize the payoff given their productivity p, preference-adjusted amenity cost shifters $\tilde{c}_g^{a,0}$, and gender wedges τ, as well as the general competitive environment. Thus, gender gaps in pay, amenities, and employment are all jointly determined.

Fifth, what is a high-paying, high-amenity, or large employer may differ across genders. While we naturally expect pay, amenities, and size to be correlated within an employer across genders, our model allows these characteristics to differ freely, so men and women climb separate firm ladders.

Sixth, even nondiscriminatory firms (i.e., those with $\tau = 0$ and $\tilde{c}_M^{a,0} = \tilde{c}_F^{a,0}$) may treat women differently than men due to the presence of other discriminatory firms (i.e., those with $\tau \neq 0$, as in

[25] By search frictions, here we refer to the combination of the vacancy cost shifter $c_g^{v,0} > 0$, the relative hazard of voluntary on-the-job offers $s_g^E < \infty$, the relative hazard of involuntary on-the-job offers $s_g^G > 0$, and the separation hazard $\delta > 0$.

Becker, 1971, or $\tilde{c}_M^{a,0} \neq \tilde{c}_F^{a,0}$). On-the-job search leads profit-maximizing firms to account for other employers' characteristics when making their own equilibrium decisions. In this sense, our model features "discrimination" spillovers arising from strategic links throughout the distribution of firms.[26]

4.7 Sources of Gender-Specific Sorting across Firms

We can group the reasons why men and women are sorted differently across firms with different amenities and wages in our model into four forces. The first force behind gendered sorting comprises *gender-specific firm heterogeneity in total compensation*. Akin to classical job-ladder models (Burdett and Mortensen, 1998; Bontemps et al., 1999, 2000), firms in equilibrium offer workers different utility levels x_{gz} as a strictly increasing function of composite productivity $\tilde{p}_{gz}$. Intuitively, men and women climb different job ladders and a firm's job-ladder rank determines its share of job offers accepted by workers of a given type (g, z).

The second force is *gender-specific firm heterogeneity in wage and amenity-valuation shares*. Conditional on a firm's utility offer x_{gz}, its shares of wages w and amenity valuations $\beta_g a$ in total compensation $x = w + \beta_g a$ vary as a function of the preference-adjusted amenity cost shifter $\tilde{c}_g^{a,0}$, while adjusting productivity net of the gender wedge, $(1 - \tau_g)p$, to hold fixed composite productivity $\tilde{p}$. Intuitively, compensating differentials imply that men and women may enjoy different total-compensation bundles of wages and amenity valuations across firms. Thus, women may have a relative preference—or induced demand due to household arrangements and other constraints—for certain workplace amenities and other employer characteristics, captured by β_g, consistent with the arguments in Bertrand et al. (2010), Goldin (2014, 2023), and Cubas et al. (2021, 2023).

The third force pertains to *gender-specific search frictions*. Due to the random nature of job offers, reaching higher-utility employers takes time. The strength of search frictions is summarized by the effective speed of climbing the job ladder, $\kappa_{gz}^E = \lambda_{gz}^E / (\delta_g + \lambda_{gz}^G)$, which differs across workers of gender g and ability z. Intuitively, men and women may be differently distributed across rungs of their respective job ladders as a result of differential search frictions (Bowlus, 1997).[27]

The fourth and final force stems from *gender-specific firm heterogeneity in vacancies*. A firm's number of vacancies, v_{gz}, relative to all other firms' vacancy postings, determines a worker's probability of

[26] Black (1995) studies a related phenomenon with degenerate wage distributions, while Flabbi (2010) considers spillovers through gender-specific values of unemployment without on-the-job search. Relatedly, Caldwell and Harmon (2019) and Caldwell and Danieli (2023) show that workers' outside employment opportunities affect current employment outcomes.

[27] While the job offer rate from nonemployment, λ_{gz}^U, matters for the nonemployment rate of workers of type (g, z), it does not affect their distribution across firms conditional on being employed.

receiving a job offer from that firm conditional on receiving any offer. Note that a firm's number of posted vacancies v_{gz} is strictly increasing in its composite productivity $\tilde{p}_{gz}$. Intuitively, for a given gender, the payoff from recruiting workers is larger for firms with greater composite productivity.

4.8 Discussion of Model Assumptions

We now discuss some of our more restrictive modeling assumptions and their implications—see Appendix C.6 for details. First, that output is linear within worker types is arguably not particularly restrictive since a firm's profit function is already concave due to convex vacancy costs.

Second, labor market segmentation allows firms to tailor wages, amenities, and vacancies to each market, which we view as a modeling device to match the empirical gender segregation across employers together with observed differences in pay and amenity utilization even within employers.[28]

Third, while our model allows for gender differences in labor market flow rates, we do not take into account workers' family status. This simplifying assumption is grounded in the empirical observation that women do not systematically sort into lower-paying firms around parental leaves in our data. Furthermore, women may be treated differently by employers even before childbirth in anticipation of future fertility events.

Fourth, we have paid special attention to gender differences while abstracting from labor demand or supply factors within genders. However, our framework is more general than its application to gender, as it can be applied to any set of population groups indexed by g, which could be either observed (e.g., parental status, education, race) or inferred in a pre-estimation step (e.g., k-means clustering on observables, as in Bonhomme et al., 2019, 2022). This makes our model a flexible tool for studying employer heterogeneity in pay, amenities, and employment across population groups.

Finally, the Brazilian labor market is characterized by many firm-level distortions (Muendler, 2004; Vasconcelos, 2017; Dix-Carneiro et al., 2024) that may be more pronounced than in other contexts such as the U.S. Therefore, what we refer to as "productivity" (p) should really be interpreted as productivity net of distortions or implicit taxes, in the spirit of Restuccia and Rogerson (2008) and Hsieh and Klenow (2009). Our notion of misallocation captures the output and welfare losses from workers across firms due to search frictions, all relative to an economy with the same distribution of

[28] For robustness, we have solved a number of alternative models. 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. A model in which firms produce an amenity vector with gender-specific utility weights is discussed in Appendix C.5. Finally, a model in which vacancies are gender-neutral, as in Appendix C.7, fails to account for the empirical dispersion of female employment shares and, hence, the between-employer pay gap in the data.

firm-level productivities and distortions but without labor market frictions.

5 Identification

To operationalize our model, we provide a constructive proof of identification of all model parameters based on linked employer-employee data. From a bird's eye view, we leverage the intuition that firms' unobserved surplus, consisting of productivity and amenity values net of wage and amenity costs, can be inferred from the employer size distribution. To anticipate our findings, our model rationalizes the observation of a highly skewed employment distribution plus substantial pay dispersion conditional on firm surplus through sizable compensating differentials due to employer amenities.

Relative to existing results on the nonparametric identification of job ladder models—notably Bontemps et al. (1999, 2000)—we achieve identification in an environment with the additional model ingredients of worker heterogeneity in gender and ability (as in Card et al., 2016), endogenous vacancy posting (as in Mortensen, 2003), and endogenous amenities (as in Hwang et al., 1998).

Our identification argument, which is summarized in Figure D.1 in Appendix D.1, takes as given three exogenous parameters and proceeds recursively in five steps.

5.1 Exogenous Parameters

There are three exogenously set parameters in our framework. First, the discount rate $\rho = 0.052$ corresponds to an annual compound real interest rate of 5.3 percent. The choice of this parameter value is innocuous as it affects only our computation of the flow value of nonemployment. Second, we impose a common normalization of the matching efficiency $\chi_g = 1.000$ for both genders g. This is without loss of generality, as our results below really concern the vacancy cost shifter $c_g^{v,0}$ relative to the matching efficiency χ_g, which is all that matters for our purposes.[29] Third, the elasticity of the matching function $\alpha = 0.500$ is an agreed-upon value in the literature—see, for example, Petrongolo and Pissarides (2001), Hall and Milgrom (2008), and Engbom and Moser (2022).

5.2 Step 1: Gender-Specific Firm Pay

In the first step, we demonstrate that our equilibrium model provides a microfoundation for the decomposition of log wages into worker FEs and gender-specific employer FEs in Card et al.'s (2016) variant of the original framework due to AKM, on which our analysis in Section 3.1 builds.

[29] To separately identify the match efficiency χ_g would require observing the total number of vacancies in the economy.

Proposition 1 (Equilibrium Wage Equation)**.** *The equilibrium wage of a worker of gender g and ability z at a firm with composite productivity $\tilde{p}_g$, preference weight β_g and preference-adjusted amenity cost shifter $\tilde{c}_g^{a,0}$ is*

$$\ln w_{gz}\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right) = \underbrace{\alpha_z}_{\text{"worker wage FE"}} + \underbrace{\psi_g^w\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right)}_{\text{"gender-firm wage FE"}}, \tag{19}$$

where

$$\alpha_z = \ln z, \tag{20}$$

$$\psi_g^w\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right) = \ln\left(\tilde{p}_g - \beta_g a_g^*\left(\tilde{c}_g^{a,0}\right) - \int_{\tilde{p}' \geq \phi_g}^{\tilde{p}_g}\left[\frac{1+\kappa_g^E\left[1-F_g\left(x_g^*\left(\tilde{p}_g\right)\right)\right]}{1+\kappa_g^E\left[1-F_g\left(x_g^*\left(\tilde{p}'\right)\right)\right]}\right]^2 d\tilde{p}'\right). \tag{21}$$

Proof. See Appendix D.2. □

Proposition 1 shows that equilibrium wages in the model are log-additive between a worker component and a gender-specific firm component, as in Card et al.'s (2016) variant of the framework originally due to AKM. The worker wage FE α_z is a strictly monotonic transformation of worker ability z. The gender-firm wage FE $\psi_g^w(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0})$ depends only on gender-firm-specific parameters—namely, a firm's composite productivity $\tilde{p}_g$, its preference weight β_g and its preference-adjusted amenity cost shifter $\tilde{c}_g^{a,0}$.

In Card et al. (2016), the interpretation of gender-firm FEs was one of gender-specific rent sharing. In contrast, our interpretation allows for both gender-specific rent sharing and compensating differentials. Gender-specific rent sharing is captured by the function $\psi_g^w(\cdot)$, which depends on gender-specific monopsony power. Compensating differentials shape equilibrium wage offers through two channels. First, directly, by substituting for a given firm's wage payments—see the term $\tilde{p}_g - \beta_g a_g^*$ in equation (21). Second, indirectly, by shaping the degree of competition between firms—see the integral term involving $x_g^*(\tilde{p}_g)$ in equation (21). Therefore, our equilibrium model lends itself to reinterpreting the wage equation with gender-specific employer pay components, as in Card et al. (2016).

Proposition 6 in Appendix D.4 shows how to impose a model-consistent normalization of firm pay across genders, used in our empirical analysis in Section 3.1, by extending the arguments in Card et al. (2016) to our environment with gender-specific amenities and compensating differentials.

In Appendix D.3, we demonstrate that equilibrium amenities also have a log-additive structure, $\ln \beta_g a_{gz}(\tilde{c}_g^{a,0}) = \alpha_z + \psi^a(\beta_g, \tilde{c}_g^{a,0})$, where $\alpha_z = \ln z$ is a worker amenity FE and $\psi^a(\beta_g, \tilde{c}_g^{a,0}) = \ln \beta_g +$

$\ln(\tilde{c}_g^{a,0})/(1-\eta^a)$ is a gender-firm amenity FE. While Card et al. (2016) studied a model of wages akin to equation (19), a formal treatment of amenities and compensating differentials was missing from their analysis. Our framework fills this gap by explicitly modeling firms' equilibrium wage and amenity choices.

In summary, our model allows us to separate worker from firm components of pay and amenities. As a result, we can abstract from heterogeneity in worker ability when discussing the sources of firm heterogeneity. This rationalizes analogous reduced-form assumptions implicitly made in environments that pool workers for the estimation of firm pay and amenities—see, for instance, Sorkin (2017, 2018). For the remainder of the analysis, we focus on gender-firm components of pay and amenities. This allows us to pool workers in the data, drop ability z from all subscripts in the model, and henceforth treat the gender-firm pay component $w_g \equiv \psi_g^w(\cdot)$ in equation (19) as known.[30]

5.3 Step 2: Employer Ranks

In the second step, we estimate employer rankings by gender based on a model-consistent revealed-preference argument. For the remainder of this section, we drop the gender subscript g when it is dispensable and refer to firms by their (gender-specific) rank r. For example, we write $w(r)$ for the firm component of wages received by workers of gender g at a firm of rank r.

Proposition 2 (Employer Ranks)**.** *All workers of the same gender share a common employer ranking $r \in [0,1]$, which is identified given employer sizes $l(r)$.*

Proof. See Appendix D.5. □

Proposition 2 allows us to estimate gender-specific revealed-preference employer ranks.[31] Our rank notion coincides with that of our model in Section 4, in which workers are less likely to endogenously separate from and more likely to accept offers at higher-utility employers, so that higher-ranked employers are larger.[32] Of course, there are many other ways to estimate employer ranks in a model-consistent way. One alternative is to exploit the pattern of worker flows between firms, for instance using poaching ranks (Bagger and Lentz, 2019) or a variant of PageRanks (Page et al., 1998; Sorkin, 2018). Appendix E.8 proves that our model is consistent with these alternative firm rank

[30]Specifically, this allows us to control for gender-specific selection based on ability (Mulligan and Rubinstein, 2008).

[31]In Appendix D.11, we demonstrate, using a combination of theory and Monte-Carlo simulations, the robustness of our estimation results to model misspecification and subgroup heterogeneity.

[32]That firms reveal their value (here, match surplus) through labor demand (here, vacancies) is a feature our model shares with a new generation of neoclassical frameworks (Lamadon et al., 2022; Berger et al., 2022; Felix, 2022; Sharma, 2023).

measures. Appendix E.9 studies in detail the empirical correlation structure between alternative employer rank measures (Table E.2), how alternative employer rank measures compare to our baseline rank measure equal to firm size across sectors (Figure E.7), and how our estimation and decomposition results change if we take the model-consistent variant of PageRanks as an alternative measure of employer ranks (Tables E.3 and E.4).

5.4 Step 3: Labor Market Objects

In the third step, we estimate labor market objects by combining employer ranks from above with information on worker flows between employers as well as between employment and nonemployment. To this end, we exploit the existence of a job ladder across firm utility ranks in our model.

Proposition 3 (Labor Market Objects)**.** *Gender-firm-specific recruiting intensities $f(r)$ and vacancies $v(r)$ as well as gender-specific separation hazards δ, job offer hazards from nonemployment λ^U, involuntary job offer hazards λ^G, voluntary on-the-job offer hazards λ^E, and aggregate vacancies V are identified given employer ranks and data on worker flows between employment states.*

Proof. See Appendix D.6. □

Proposition 3 states that ordinal employer ranks—rather than cardinal utility levels—are sufficient to identify key labor market objects in our model. Given that we have already estimated model-consistent revealed-preference ranks for each gender across firms, the flow pattern of workers between firms as well as between employment and nonemployment pins down the stated objects.

To gain some intuition, consider the identification of the involuntary job offer hazard, λ^G. With data on firm pay alone, this hazard is impossible to identify in our framework. The reason is that wage cuts between employers may be due to any one or a combination of two reasons. First, a worker may voluntarily switch to a higher-utility firm that offers higher utility but a disproportionately higher share of it is delivered in the form of amenities, leading to lower wages due to compensating differentials. Second, a worker may switch to a a firm that offers lower utility, pay, and amenities due to an involuntary job offer. In contrast, given gender-specific firm ranks already estimated by virtue of Proposition 2 in the previous step, we can infer the involuntary job offer hazard by simply counting the incidence of transitions up versus down the firm ranks at different rungs of the job ladder.

As a result of Proposition 3, going forward we can treat gender-firm-specific recruiting intensities, $f(r)$, and vacancies, $v(r)$, as well as gender-specific separation hazards, δ, job offer hazards from nonemployment, λ^U, involuntary job offer hazards, λ^G, voluntary on-the-job offer hazards, λ^E, and aggregate vacancies, V as known.

5.5 Step 4: Firm-Level Parameters

In the fourth step, we identify gender-specific parameters for each firm. In doing so, we treat as known some economy-wide parameters, the identification of which we discuss next, in Section 5.6.[33]

We now state an important identification result.

Proposition 4 (Firm-Level Parameters)**.** *The following gender-firm-specific parameters as functions of r are point identified given employer ranks and labor market objects: productivity $p(r)$, the gender wedge $\tau(r)$, the amenity valuation $\beta(r)a^*(r)$, and amenity costs $c^a(a^*(r))$.*

Proof. See Appendix D.7. □

Our identification argument leverages the insight that unobserved firm flow profits per worker can be inferred from equilibrium recruiting intensities or, equivalently, from firm sizes. Intuitively, if a firm makes higher profits per matched worker, it will optimally post more vacancies and attain a larger size. In turn, by leveraging the equilibrium structure of our model, the distribution of profits per matched worker across firm ranks allows us to infer utility levels offered to workers up to a constant.[34] In combination with the observed wage components, these yield firm-specific amenity values, which—given our assumed amenity cost function—is equivalent to identifying amenity cost parameters. Finally, we identify firm productivities and gender wedges by combining observed wages and identified amenity costs with firms' profits per matched worker.[35] For this, recall that the exact functional form by which τ_g enters firms' payoff function (e.g., multiplicatively or additively) is immaterial for identification and estimation.

It is worth noting that our result is more general than the exact statement in Proposition 4, in the sense that our assumptions on the functional form of the amenity cost in equation (8) and of the vacancy cost in equation (9) can be significantly relaxed. All that is required is that the amenity and

[33] These are the vacancy cost intercept $c^{v,0}$, the elasticity of vacancy costs η^v, and the elasticity of amenity costs η^a.

[34] In particular, utility levels offered across firms can be shifted up or down by a constant without affecting the observed behavior of firms or workers in our model. We choose a normalization of utility levels to achieve a minimum amenity valuation that is positive but arbitrarily close to zero at some firm, which minimizes the share of total utility due to amenity valuations.

[35] In the data, separation rates are decreasing in firm size: this can be viewed as a "job-stability amenity" that our estimation can attribute to larger firms offering higher utility. Thus, after solving a version of the model that allows for nonincreasing separation rates across firm ranks in Appendix C.9, Appendix D.12 details how our identification argument changes when separation rates are heterogeneous across firms. We thank an anonymous referee for this suggestion.

vacancy cost functions are increasing and convex, as is commonly assumed in many applications.[36] Furthermore, our choice of the amenity cost function has no bearing on the estimated distribution of firm-level amenities—in particular, our choice of the elasticity of the amenity cost function η^a is irrelevant for the estimated distribution of amenities $\beta(r)a(r)$.[37] In comparison to related work by Sorkin (2017, 2018), we leverage the equilibrium nature of our model—specifically, firms' endogenous vacancy posting subject to a convex increasing cost function—to achieve identification of *levels* of utility across firms, which in the context of Sorkin (2017, 2018) are arbitrarily normalized due to the scale parameter of random utility shocks being not identified in his partial-equilibrium framework.[38] Our approach to identifying firm-level amenities relies on a choice of an integration constant K—see equation (D.45) in Appendix D.7. We argue that choosing a low enough value of the constant K to guarantee that the minimum level of amenity valuations across firms is close to zero is innocuous. Specifically, choosing a larger value of the constant K would lead us to infer an even greater share of amenities in total compensation and for the gender gap in log total compensation to be even smaller.

5.6 Step 5: Economy-Wide Parameters

In the final step, we pin down three remaining economy-wide parameters given aggregate statistics.

Proposition 5 (Economy-Wide Parameters)**.** *(i) The vacancy cost shifter $c^{v,0}$ is identified based on the aggregate labor share; (ii) the elasticity of the vacancy cost function η^v is identified based on the firm pay-profit gradient; (iii) the elasticity of the amenity cost function η^a is identified based on the aggregate amenity cost share in the data.*

Proof. See Appendix D.9. □

The intuition for part (i) relies on the fact that vacancy costs introduce concavity with respect to the vacancy choice into the firm's objective, which uniquely pins down a firm's size and profits. Firm optimality with respect to its vacancy choice implies that the same vacancy posting behavior can be

[36] Examples of such amenity cost functions are Hwang et al. (1998), Lang and Majumdar (2004), and Lavetti and Schmutte (2018). The model in Sorkin (2018) is isomorphic to one with convex increasing amenity costs. Examples of such vacancy cost functions are Mortensen (2003), Kaas and Kircher (2015), Lise and Robin (2017), Bilal et al. (2022), Engbom and Moser (2022), Bilal and Lhuillier (2023), Bloesch and Larsen (2023), Heise and Porzio (2023), and Lindenlaub et al. (2023).

[37] For instance, we would recover the identical distribution of amenities $\beta(r)a(r)$ under the assumption of exogenous amenities. However, the estimated value of the elasticity of the amenity cost function η^a will matter for the inferred distribution of firm productivity $p(r)$ and thus composite productivity $\tilde{p}(r)$.

[38] In Sorkin (2017, 2018), the *levels* of utility across discrete choices of firms could be pinned down if the elasticity of labor supply with respect to the wage were known. Absent a credibly identified change in firm pay, holding fixed firm amenity values and all other features of the environment, our model provides an alternative approach to pinning down the scale of firm utilities using an equilibrium model and functional forms commonly used in the macro-labor literature, the cost parameters of which we identify and estimate based on linked employer-employee data and aggregate statistics.

rationalized by any combination of profits and the vacancy cost shifter that keeps the ratio between the two constant. However, as we scale up profits, we mechanically lower the labor share because the data pin down the level of wages. Therefore, we can estimate $c^{v,0}$ by finding the scale of profits that matches the empirical labor share.

The intuition for part (ii) is that as the elasticity of the vacancy posting cost becomes higher, the number of vacancies becomes more similar across two firms with given levels of profits per worker. In the limit as $\eta^v \to \infty$, all firms post a constant number of vacancies. Therefore, to rationalize the observed dispersion in vacancy posting, the model needs to generate more dispersion in profits per matched worker, yielding greater profit dispersion. For fixed levels of pay observed in the data, this implies that the pay-profit gradient varies directly with the elasticity of the vacancy cost function η^v.

Part (iii) follows from our assumptions on the amenity cost function in equation (8). The cost of creating the optimal amount of amenities is $c^a(a^*) = \beta a^* / \eta^a$, which is inversely proportional to η^a. Therefore, we can identify η^a by using it to match the aggregate amenity cost share.

5.7 Interpretation of Results

To summarize, we have identified gender-firm-specific parameters $(p(r), \tau(r), c^{a,0})$, gender-specific labor market objects $(\delta, \lambda^U, \lambda^E, \lambda^G)$, and economy-wide parameters $(c^{v,0}, \eta^v, \eta^a)$. Combining the model structure with the data, this yielded firm-level estimates of $w(r), a(r), v(r)$ across firm ranks r.

At this point, it may be helpful to take a step back and ask: What can we (not) learn from these identification results? Recall the utility of a worker of gender g and ability z is $x = w + \beta a$, which consists of the sum of the monetary wage w and the nonwage amenity valuation βa. Thus, for each gender, our estimates of amenity valuations βa and total compensation x are *in units of wages*, w. In other words, we have identified the relative importance of wages versus amenity valuations relative to the sum of the two.

Nevertheless, since wages and amenities are in the same units, our identification results allow us to quantify the relative importance of utility x and amenity valuations βa in wages w across firms r as

$$\underbrace{\frac{w_{gz}(r)}{x_{gz}(r)}}_{\text{wage share in total compensation}} + \underbrace{\frac{\beta_g a_{gz}(r)}{x_{gz}(r)}}_{\text{amenity share in total compensation}} = 1. \tag{22}$$

They also allow us to quantify the shares of variation in utility $x_{gz}(r)$ across firms r due to variation in wages, amenities, and their covariance:

$$\underbrace{\frac{Var\left(x_{gz}(r)\right)}{Var(w_{gz}(r))}}_{\text{variance share of wages due to } x_{gz}(r)} + \underbrace{\frac{Var\left(\beta_g a_{gz}(r)\right)}{Var(w_{gz}(r))}}_{\text{variance share of wages due to } \beta_g a_{gz}(r)} - \underbrace{\frac{2Cov\left(x_{gz}(r), \beta_g a_{gz}(r)\right)}{Var(w_{gz}(r))}}_{\text{variance share of wages due to covariance}} = 1. \tag{23}$$

That all the statistics in equations (22) and (23) are identified in our equilibrium framework is an important contribution over existing work on compensating differentials. For example, the partial-equilibrium model by Sorkin (2018) *"can identify the variation in amenities that comes from the 'pure' Rosen motive"* (p. 1373) but, without further assumptions, *"cannot identify the variation in amenities that contribute to utility dispersion, that is, those that come from the Mortensen motive"* (p. 1374). By making additional assumptions on the cost structure of firms' hiring decisions and other features of the environment, our equilibrium model sheds new light on compensating differentials across firms by recovering the joint distribution of amenities and wages, which allows us to quantify both the *"Rosen motive"* and the *"Mortensen motive,"* in the language of Sorkin (2018).

In addition to decomposing total compensation into wages and amenity valuations, our model also allows us to identify the underlying sources of (gender differences in) wages and amenities by associating with each wage-amenity bundle $(w, \beta_g a)$ a firm type $(p, (\beta_g)_g, \tau, (\tilde{c}_g^{a,0})_g)$. In classical job ladder models à la Burdett and Mortensen (1998), productivity p is the only source of firm heterogeneity and wages are the only form of compensation, so more productive firms offer higher wages and employ more workers in equilibrium. The identification results within our richer framework allow us to empirically test these relationships in a world with endogenous amenity provision and hiring.

5.8 Implementation

To implement our identification results, we proceed sequentially by gender. We first trim the top and bottom one percent of the gender-specific pay FE distributions to eliminate outliers. Then, we start by studying men, for whom gender wedges are normalized to zero. We estimate men's firm pay $w_M(r_M)$ using Proposition 1 and men's firm ranks r_M using Proposition 2. Given the distribution of firm ranks r_M, Proposition 3 allows us to estimate all men's labor market objects. We regularize the estimated hiring intensities with a kernel smoother to limit the extent of measurement error and then use Proposition 4 to identify the distribution of productivity $p(r_M)$ and amenity cost shifters $c_M^{a,0}(r_M)$

for men.[39] We then iterate over this process until we find the economy-wide parameters that allow us to match a set of aggregate moments detailed in Proposition 5.

Second, we study women, for whom gender wedges represent the implicit tax relative to the productivity level estimated for men. In other words, gender wedges are the only remaining reason for unequal treatment of men and women at the same employer, all else (i.e., estimates of worker ability, firm productivity, preference weights, and the firm's gender-specific amenity costs) equal. The identification proceeds analogously to that for men except that we take the values of firm productivity $p(r)$ as given, and economy-wide parameters are already estimated based on men. As part of this procedure, we recover firms' gender wedges for women such that women's productivity net of the gender wedge at a firm with rank r_F is $(1 - \tau(r_F))p(r_F)$.

A final remark is in order. While our model features a large number of parameters, our identification results demonstrate that these parameters are point-identified and thus pinned down empirically. In this sense, despite its richness, our model has no remaining degrees of freedom.[40]

6 Estimation Results

In Section 3, we estimated gender-specific firm pay components. Now, we present estimates of gender-specific firm ranks, labor market parameters, firm types, and economy-wide parameters based on the identification proof in Section 5. A novel aspect of our approach is that we impose no parametric restrictions on the distribution of gender-specific firm types, which the following analysis exploits.

6.1 Estimates of Employer Ranks

While the average firm in our sample is small, the average worker is employed at a large firm. Appendix Figure E.1 shows that employment is highly skewed toward large, high-ranked employers for both genders: the employment-weighted mean rank for men is 0.846 and that for women is 0.845.

6.2 Estimates of Labor Market Parameters

Our estimates of labor market parameters are shown in Table 4. The gender-specific population shares $\mu_g \equiv \int_z \mu_{gz}\, dz$ that we read off the data are $\mu_M = 59.9\%$ for men and $\mu_F = 40.1\%$ for women. Women receive fewer job offers from nonemployment ($\lambda_F^U = 9.1\%$ compared to $\lambda_M^U = 10.4\%$) and have a lower

[39] To translate our results from continuous firm types in theory to discrete firm types in the data, see Appendix D.8.

[40] Appendix D.10 demonstrates in a sequence of Monte Carlo simulations that our identification procedure perfectly recovers the model parameters under a range of parameterizations.

job destruction rate ($\delta_F^U = 2.8\%$ compared to $\delta_M^U = 3.6\%$). For both men and women, involuntary job offers ($\lambda_M^G = 1.1\%$ and $\lambda_F^G = 0.8\%$) are about as frequent as voluntary ones ($\lambda_M^E = 0.9\%$ and $\lambda_F^E = 0.7\%$). Overall, men receive a greater total of job offers than women. The prevalence of involuntary job offers indicates substantial undirectedness of job search across utility ranks, not just pay ranks (Jolivet et al., 2006).[41] The implied nonemployment rates ($u_M = 23.6\%$ and $u_F = 21.9\%$) reflect the presence of a large informal sector for both men and women in Brazil.[42] The flow value of nonemployment for men ($b_M = 2.281$) is slightly higher than that for women ($b_F = 2.234$), as is the implied reservation utility ($\phi_M = 2.353$ compared to $\phi_F = 2.274$).[43] It is important to keep in mind that these estimates reflect widespread unregistered employment in Brazil (Meghir et al., 2015). Compared to a typical high-income country's labor market (Taber and Vejlin, 2020), these estimates suggest substantial labor market imperfections in Brazil.

Table 4. Job offer arrival rates, job destruction rates, and flow values of nonemployment

Parameter	Description	Men	Women
μ_g	Population shares	0.599	0.401
λ_g^U	Offer arrival rate from nonemployment	0.104	0.091
δ_g	Job destruction rate	0.035	0.028
s_g^E	Relative arrival rate of voluntary on-the-job offers	0.090	0.075
s_g^G	Relative arrival rate of involuntary on-the-job offers	0.101	0.081
b_g	Flow value of nonemployment	2.282	2.223

Note: This table shows the estimated values of all labor market parameters—specifically, the offer arrival rate from nonemployment λ_g^U, the job destruction rate δ_g, the relative arrival rate of voluntary on-the-job offers s_g^E, the relative arrival rate of involuntary job offers s_g^G, and the flow value of nonemployment b_g—separately by gender $g \in \{M, F\}$. All rates are monthly. *Source:* Model estimates based on RAIS, 2007–2014.

6.3 Estimates of Firm Types

Productivity. Our estimates of firm productivity p in Figure E.3 in Appendix E.6 display substantial dispersion and a long right tail. The employment-weighted mean log productivity is 0.864 for men and 0.781 for women, implying a gender productivity gap of 8.3 log points. The standard deviation of log productivity is 0.573 for men and 0.601 for women, more than double that of firm pay. To test whether our model estimates of firm productivity p captures real-world firm productivity, we

[41] Figure E.2 in Appendix E.5 shows large dispersion in estimated firm-level recruiting intensities, and more so for women.

[42] While women are more likely to be informally employed (Engbom et al., 2022) and out of the labor force (World Bank, 2021b), these estimates suggest that both sexes are similarly attached to Brazil's formal sector conditional on participating. For the U.S., Albanesi and Şahin (2018) also find that men's unemployment rate has exceeded women's in the recent past.

[43] Le Barbanchon et al. (2020) also find that unemployed men have a higher reservation wage than women in France.

estimate

$$m_j = \eta \widehat{p}_j + \iota_j, \tag{24}$$

where m_j is our empirical measure of firm productivity, which we take to be log revenues per worker from Bureau van Dijk's Orbis Historical data for a subset of firms in RAIS, $\widehat{p}_j$ is our model's estimated firm productivity, and ι_j is an error term. Note that we put the empirical measure of firm productivity, m_j, on the left-hand side—rather than the right-hand side—of the regression to minimize the influence of measurement error and transitory fluctuations in the data.

Table 5 shows the results from this exercise. We find an estimated coefficient $\hat{\eta}$ of 0.897. This estimate suggests that we cannot reject, at conventional levels, the null hypothesis that $\eta = 1$.[44]

Table 5. Regressing empirical productivity measures on model estimates of firm productivity

	Coefficient	(std. err.)
Revenue per worker	0.897***	(0.155)
R^2	0.120	

Note: This table reports the estimated coefficient from regressing empirical revenue per worker on structurally estimated firm productivity p—see equation (24). Details of the variable construction are presented in Appendix E.2. Standard errors are clustered at the employer level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. *Source:* Model estimates based on RAIS, 2007–2014.

Gender Wedges. In Section 5.5, we discussed the identification of gender wedges τ separately from other (gender-specific) firm-level parameters. Here, we briefly exhibit some properties of the estimates of gender wedges τ. The distribution of estimated gender wedges τ is shown in Appendix Figure E.4. Women are less likely to work at firms with high gender wedges, with an employment-weighted mean of 0.059 for women compared to 0.235 for men. Gender wedges in our model may stand in for taste-based discrimination (Becker, 1971) or employer-level comparative advantages across genders (Goldin, 1992). To understand what they capture in the real world, we estimate

$$\ln(1 - \widehat{\tau}_j) = Z_j \eta + \iota_j, \tag{25}$$

where $\widehat{\tau}_j$ is the estimated gender wedge for employer j, Z_j is a vector of employer covariates, and ι_j is an error term. We include as covariates in Z_j a total of six variables constructed in the RAIS data.

[44] We find similar conclusions for other empirical measures of firm productivity, including profits per worker, net income per worker, and cash flow per worker.

Table 6 shows that our estimated gender wedges significantly load onto factors related to the female-friendliness of a workplace. For example, having a female manager is associated with lower gender wedges. Employers with longer working hours, major financial stakeholders, and larger workforces have higher gender wedges.[45] Altogether, we explain 63.2% of the variation in estimated gender wedges across firms. At the same time, it is worth noting that location- and industry-specific characteristics explain a substantial share of our model estimates of gender wedges. This suggests that gender wedges in our model capture important real-world employer differences.

Table 6. Regressing estimates of transformed gender wedges on employer characteristics

	Coefficient	(std. err.)
Female manager	0.006***	(0.002)
Nonroutine manual task intensity	−0.001	(0.007)
Nonroutine interpersonal task intensity	−0.002	(0.006)
Mean working hours	−0.010***	(0.004)
No major financial stakeholders	−0.010***	(0.002)
Log size	−0.155***	(0.007)
R^2	0.632	
Within-R^2	0.089	

Note: This table reports estimated coefficients from regressing log transformed gender wedges, $\ln(1-\tau)$, on observable employer characteristics—see equation (25). Estimates are conditional on municipality and sector FEs. All covariates are standardized and expressed in z-scores relative to the population of all firms for each gender. Details of all covariates are presented in Appendix E.3. Standard errors are clustered at the employer level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. *Source:* Model estimates based on RAIS, 2007–2014.

Amenity Cost Shifters. In Section 5.5, we discussed the identification of amenity cost shifters $c_g^{a,0}$ separately from other (gender-specific) firm-level parameters. The firm-level estimates of $c_g^{a,0}$ are dispersed with long right tails, see Appendix Figure E.5. We relate the implied amenity values to real-world amenity proxies as

$$\log \widehat{a}_{gj} = Z_{gj}\eta_g + \iota_{gj}, \tag{26}$$

where $\widehat{a}_{gj}$ is the estimated amenity value for gender g at employer j, Z_{gj} is a vector of gender-specific employer covariates, and ι_{gj} is an error. We include in Z_{gj} eight variables based on the RAIS data.

Table 7 shows that, for both men and women, employers with more generous parental leave policies, more stable income and employment, and larger workforces are associated with higher amenity values. For women, but not significantly for men, greater working hours flexibility is valued as a posi-

[45]We conservatively cluster standard errors at the employer level to address *experimental design issues* (Abadie et al., 2022).

tive amenity.[46] Altogether, we explain 70.4% of the variation in estimated amenities for men and 44.0% of that for women. As before, it is worth noting that location- and industry-specific characteristics explain a substantial share of our model estimates of amenities. Again, this suggests that the model amenities capture empirically relevant differences in employer characteristics. In order to reliably estimate the coefficients of specific employer characteristics, including location- and industry-specific fixed effects is important. Otherwise, the estimated coefficients may be biased upwards (downwards) if unobservable amenities are positively (negatively) correlated with our explanatory variables.

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

	Men		Women	
	Coefficient	(std. err.)	Coefficient	(std. err.)
Part-time work incidence	−0.006	(0.012)	0.010	(0.007)
Working hours flexibility	0.008	(0.013)	0.020***	(0.006)
Parental leave generosity	0.093***	(0.024)	0.023***	(0.007)
Income fluctuations	−0.034	(0.032)	−0.002	(0.007)
Workplace hazards	0.016	(0.015)	−0.002	(0.005)
Incidence of unjust firings	−0.028**	(0.014)	−0.020**	(0.009)
Incidence of workplace deaths	−0.034***	(0.011)	−0.047***	(0.010)
Log size	0.201***	(0.018)	0.139***	(0.021)
R^2	0.704		0.440	
Within-R^2	0.238		0.090	

Note: This table reports estimated coefficients from regressing log amenity valuations, $\ln(\beta_g a_g)$, on observable gender-specific employer characteristics—see equation (26). Estimates are conditional on municipality and sector FEs. All covariates are standardized and expressed in z-scores relative to the population of all firms for each gender. Details of all covariates are presented in Appendix E.4. Standard errors are clustered at the employer level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. *Source:* Model estimates based on RAIS, 2007–2014.

Correlation Structure. Appendix Table E.1 correlates gender-specific pay w_g, amenities a_g, productivity net of gender wedges $(1-\tau_g)p$, employment l_g, and employer ranks r_g. At this stage, the most interesting takeaway is that pay (0.909), ranks (0.576), and amenities (0.884) are positively but imperfectly correlated within employers across genders. For both genders, amenities are negatively related to pay, which suggests the presence of compensating differentials.

6.4 Estimates of Economy-Wide Parameters

Table 8 shows our estimated elasticity of the vacancy cost function, $\eta^v = 2.062$, which is in the range of existing estimates for Brazil (Engbom and Moser, 2022). Our estimated elasticity of the amenity

[46] Again, we conservatively cluster standard errors at the employer level.

cost function, $\eta^a = 5.726$, suggests that employer amenities are provided relatively inelastically.[47]

Table 8. Estimates of economy-wide parameters

Elasticity	Cost function	Value	Moment	Data	Model
η^v	Vacancies	2.063	Slope of log pay on log value added	0.179	0.179
η^a	Amenities	5.728	Cost share of amenities	0.080	0.080

Note: This table reports the estimated values of the elasticity of the vacancy cost function η^v and the elasticity of the amenity cost function η^a. Targeted moments are the elasticity of pay with respect to value added per worker from Alvarez et al. (2018) and the cost share of amenities in value added based on Bieri et al. (2023). See Appendix E.1 for details on the construction of the aggregate statistics. *Source:* Model estimates based on RAIS, 2007–2014.

6.5 Model Fit

Table 9 shows the model fit based on a range of moments relating to employer pay and worker transitions. Overall, the model fits the data well. The model understates the gender pay gap by 0.7 log points but broadly matches the empirical variances of gender-specific pay and the gender pay gap, job-to-job transition rates, the share of transitions with a pay cut, and the correlation between men's and women's pay within firms. Notice that we fit gender-specific firm pay and ranks by construction but discrepancies arise from the model not perfectly replicating the smoothed employment distribution.

Table 9. Model fit

Moment	Description	Data	Model
$\mathbb{E}[\psi_M - \psi_F]$	Gender log pay gap	0.115	0.110
$\mathbb{E}[\psi_F \mid g = M] - \mathbb{E}[\psi_F \mid g = F)]$	Gender log pay gap between employers	0.089	0.082
$\mathbb{E}[\psi_F - \psi_M \mid g = F]$	Gender log pay gap within employers	0.026	0.028
$Var[\psi_M]$	Variance of men's pay	0.054	0.053
$Var[\psi_F]$	Variance of women's pay	0.044	0.044
$Var[\psi_M - \psi_F]$	Variance of gender pay gap	0.009	0.010
$\mathbb{E}[\lambda_M^E(1 - F_M(x)) + \lambda_M^G]$	Job to job transition rate for men	0.013	0.015
$\mathbb{E}[\lambda_F^E(1 - F_F(x)) + \lambda_F^G]$	Job to job transition rate for women	0.010	0.011
$\mathbb{P}[\psi_M' < \psi_M]$	Wage decline probability after job to job for men	0.416	0.479
$\mathbb{P}[\psi_F' < \psi_F]$	Wage decline probability after job to job for women	0.430	0.498
$Corr(\psi_M, \psi_F)$	Correlation between men's and women's pay	0.921	0.956

Note: This table reports the fit of the model in terms of data-based and model-based moments across genders $g \in \{M, F\}$. All statistics are employment-weighted. *Source:* Model estimates based on RAIS, 2007–2014.

As a holistic measure of model fit, Figure 2 below compares the entire distribution of log firm pay (Panel A) and log firm size (Panel B). Short of perfect, the model fits the data remarkably well.

[47] Appendix E.11 analyzes in detail the sensitivity of our estimation results to alternative amenity cost shares. There, we demonstrate that our estimates of employer amenities are exactly invariant to the chosen target amenity cost share, for which a reliable estimate is difficult to find, and thus the associated amenity cost elasticity. We show that, as a result, our main decomposition results are unaffected by this choice as well.

Figure 2. Model fit in terms of the distributions of log firm pay and log firm size, by gender

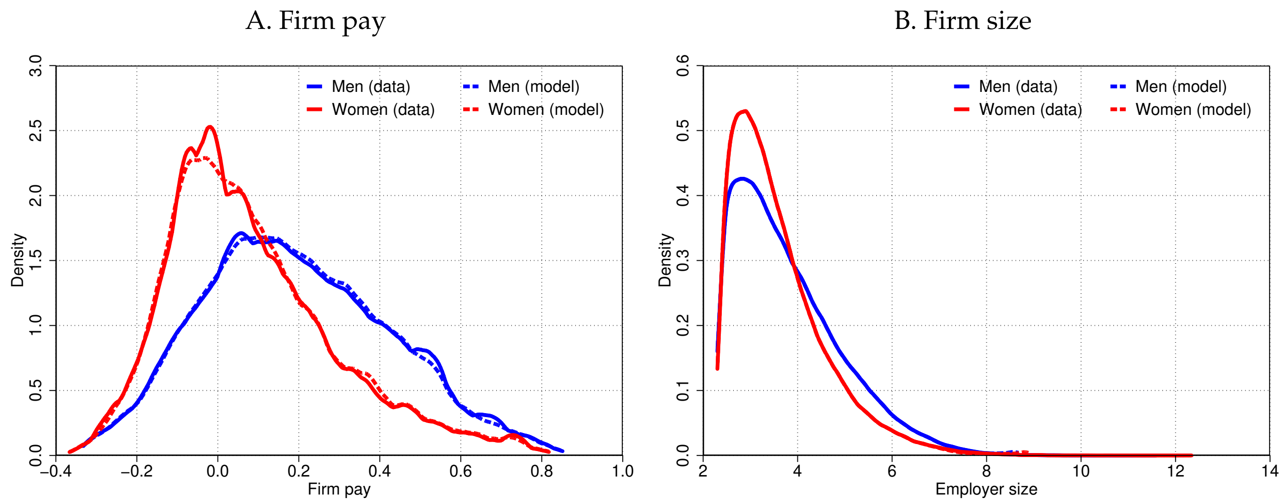

Note: This figure plots the model fit in terms of the data versus model distributions of log firm pay (Panel A) and log firm size (Panel 2B). Blue lines show the results for men, while red lines show the results for women. Solid lines show the data, while dashed lines show the model. *Source:* Model estimates based on RAIS, 2007–2014.

7 Gender-Specific Compensation Structures Across Employers

We now use the estimated model to shed light on the structure of employer compensation by gender.

7.1 Intersectoral Differences in Employer Pay, Amenity Valuations, and Ranks

Grouping our estimates into 25 sectors, Figure 3 plots mean employer ranks against mean pay ranks for men in Panel A and for women in Panel B. The rank-pay relationship is upward-sloping for men but downward-sloping for women. For men, the highest-utility employers—e.g., the automobile sector—are also among the highest-paying ones. For women, however, there is a trade-off between higher pay and higher overall utility—e.g., the public sector offers average pay rank but has the highest utility rank. There are also notable gender differences across sectors. For example, the textile sector has a higher pay rank for men but a higher utility rank for women.[48] Overall, this suggests that preferred employers for men are those with higher pay while the same is not true for women.

Figure 4 plots the estimated employer rank-amenity relationship across the same 25 sectors for men in Panel A and for women in Panel B. For men, mean employer ranks are approximately flat across amenity ranks, suggesting that amenities are not a key determinant of utility for men. For women, however, the rank-amenity relationship is steeply increasing. Therefore, for women more so than for men, employer amenities are a key determinant of overall utility.

[48]See also Sharma (2023) for a comprehensive study of gender-specific monopsony power in the Brazilian textile sector.

Figure 3. Sectoral employer ranks against employer pay ranks, by gender

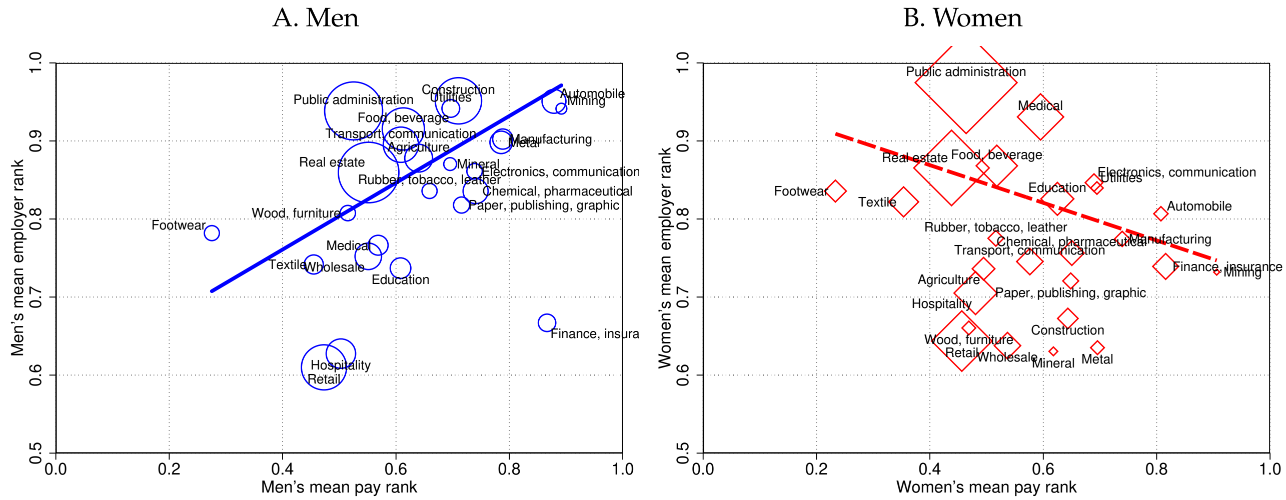

Note: This figure shows ranks of employer utility (x_g) against ranks of employer pay (w_g) for men in Panel A and for women in Panel B across 25 sectors. Marker sizes represent employment weights on which the linear fit lines are based. *Source:* Model estimates based on RAIS, 2007–2014.

Figure 4. Sectoral employer ranks against employer amenity ranks, by gender

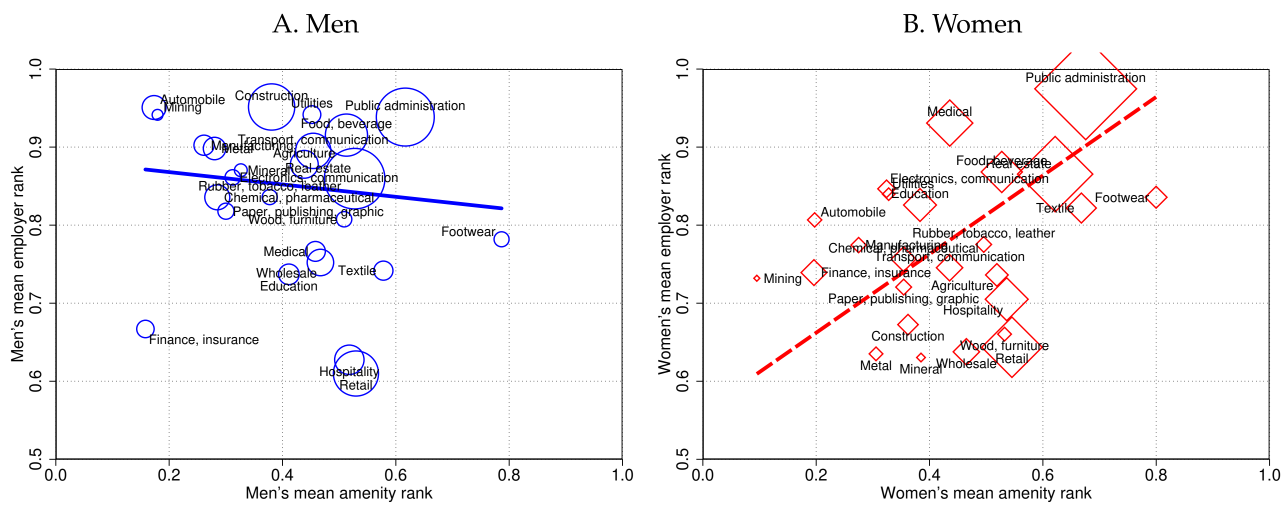

Note: This figure shows ranks of employer utility (x_g) against ranks of employer amenity valuations ($\beta_g a_g$) for men in Panel A and for women in Panel B across 25 sectors. Marker sizes represent employment weights on which the linear fit lines are based. *Source:* Model estimates based on RAIS, 2007–2014.

7.2 The Importance of Amenity Valuations in Total Compensation

What is the relative importance of wages versus amenity valuations in total compensation, and how does the compensation structure vary across employers for men and women? Figure 5 shows the gender-specific distributions of pay and amenity valuations. The distribution of pay in Panel A is repeated, for comparison, from the empirical analysis in Section 3.1. The distribution of amenity values in Panel B shows that women are more concentrated than men at employers offering high amenity values. Women's mean log amenity value is 0.164, while men's is 0.138, implying a gender amenity gap of −2.6 log points.[49]

Figure 5. Estimated distributions of pay and amenity valuations, by gender

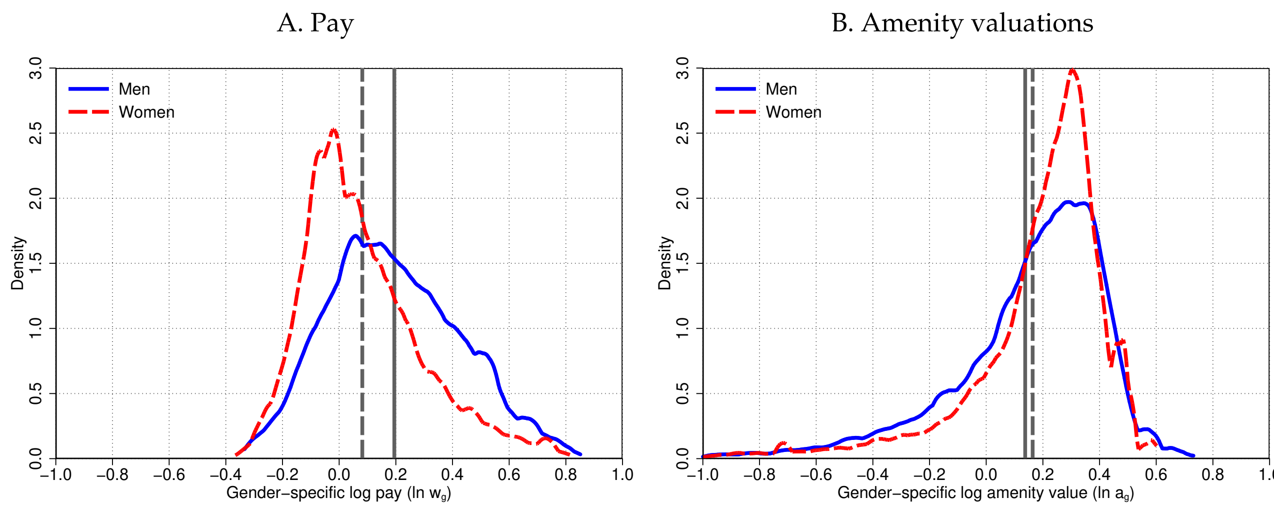

Note: This figure shows the gender-specific employment-weighted distributions of log pay ($\ln w_g$) in Panel A and of log amenity valuations ($\ln(\beta_g a_g)$) in Panel B. The grey patterned vertical lines show the population means for the corresponding gender $g \in \{M, F\}$. *Source:* Model estimates based on RAIS, 2007–2014.

The structure of pay and amenity valuations has important implications for gender-specific firm rankings. Figure 6 shows the relative values of pay, amenity valuations, and total compensation throughout the firm ladder. Relative to bottom-ranked firms, total compensation increases by 18 log points for men and 12 log points for women toward top-ranked firms. For men, more than this amount is due to higher pay at higher firm ranks, while amenity valuations actually decline across the bottom four rank quintiles. For women, pay is relatively flat throughout most of the distribution, so most of the increase in total compensation across firm ranks is due to higher amenity valuations.

Next, we inspect amenity shares in total compensation, $\beta_g a_g / (w_g + \beta_g a_g)$. Panel A of Figure 7

[49]Throughout, we compare pay, amenity valuations, and total compensation in logarithms, consistent with the empirical analysis of log pay. The conversion between levels and logs is purely for presentation purposes and inconsequential to our analysis.

Figure 6. Pay, amenity valuations, and total compensation across employer ranks, by gender

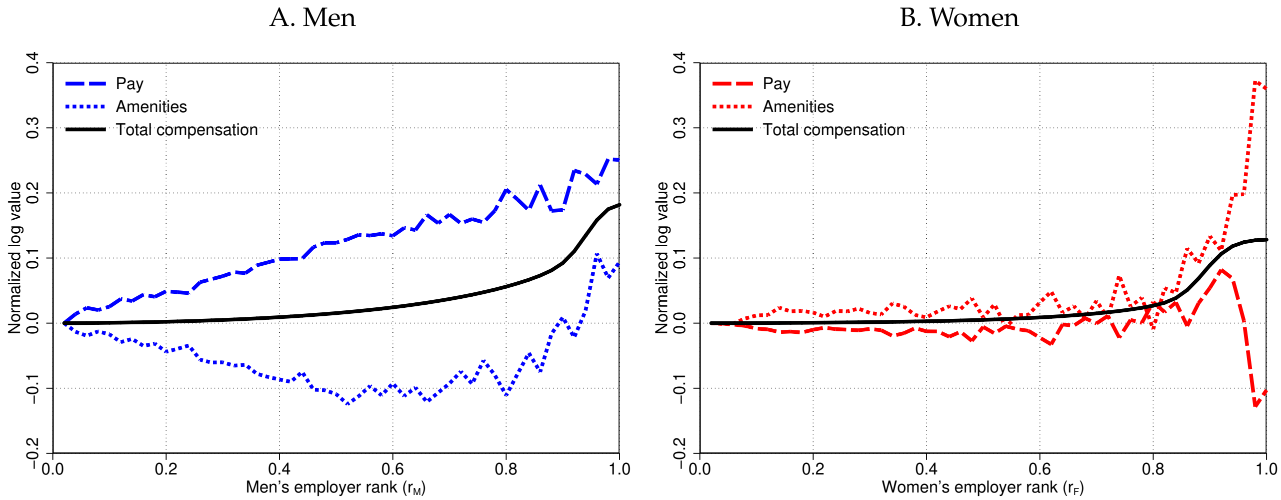

Note: This figure shows the relative values of log pay ($\ln w_g$), log amenity valuations ($\ln(\beta_g a_g)$), and log total compensation ($\ln x_g$) across firm ranks (r_g) separately for men in Panel A and for women in Panel B. All log values are normalized to zero at the gender-specific group of bottom-ranked employers. *Source:* Model estimates based on RAIS, 2007–2014.

shows the density of amenity shares. For both men and women, amenity shares ranges from zero at the low end all the way up to approximately three quarters at the high end. The mean amenity share is 48.8% for men and 52.2% for women. Overall dispersion in amenity shares is lower for women than for men. Panel B shows the estimated amenity shares across employer ranks r_g. For men, the amenity share is decreasing across most employer ranks. For women, the amenity share is mostly flat and then spikes up in the top decile.

A potentially important consideration in interpreting these results is sectoral heterogeneity, in particular the role of the public sector. In Appendix F.3, we conduct a series of analyses to investigate the sensitivity of our results to the in-sample presence of public-sector employers, which form an important part of the Brazilian labor market in general, and more so for women than for men. However, we estimate relatively stable magnitudes of pay gaps, amenity-valuation gaps, and total-compensation gaps across genders in the private sector as well as in the public sector. In this sense, our main results are not sensitive to the presence of the public sector.

7.3 Employer Pay Dispersion Reflecting Utility Dispersion

Our framework's ability to account for differences in amenity valuations across employers allows us to revisit hitherto documented facts about labor market inequality. As shown in Table 10, the variance of log employer pay is 0.054 for men and 0.044 for women. Taken at face value, such dispersion in pay for identical workers across employers suggests significant labor market imperfections for both

Figure 7. Distribution of amenity shares in total compensation, by gender

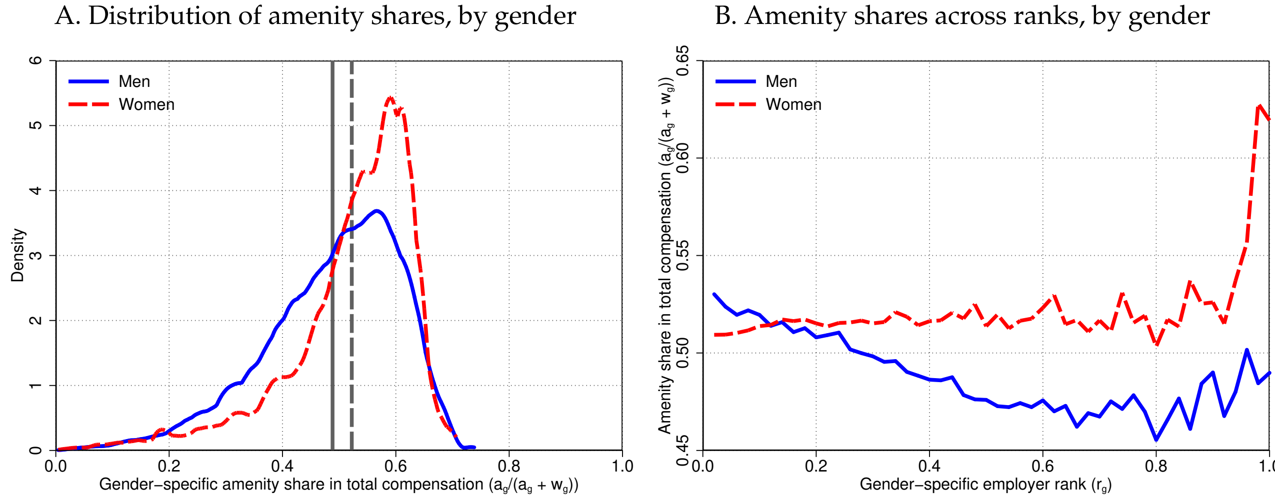

Note: Panel A of this figure shows the gender-specific employment-weighted distribution of amenity shares in total compensation ($a_g/(w_g + a_g)$). The grey patterned vertical lines show the population means for the corresponding gender $g \in \{M, F\}$. Panel B of this figure shows the amenity share in total compensation ($a_g/(w_g + a_g)$) across employer ranks (r_g) separately by gender $g \in \{M, F\}$. *Source:* Model estimates based on RAIS, 2007–2014.

genders. However, we find that the variance of log total compensation across employers is 0.002 (i.e., 4.4% of pay dispersion) for men and 0.002 (i.e., 3.6% of pay dispersion) for women. Thus, the lion's share of firm pay differences are explained by compensating differentials due to firm amenities, with only a small fraction reflecting utility differences. As a result, looking only at differences in pay across employers vastly overstates labor market inequality for both men and women.[50]

Table 10. Decomposition of employer pay dispersion into utility and amenity terms

	Men		Women	
Variances	Level	Share (%)	Level	Share (%)
Variance of log pay	0.054		0.044	
Variance components of log pay:				
Log utility	0.002	4.4	0.002	3.6
Log amenities	0.051	94.3	0.045	102.8
Covariance between log utility and log amenities	0.001	1.3	−0.003	−6.4
Covariance components of log pay:				
Covariance between log utility and log pay	0.003	5.1	0.000	0.4
Covariance between log amenities and log pay	0.052	94.9	0.044	99.6

Note: This table shows the variance and covariance components of log pay ($\ln w$). To this end, we define "amenities" as the utility-to-wage ratio $\tilde{a} = x/w$ so that $\ln w + \ln \tilde{a} = \ln x$. The variance components correspond to the variance decomposition $Var(\ln w) = Var(\ln x) + Var(\ln \tilde{a}) - 2Cov(\ln x, \ln \tilde{a})$. The covariance components correspond to the covariance decomposition $Var(\ln w) = Cov(\ln w, \ln x) - Cov(\ln w, \ln \tilde{a})$. *Source:* Model estimates based on RAIS, 2007–2014.

[50] This result mirrors a similar finding by Lamadon et al. (2022) who estimate a model without search frictions for the U.S. labor market. What is striking is that we find a similarly small role for utility dispersion across firms using a model with search frictions for a labor market characterized by significant imperfections in a developing-country context.

7.4 Gender Gaps in Pay, Amenity Valuations, and Total Compensation

While there is a gender pay gap of 11.3 log points, we find a gender amenity-valuation gap of -6.7 log points in favor of women. As a result, the gender gap in total compensation is 4.6 log points (i.e., 40.9% of the pay gap). That the total compensation gap is lower than the pay gap is a direct consequence of the fact that women work at employers with higher amenity-valuation shares in total compensation.

To shed light on the gender gaps in pay, amenity valuations, and total compensation, Table 11 shows Kitagawa-Oaxaca-Blinder decompositions into gaps between versus within employers.[51] The first row, repeated from the empirical analysis in Section 3, shows that the majority share of the gender pay gap is between employers. The second row shows that the gender amenity-valuation gap of -6.7 log points is due to a larger between-employer component of -8.7 log points, reflecting the sorting of women into high-amenity-valuation firms, and an offsetting within-employer component of 2.0 log points, reflecting the amenity premium that men enjoy over women at the same firm. Altogether, the gender gap in total compensation of 4.6 log points is almost entirely accounted for by the within-employer gap. This suggests that in order to close the gender utility gap, we need to inspect factors leading to unequal treatment of men and women within the same employer—such as the gender wedges τ in our model.[52]

Table 11. Kitagawa-Oaxaca-Blinder decompositions of gaps in pay, amenity valuations, and total compensation

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.089	78.7	0.024	21.3
Amenity-valuation	−0.067	−0.087	130.0	0.020	−30.0
Total compensation	0.046	0.002	4.6	0.044	95.4

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) into a between-employer and a within-employer gap. To this end, we define "amenity valuations" as the utility-to-wage ratio $\tilde{a} = x/w$ so that $\ln w + \ln \tilde{a} = \ln x$. All decompositions are based on equation (2). Table B.1 in Appendix B.2 as well as Tables F.1–F.2 in Appendix F.1 show alternative decompositions using men's compensation for computing the between-employer component of pay, amenity valuations, and total compensation, respectively. *Source:* Model estimates based on RAIS, 2007–2014.

At this point, it is worth taking a step back and highlighting the significance of the results in Table

[51] We obtain similar decomposition results for the whole economy (Table 11), the private sector only, or the public sector only.

[52] Appendix E.10 compares the estimation results between our baseline model and the model extension where we allow separation rates to be heterogeneous and nonincreasing in firm ranks. We demonstrate that our estimates are comparable and that the decomposition results are largely similar to those presented here.

11 in light of existing work on gender pay differences within and between firms (e.g., Card et al., 2016). Like Card et al. (2016) for Portugal—see Table III on p. 667 of their published paper—we find a significant gender gap in firm pay that is mostly between employers (i.e., due to "sorting") rather than within employers (i.e., due to "bargaining"). At the same time, our estimated model suggests that there is a gap in firm amenity valuations in favor of women, which leads the total-compensation gap to be less than half of the pay gap. Strikingly, women sort into employers with relatively high amenity levels compared to the average amenity level enjoyed by men, giving rise to a large negative between-employer amenity-valuation gap. This is partly offset by women enjoying lower amenity levels than men within the same employers, giving rise to a positive within-employer amenity-valuation gap. In other words, women gravitate toward employers with relatively high amenity valuations but they still enjoy lower amenity valuations than their male coworkers at those employers. As a result, the total-compensation gap is almost entirely within, rather than between, employers. These results suggest a novel interpretation of reduced-form wage regressions akin to those in Card et al. (2016) and our paper's Section 3. Our structural estimates reinforce the finding by Card et al. (2016) that women get lower rents at highly productive firms, which they interpreted as being due to (lack of) bargaining but which our structural model rationalizes through equilibrium behavior reflected in firm posting and worker job mobility patterns that differ across genders.

For brevity, we defer to the supplementary materials some additional results on different margins of gender discrimination (Appendix F.2), the implications of our framework for estimates of firm productivity (Appendix F.4), and the effects of switching employment and compensation policies across genders (Appendix F.5).

8 Equilibrium Counterfactuals

In this section, we use the estimated equilibrium model to conduct a series of counterfactual experiments. The equilibrium nature of our model is key because we study counterfactual economies in which the removal of certain model ingredients or the imposition of certain policies changes the optimizing behavior of workers and firms. We rely on the rich, nonparametrically identified distribution of gender-specific employer characteristics to quantify both microeconomic (e.g., the gender pay gap) and macroeconomic (e.g., aggregate output) effects due to these counterfactuals.

8.1 The Role of Gender-Specific Compensating Differentials.

What would gender inequality be in a world without amenity differences across employers? To answer this question, we set all amenities to the mean amenity value for each gender, $\beta_g a_g \mapsto \mathbb{E}[\beta_g a_g]$. The goal is to investigate the role that gender-specific compensating differentials (Rosen, 1986) play in determining gender gaps in pay and utility.

Table 12. Effects of eliminating firm heterogeneity in amenities

	Baseline	Same amenities
Gender log pay gap	0.109	0.057
between employers	0.082	0.020
within employers	0.027	0.037
Gender log amenities gap	−0.066	−0.011
between employers	−0.075	−0.010
within employers	0.009	−0.002
Gender log utility gap	0.042	0.046
between employers	0.007	0.010
within employers	0.035	0.035
Output	1.000	1.016
Worker welfare	1.000	1.014
for men	1.000	1.014
for women	1.000	1.013
Total employment	0.771	0.783
for men	0.764	0.772
for women	0.781	0.799

Note: This table reports results from equilibrium counterfactuals. Log amenities are defined as the log amenity valuations ($\ln \tilde{a} = \ln(x/w)$). The baseline economy (column 0) is compared to a counterfactuals without differences in amenities across employers (column 1) . *Source:* Model estimates based on RAIS, 2007–2014.

As a result, the gender pay gap declines by 5.2 log points, or 47.6 percent of the baseline, due to two channels. First, women relocate away from low-productivity, high-amenity employers, which lowers the between-employer pay gap by 6.2 log points. Second, higher-productivity employers tend to have larger gender wedges and pay women less, which increases the within-employer gap by 1.0 log points. While the gender pay gap closes substantially, the gender utility gap increases because women lose more in amenities than they gain in pay. However, employment rises by 1.2 percentage points, and output rises by 1.6 percent due to the increased competitiveness of high-productivity firms. In net, worker welfare increases by 1.4 percent. Overall, this counterfactual highlights that amenities play an important role in that they shape equilibrium labor market competition by disproportionately helping lower-productivity firms attract and retain workers—especially women.

In Appendix G.3, we perform three additional equilibrium counterfactuals that allow us to shed light on the role of taste-based discrimination (Becker, 1971), labor market frictions (Manning, 2003), and their interplay. We have three main findings. First, amenities interact with gender wedges: low-paying low-productivity firms require both higher amenities and lower gender wedges in order to attract a substantial female workforce. Second, differences in labor market frictions are not an important determinant of gender pay gaps in Brazil. Third, moving to a gender-neutral world implies substantial welfare and output gains.

8.2 The Unintended Effects of Equal-Treatment Policies

Many countries have recently considered or already imposed legislation requiring the equal treatment of men and women in their workforce and their applicant pool.[53] Such policies may be justified by a planner's concern with gender inequity, possibly rooted in firms' unequal treatment of men and women. At the same time, our model's equilibrium may not be efficient due to congestion and business-stealing externalities, so policies can also have (positive or negative) efficiency effects. Given that equal-treatment policies affect firms' incentives to pay, provide amenities, and recruit workers, our equilibrium framework is ideally suited to evaluating such policies. Thus, our model provides an ideal laboratory to study the equilibrium consequences of such policies.[54] To this end, we simulate the effect of imposing an *equal-pay policy* and, separately, an *equal-hiring policy*. For these policy simulations, it matters whether amenities vary due to preference weights β_g or due to amenity cost shifters $c_g^{a,0}$. Here, we assume all variation in amenities is due to variation in amenity cost shifters and that $\beta_g = 1$ for all firms. Our results are summarized in Table 13, which shows the baseline economy in column 0, the equal-pay policy in column 1, and the equal-hiring policy in column 2.[55]

Equal-Pay Policy. Our first policy experiment simulates an equal-pay policy requiring all dual-gender firms to offer the same pay to workers of identical ability, regardless of gender. At first glance, this policy is effective at reducing gender inequality. By construction, the within-employer gender pay gap disappears. In addition, the between-employer gender pay gap also reduces substantially

[53] In the Brazilian context, on July 3, 2023, Brazilian president Luiz Inácio Lula da Silva passed *Law Number 14.611*, amending *Article 461* of the Labor Code, which requires equal pay for men and women performing substantially equal work. In the U.S. context, the *Equal Pay Act of 1963* is an amendment to the *Fair Labor Standards Act* that prohibits employers from paying different wages across the sexes for substantially equal work.

[54] In related work, Lamadon et al. (2022) assume that firms are endowed with a fixed set of amenities but note that *"it would be interesting to extend this analysis to allow for firms to adjust amenities in response to [policy] counterfactuals"* (p. 208).

[55] The fact that equal-treatment policies link the markets for men and women significantly complicates the computation of an equilibrium. Appendices G.1 and G.2 describe the baseline and alternative numerical solution algorithms, respectively.

since firms, which initially paid men more than women, are now forced to offer suboptimally higher wages for women and suboptimally lower wages for men.

However, firms' endogenous responses offset the closing of the gender gap. To partially counteract the equal-pay mandate, firms now offer higher amenities to men, increasing the within-employer amenities gap. Hiring workers of both genders is reduced as firms make lower profits for women and become less attractive employers to men. Consequently, high-gender wedge firms shrink, employment declines, and average pay increases. The latter is because surviving firms with negative gender wedges absorb some of the lost jobs. Aggregate output declines by 1.4 percent, and worker welfare declines slightly in approximately equal proportions for men and women. The bottom line is that an equal-pay policy closes 74.5 percent of the overall pay gap but is detrimental to worker welfare and output due to its negative employment effects and due to firms counteracting the policy through amenity creation.[56]

Equal-Hiring Policy. Our second policy experiment simulates an equal-hiring policy that requires all dual-gender firms to post an equal number of vacancies across the sexes. We find that the equal-hiring policy also reduces gender inequality in some dimensions. Such a policy has large effects on worker reallocation, leading to a reduction in the gender pay gap by 7.5 log points due to a near-elimination of the between-employer pay gap. However, worker welfare actually declines slightly. The reason is that firms, constrained by the policy to make suboptimal hiring decisions, now pay less to both men and women. Specifically, high-productivity firms with positive gender wedges reluctantly hire more women, while low-productivity firms with large amenities hire more men. As a result, the firm ladders of both genders become less directed in the utility space. In particular, women receive relatively more pay but relatively fewer valued amenities. Consequently, the welfare of men and women, along with output, declines. Employment of men declines, while that of women increases substantially. The equal-hiring policy simulation shows that such a heavy-handed policy has large distortionary effects on the allocation of men and women in the labor market, related to the fact that men and women have different valuations of pay versus nonpay attributes across employers.[57]

[56] We also simulate the impact of a policy that requires all dual-gender firms to offer the same amenities to workers of identical ability and report our results in Appendix G.4. We find that this equal-amenities policy has similarly detrimental effects on output and welfare, particularly for women.

[57] A related question is whether a policy targeting only one sector of the economy, for instance, the public sector, may have similar effects through labor market competition. To investigate this, in Appendix G.5, we simulate an equal-hiring policy restricted exclusively to the public sector. We find this policy to have smaller effects on pay gaps than the generalized equal-hiring mandate and no discernible effect on the gender gap in utility.

Table 13. Effects of simulated equal-pay and equal-hiring policies

	Baseline	Equal-pay policy	Equal-hiring policy
	(0)	(1)	(2)
Gender log pay gap	0.109	0.028	0.034
between employers	0.082	0.028	0.006
within employers	0.027	0.000	0.028
Gender log amenities gap	−0.066	0.003	0.011
between employers	−0.075	−0.027	−0.006
within employers	0.009	0.030	0.017
Gender log utility gap	0.042	0.031	0.045
between employers	0.007	0.000	0.000
within employers	0.035	0.030	0.045
Output	1.000	0.986	0.997
Worker welfare	1.000	0.996	0.992
for men	1.000	0.996	0.991
for women	1.000	0.996	0.993
Total employment	0.771	0.763	0.764
for men	0.764	0.760	0.722
for women	0.781	0.767	0.825

Note: Table reports results from two counterfactual policy experiments. Log amenities are defined as the log amenity valuations ($\ln \tilde{a} = \ln(x/w)$). Baseline results (column 0) are compared against the economy with an equal-pay policy (column 1) and the economy with an equal-hiring policy (column 2). *Source:* Model estimates based on RAIS, 2007–2014.

9 Conclusion

In this paper, we studied the micro sources and macro consequences of the gender pay gap. To interpret the empirical fact that men are sorted into employers with higher pay, we developed an equilibrium model of firm pay, amenities, and employment. A notable feature of our model was that it allowed for the allocation of and transfers to workers of both genders to reflect compensating differentials (Rosen, 1986), taste-based discrimination (Becker, 1971), and labor market frictions (Manning, 2003). We provided a constructive proof of identification of all model parameters based on linked employer-employee data. We then used the estimated framework to shed light on the structure of employer compensation by gender and to simulate counterfactual experiments, including equal-treatment policies.

It should be evident that our methodology and quantitative results have many interesting implications beyond the application to gender. Here, we focus on three. First, our empirical and structural estimates suggest that men and women do not share a common employer ranking. Consequently, we should not expect all workers to climb the same job ladder. In allowing for separate job ladders by gender, we have taken but a first step in this direction. Our methodology can be applied to study employer heterogeneity across other population groups (e.g., parental status, education, race).

Second, our finding of a large role for compensating differentials demonstrates that measured inequality in labor market outcomes may overstate—or possibly understate, in other contexts—true inequality. We documented stark differences in the structure of total compensation across men and women, which suggests it may be fruitful to revisit other distributional phenomena, such as cross-country differences in welfare and the uneven evolution of welfare within countries over time.

Third, employers in our framework have more than one margin to adapt to the economic environment. When considering the impact of equal-treatment policies, minimum wages, or income taxation, for example, amenities may move in the opposite direction from pay. Thus, those policies' consequences for equity and efficiency may differ from what data on pay alone might suggest.

A cautionary note regarding the interpretation of our results is in order. Workers' preferences over amenities, such as those relating to work-life balance, plausibly capture the induced demand due to household arrangements and societal norms that differ across men and women (Goldin, 2014, 2023; Cubas et al., 2021, 2023). In this sense, the gender differences we document may reflect not preferences but constraints imposed by nonmarket factors.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge, "When Should You Adjust Standard Errors for Clustering?," *Quarterly Journal of Economics*, 2022, *138* (1), 1–35.

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

_ , Robert H. Creecy, and Francis Kramarz, "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," *U.S. Census Bureau Technical Report 2002-06*, 2002.

Abraham, Lisa, Johannes Hallermeier, and Alison Stein, "Words matter: Experimental Evidence from Job Applications," *Journal of Economic Behavior and Organization*, 2024, *225*, 348–391.

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 and Ayşegül Şahin, "The Gender Unemployment Gap," *Review of Economic Dynamics*, 2018, *30*, 47–67.

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.

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.

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," *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.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading off Commute Against Wage," *Quarterly Journal of Economics*, 2020, *136* (1), 381–426.

Barth, Erling, Sari Pekkala Kerr, and Claudia Olivetti, "The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data," *European Economic Review*, 2021, *134*, 103713.

Becker, Gary S., *The Economics of Discrimination*, second ed., Chicago: University of Chicago Press, 1971.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–1193.

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz, "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2010, 2 (3), 228–255.

Bieri, David S., Nicolai V. Kuminoff, and Jaren C. Pope, "National Expenditures on Local Amenities," *Journal of Environmental Economics and Management*, 2023, *117*, 102717.

Bilal, Adrien and Hugo Lhuillier, "Outsourcing, Inequality and Aggregate Output," *Working Paper*, 2023.

_ , Niklas Engbom, Simon Mongey, and Giovanni L. Violante, "Firm and Worker Dynamics in a Frictional Labor Market," *Econometrica*, 2022, *90* (4), 1425–1462.

Black, Dan A., "Discrimination in an Equilibrium Search Model," *Journal of Labor Economics*, 1995, *13* (2), 309–334.

Bloesch, Justin and Birthe Larsen, "When do Firms Profit from Wage Setting Power? New vs. Classical Monopsony," *Working Paper*, 2023.

Boerma, Job and Loukas Karabarbounis, "Inferring Inequality With Home Production," *Econometrica*, 2021, *89* (5), 2517–2556.

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?," *Journal of Labor Economics*, 2023, *41* (2), 291–322.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa, "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 2019, *87* (3), 699–739.

— **,** — **, and** — , "Discretizing Unobserved Heterogeneity," *Econometrica*, 2022, *90* (2), 625–643.

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," *NBER Working Paper No. 24074*, 2020.

Bowlus, Audra J., "A Search Interpretation of Male-Female Wage Differentials," *Journal of Labor Economics*, 1997, *15* (4), 625–657.

Bruns, Benjamin, "Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap," *American Economic Journal: Applied Economics*, 2019, *11* (2), 74–113.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–273.

Caldwell, Sydnee and Nikolaj Harmon, "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks," *Working Paper*, 2019.

— **and Oren Danieli**, "Outside Options in the Labor Market," *Working Paper*, 2023.

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.

Casarico, Alessandra and Salvatore Lattanzio, "What Firms Do: Gender Inequality in Linked Employer-Employee Data," *Working Paper*, 2022.

Chen, Kuan-Ming, Claire Ding, John A List, and Magne Mogstad, "Reservation Wages and Workers' Valuation of Job Flexibility: Evidence from a Natural Field Experiment," *NBER Working Paper No. 27807*, 2020.

Cortés, Patricia, Jessica Pan, Laura Pilossoph, and Basit Zafar, "Gender Differences in Job Search and the Earnings Gap: Evidence from Business Majors," *Working Paper*, 2023.

Coudin, Elise, Sophie Maillard, and Maxime Tô, "Family, Firms and the Gender Wage Gap in France," *IFS Working Paper No. W18/01*, 2018.

Cruz, Gabriel and Tomás Rau, "The Effects of Equal Pay Laws on Firm Pay Premiums: Evidence from Chile," *Labour Economics*, 2022, *75*, 102135.

Cubas, German, Chinhui Juhn, and Pedro Silos, "Work-Care Balance over the Day and the Gender Wage Gap," *AEA Papers and Proceedings*, 2021, *111*, 149–153.

— **,** — **, and** — , "Coordinated Work Schedules and the Gender Wage Gap," *Economic Journal*, 2023, *133* (651), 1036–1066.

de Souza, Gustavo, "The Labor Market Consequences of Appropriate Technology," *Federal Reserve Bank of Chicago Working Paper N. 2022-53*, 2022.

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*, 2021.

— **, Pinelopi Koujianou Goldberg, Costas Meghir, and Gabriel Ulyssea**, "Trade and Domestic Distortions: The Case of Informality," *Cowles Foundation Discussion Paper No. 2384*, 2024.

Engbom, Niklas and Christian Moser, "Earnings Inequality and the Minimum Wage: Evidence from Brazil," *American Economic Review*, 2022, *112* (12), 3803–3847.

__ , Gustavo Gonzaga, Christian Moser, and Roberta Olivieri, "Earnings Inequality and Dynamics in the Presence of Informality: The Case of Brazil," *Quantitative Economics*, 2022, *13* (4), 1405–1446.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–3182.

Felix, Mayara, "Trade, Labor Market Concentration, and Wages," *Working Paper*, 2022.

Firpo, Sergio and Alysson Portella, "Decline in Wage Inequality in Brazil: A Survey," *World Bank Policy Research Working Paper No. 9096*, 2019.

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.

Flinn, Christopher, Petra E. Todd, and Weilong Zhang, "Labor Market Returns to Personality: A Job Search Approach to Understanding Gender Gaps," *Working Paper*, 2023.

Folke, Olle and Johanna Rickne, "Sexual Harassment and Gender Inequality in the Labor Market," *Quarterly Journal of Economics*, 2022, *137* (4), 2163–2212.

Fukui, Masao, Emi Nakamura, and Jón Steinsson, "Women, Wealth Effects, and Slow Recoveries," *American Economic Journal: Macroeconomics*, 2023, *15* (1), 269–313.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative Matching or Exclusionary Hiring? The Impact of Employment and Pay Policies on Racial Wage Differences in Brazil," *American Economic Review*, 2021, *111* (10), 3418–3457.

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.

__ , "An Evolving Economic Force," *Nobel Lecture*, 2023.

__ and Cecilia Rouse, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *American Economic Review*, 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.

__ and Paul R. Milgrom, "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 2008, *98* (4), 1653–1674.

Hamermesh, Daniel S., "Changing Inequality in Markets for Workplace Amenities," *Quarterly Journal of Economics*, 1999, *114* (4), 1085–1123.

Heise, Sebastian and Tommaso Porzio, "Spatial Wage Gaps in Frictional Labor Markets," *Federal Reserve Bank of Minneapolis OIGI Working Paper 29*, 2023.

Hotz, V. Joseph, Per Johansson, and Arizo Karimi, "Parenthood, Family Friendly Workplaces, and the Gender Gaps in Early Work Careers," *Working Paper*, 2018.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *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," *Econometrica*, 2023, *91* (3), 903–942.

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 U.S.," *European Economic Review*, 2006, *50* (4), 877–907.

Juhn, Chinhui and Kristin McCue, "Specialization Then and Now: Marriage, Children, and the Gender Earnings Gap across Cohorts," *Journal of Economic Perspectives*, 2017, *31* (1), 183–204.

Jung, Philip and Moritz Kuhn, "Earnings Losses and Labor Mobility over the Life Cycle," *Journal of the European Economic Association*, 2019, *17* (3), 678–724.

Kaas, Leo and Philipp Kircher, "Efficient Firm Dynamics in a Frictional Labor Market," *American Economic Review*, 2015, *105* (10), 3030–3060.

Kalemli-Özcan, Şebnem, Bent E. Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş, "How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration," *American Economic Journal: Macroeconomics*, 2024, *16* (2), 353–374.

Kim, Hyuncheol Bryant, Hyunseob Kim, and John Y. Zhu, "The Selection Effects of Part-Time Work: Experimental Evidence from a Large-Scale Recruitment Drive," *Working Paper*, 2024.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten, "Leave-Out Estimation of Variance Components," *Econometrica*, 2020, *88* (5), 1859–1898.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market," *American Economic Review*, 2022, *112* (1), 169–212.

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*, 2018.

Lentz, Rasmus, "Models of Wages and Mobility in Frictional Labor Markets with Random Search," *Revue Economique*, 2024, *51* (1), 113–145.

_ and Dale T. Mortensen, "Labor Market Models of Worker and Firm Heterogeneity," *Annual Review of Economics*, 2010, *2* (1), 577–602.

_ , Suphanit Piyapromdee, and Jean-Marc Robin, "The Anatomy of Sorting: Evidence from Danish Data," *Working Paper*, 2023.

Lindenlaub, Ilse, Ryungha Oh, and Michael Peters, "Firm Sorting and Spatial Inequality," *NBER Working Paper No. 30637*, 2023.

Lise, Jeremy and Jean-Marc Robin, "The Macrodynamics of Sorting between Workers and Firms," *American Economic Review*, 2017, *107* (4), 1104–1135.

Luo, Mi and Simon Mongey, "Assets and Job Choice: Student Debt, Wages and Job Satisfaction," *NBER Working Paper No. 25801*, 2019.

Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger, "The Value of Working Conditions in the United States and the Implications for the Structure of Wages," *American Economic Review*, 2023, *113* (7), 2007–2047.

Manning, Alan, *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton, NJ: Princeton University Press, 2003.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017, *107* (12), 3722–3759.

_ and _ , "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field," *American Economic Review: Insights*, 2019, *1* (1), 111–126.

Meghir, Costas, Renata Narita, and Jean-Marc Robin, "Wages and Informality in Developing Countries," *American Economic Review*, 2015, *105* (4), 1509–1546.

Mortensen, Dale T., *Wage Dispersion: Why Are Similar Workers Paid Differently?*, MIT Press, 2003.

Moscarini, Giuseppe and Fabien Postel-Vinay, "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.

Muendler, Marc-Andreas, "Trade, Technology, and Productivity: A Study of Brazilian Manufacturers, 1986–1998," *CESifo Working Paper No. 1148*, 2004.

Mulligan, Casey B. and Yona Rubinstein, "Selection, Investment, and Womens Relative Wages Over Time," *Quarterly Journal of Economics*, 2008, *123* (3), 1061–1110.

Page, Lawrence, Sergey Brin, Rajeev Motwani, and Terry Winograd, "The PageRank Citation Ranking: Bringing Order to the Web," *Working Paper*, 1998.

Palladino, Marco G., Alexandra Roulet, and Mark Stabile, "Understanding the Role of Firms in the Gender Wage Gap over Time and Across Worker Types," *CEPR Discussion Paper No. 16671*, 2022.

Petrongolo, Barbara and Christopher A. Pissarides, "Looking into the Black Box: A Survey of the

Matching Function," *Journal of Economic Literature*, 2001, *39* (2), 390–431.

Pierce, Brooks, "Compensation Inequality," *Quarterly Journal of Economics*, 2001, *116* (4), 1493–1525.

Restuccia, Diego and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, 2008, *11* (4), 707–720.

Robinson, Joan, *The Economics of Imperfect Competition*, London: Macmillan, 1933.

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.

Sharma, Garima, "Monopsony and Gender," *Working Paper*, 2023.

Sockin, Jason, "Show Me the Amenity: Are Higher-Paying Firms Better All Around?," *Working Paper*, 2022.

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, "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," *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.

Ulyssea, Gabriel, "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review*, 2018, *108* (8), 2015–2047.

Vasconcelos, Rafael, "Misallocation in the Brazilian Manufacturing Sector," *Brazilian Review of Econometrics*, 2017, *37* (2), 191–232.

Vattuone, Giulia, "Worker Sorting and the Gender Wage Gap," *Working Paper*, 2023.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *Quarterly Journal of Economics*, 2017, *133* (1), 457–507.

World Bank, "Birth Rate, Crude (per 1,000 People) - Brazil," 2021. Online database accessed on August 3, 2023, at https://data.worldbank.org/indicator/SP.DYN.CBRT.IN?locations=BR.

—, "Labor Force, Female (% of Total Labor Force) - Brazil," 2021. Online database accessed on August 3, 2023, at https://genderdata.worldbank.org/indicators/sp-dyn-cbrt-in.

Xiao, Pengpeng, "Equilibrium Sorting and the Gender Wage Gap," *Working Paper*, 2023.

The Gender Pay Gap:
Micro Sources and Macro Consequences

Online Appendix—Not for Publication

Iacopo Morchio
University of Bristol

Christian Moser
Columbia University, NBER, CEPR, and IZA

Contents of the Online Appendix

A Data Description Appendix

A.1 Comparison of Actual versus Potential Experience

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

A. Men B. Women

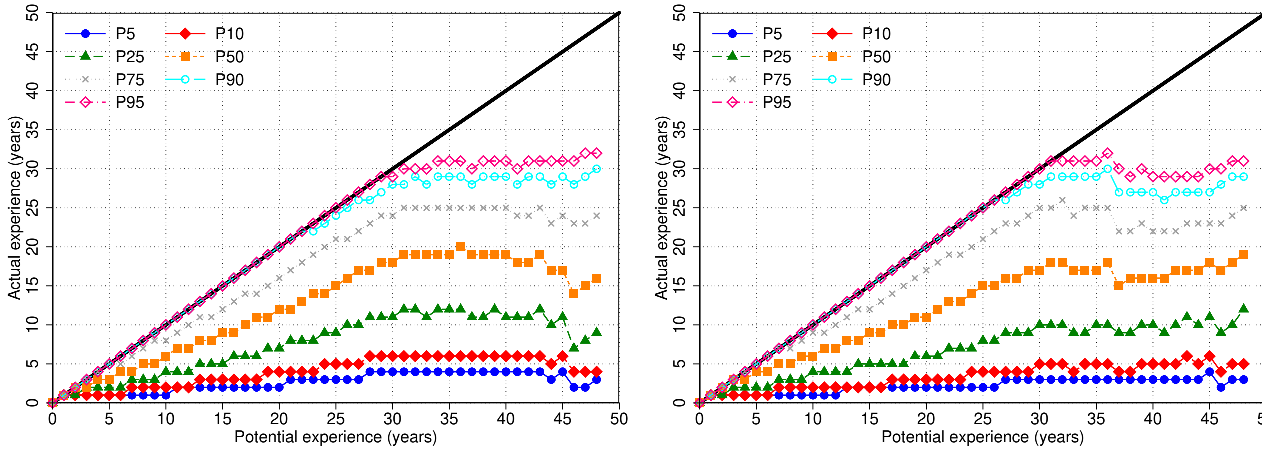

Note: Figure shows percentiles of actual against potential experience separately for men (Panel A) and women (Panel B). Actual experience is constructed from panel data for 1985–2014, while potential experience = age − years of education + 6. Solid line represents the 45-degree line, for which actual experience equals potential experience. *Source:* RAIS, 2007–2014.

A.2 Additional Summary Statistics

Table A.1. Summary statistics for the raw data

	2007			2014			Pooled 2007–2014		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share Nonwhite	0.338	0.369	0.288	0.390	0.420	0.347	0.365	0.397	0.318
Share primary school	0.107	0.143	0.051	0.067	0.093	0.031	0.085	0.115	0.039
Share middle school	0.221	0.268	0.146	0.180	0.223	0.119	0.199	0.245	0.130
Share high school	0.496	0.470	0.538	0.589	0.569	0.617	0.548	0.524	0.584
Share college	0.176	0.119	0.265	0.164	0.114	0.233	0.168	0.115	0.246
Mean years of education	10.8	10.1	11.9	11.2	10.7	12.1	11.0	10.4	12.0
(std. dev.)	(3.4)	(3.5)	(3.0)	(3.0)	(3.1)	(2.6)	(3.2)	(3.3)	(2.8)
Mean years of age	32.8	32.8	32.8	33.5	33.5	33.5	33.0	33.1	33.0
(std. dev.)	(9.4)	(9.4)	(9.4)	(9.5)	(9.5)	(9.4)	(9.4)	(9.4)	(9.4)
Mean employer size	2,566	1,401	4,421	2,436	1,694	3,479	2,406	1,518	3,742
(std. dev.)	(18,390)	(12,175)	(25,204)	(16,366)	(12,398)	(20,651)	(16,264)	(11,559)	(21,410)
Mean gender-employer size	1,600	613	3,171	1,503	896	2,357	1,495	761	2,599
(std. dev.)	(11,892)	(3,325)	(18,577)	(9,927)	(4,978)	(14,178)	(10,194)	(4,173)	(15,231)
Mean employer age	23.3	21.6	25.8	23.4	22.0	25.5	23.1	21.6	25.4
(std. dev.)	(21.3)	(19.9)	(23.1)	(22.3)	(20.9)	(23.8)	(21.8)	(20.4)	(23.5)
Mean months employed during year	9.6	9.5	9.7	9.7	9.6	9.8	9.6	9.5	9.7
(std. dev.)	(3.3)	(3.3)	(3.3)	(3.2)	(3.1)	(3.2)	(3.2)	(3.2)	(3.3)
Mean contractual work hours	41.8	42.6	40.4	41.8	42.6	40.8	41.9	42.6	40.7
(std. dev.)	(5.7)	(4.5)	(7.0)	(5.4)	(4.3)	(6.4)	(5.4)	(4.3)	(6.6)
Mean years of tenure	3.4	3.1	3.9	3.2	3.0	3.4	3.2	2.9	3.5
(std. dev.)	(5.1)	(4.8)	(5.6)	(5.0)	(4.8)	(5.2)	(5.0)	(4.7)	(5.4)
Mean log real monthly earnings	6.984	7.032	6.908	7.222	7.283	7.136	7.112	7.166	7.032
(std. dev.)	(0.680)	(0.689)	(0.659)	(0.631)	(0.640)	(0.606)	(0.656)	(0.666)	(0.633)
Number of worker-years	46,347,012	28,474,352	17,872,660	62,571,376	36,561,824	26,009,552	449,390,272	269,897,824	179,492,464
Number of unique workers	38,412,504	23,362,020	15,050,485	50,872,600	29,455,332	21,417,268	77,397,648	44,436,752	32,960,898
Number of unique employers	2,775,579	1,710,549	1,065,030	3,677,552	2,131,329	1,546,223	6,044,593	3,420,973	2,623,620
Share female	0.386			0.416			0.399		
Mean log gender earnings gap	0.124			0.147			0.134		

Note: Table shows summary statistics for the raw data for 2007, 2014, and the pooled years 2007–2014. *Source:* RAIS, 2007–2014.

Table A.2. Summary statistics for selected sample

	2007			2014			Pooled 2007–2014		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share Nonwhite	0.351	0.381	0.301	0.398	0.429	0.350	0.378	0.409	0.329
Share primary school	0.106	0.138	0.053	0.072	0.097	0.034	0.088	0.117	0.043
Share middle school	0.219	0.263	0.144	0.185	0.226	0.122	0.203	0.245	0.134
Share high school	0.488	0.472	0.515	0.562	0.553	0.577	0.531	0.516	0.554
Share college	0.187	0.126	0.289	0.180	0.124	0.267	0.179	0.122	0.269
Mean years of education	10.9	10.2	12.0	11.3	10.7	12.2	11.1	10.4	12.1
(std. dev.)	(3.4)	(3.5)	(3.1)	(3.1)	(3.2)	(2.8)	(3.3)	(3.3)	(2.9)
Mean years of age	33.0	32.9	33.3	34.6	34.5	34.8	33.6	33.5	33.8
(std. dev.)	(9.2)	(9.2)	(9.3)	(9.3)	(9.3)	(9.3)	(9.4)	(9.4)	(9.4)
Mean employer size	2,784	1,528	4,875	2,961	2,064	4,332	2,775	1,757	4,402
(std. dev.)	(17,415)	(11,393)	(24,188)	(16,617)	(12,450)	(21,404)	(16,298)	(11,448)	(21,828)
Mean gender-employer size	1,744	698	3,485	1,852	1,154	2,920	1,743	926	3,048
(std. dev.)	(11,287)	(3,137)	(17,840)	(10,211)	(5,419)	(14,726)	(10,314)	(4,413)	(15,574)
Mean employer age	27.9	25.4	31.9	29.7	27.3	33.5	28.0	25.7	31.7
(std. dev.)	(22.1)	(20.8)	(23.7)	(23.4)	(22.2)	(24.8)	(22.8)	(21.5)	(24.3)
Mean months employed during year	9.9	9.7	10.1	10.2	10.0	10.3	9.9	9.8	10.0
(std. dev.)	(3.2)	(3.3)	(3.2)	(3.0)	(3.0)	(2.9)	(3.2)	(3.2)	(3.1)
Mean contractual work hours	41.6	42.6	40.0	41.6	42.5	40.3	41.7	42.6	40.3
(std. dev.)	(5.4)	(4.1)	(6.8)	(5.1)	(3.9)	(6.3)	(5.1)	(3.9)	(6.4)
Mean years of tenure	4.0	3.5	4.8	4.2	3.9	4.6	3.9	3.6	4.5
(std. dev.)	(5.6)	(5.2)	(6.2)	(5.7)	(5.4)	(6.1)	(5.6)	(5.2)	(6.1)
Mean log real monthly earnings	7.072	7.119	6.995	7.348	7.407	7.259	7.206	7.259	7.122
(std. dev.)	(0.709)	(0.713)	(0.696)	(0.674)	(0.675)	(0.662)	(0.693)	(0.697)	(0.678)
Number of worker-years	27,609,184	17,248,420	10,360,763	35,134,432	21,238,264	13,896,165	271,400,512	166,964,240	104,436,280
Number of unique workers	27,609,184	17,248,420	10,360,763	35,134,432	21,238,264	13,896,165	56,868,704	33,975,716	22,892,988
Number of unique employers	425,173	295,134	130,039	492,155	334,905	157,250	640,639	420,451	220,188
Share female	0.375			0.396			0.385		
Mean log gender earnings gap	0.124			0.149			0.137		

Note: Table shows summary statistics for the selected sample for 2007, 2014, and the pooled years 2007–2014. *Source:* RAIS, 2007–2014.

Table A.3. Summary statistics for the connected set

	2007			2014			Pooled 2007–2014		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share Nonwhite	0.350	0.380	0.299	0.397	0.429	0.349	0.378	0.409	0.327
Share primary school	0.105	0.137	0.052	0.072	0.096	0.034	0.088	0.116	0.042
Share middle school	0.217	0.262	0.142	0.185	0.226	0.122	0.202	0.245	0.133
Share high school	0.487	0.473	0.512	0.562	0.553	0.576	0.531	0.516	0.554
Share college	0.190	0.128	0.294	0.181	0.124	0.268	0.179	0.122	0.272
Mean years of education	10.9	10.2	12.1	11.3	10.7	12.2	11.1	10.4	12.1
(std. dev.)	(3.4)	(3.5)	(3.1)	(3.1)	(3.2)	(2.8)	(3.3)	(3.3)	(2.9)
Mean years of age	33.1	32.9	33.4	34.6	34.5	34.8	33.6	33.5	33.8
(std. dev.)	(9.3)	(9.2)	(9.3)	(9.3)	(9.3)	(9.3)	(9.4)	(9.4)	(9.4)
Mean employer size	2,881	1,573	5,083	2,987	2,075	4,393	2,815	1,774	4,497
(std. dev.)	(17,749)	(11,589)	(24,733)	(16,687)	(12,481)	(21,550)	(16,418)	(11,509)	(22,059)
Mean gender-employer size	1,805	719	3,635	1,868	1,160	2,961	1,768	936	3,114
(std. dev.)	(11,505)	(3,190)	(18,244)	(10,254)	(5,432)	(14,827)	(10,390)	(4,436)	(15,739)
Mean employer age	28.3	25.8	32.6	29.8	27.3	33.6	28.1	25.7	32.0
(std. dev.)	(22.2)	(20.8)	(23.7)	(23.5)	(22.2)	(24.8)	(22.8)	(21.5)	(24.3)
Mean months employed during year	9.9	9.8	10.1	10.1	10.0	10.3	9.9	9.8	10.0
(std. dev.)	(3.2)	(3.2)	(3.2)	(3.0)	(3.0)	(2.9)	(3.2)	(3.2)	(3.1)
Mean contractual work hours	41.6	42.5	39.9	41.6	42.5	40.3	41.7	42.6	40.3
(std. dev.)	(5.5)	(4.1)	(6.9)	(5.1)	(3.9)	(6.3)	(5.1)	(3.9)	(6.4)
Mean years of tenure	4.0	3.6	4.8	4.2	3.9	4.6	3.9	3.6	4.5
(std. dev.)	(5.6)	(5.2)	(6.3)	(5.7)	(5.4)	(6.1)	(5.6)	(5.2)	(6.1)
Mean log real monthly earnings	7.081	7.125	7.007	7.351	7.408	7.263	7.211	7.262	7.129
(std. dev.)	(0.712)	(0.715)	(0.699)	(0.674)	(0.675)	(0.662)	(0.693)	(0.697)	(0.679)
Number of worker-years	26,545,820	16,656,529	9,889,290	34,829,856	21,129,484	13,700,373	267,318,328	165,149,632	102,168,696
Number of unique workers	26,545,820	16,656,529	9,889,290	34,829,856	21,129,484	13,700,373	56,297,308	33,761,656	22,535,652
Number of unique employers	396,269	278,455	117,814	477,942	328,628	149,314	607,029	403,585	203,444
Share female	0.373			0.393			0.382		
Mean log gender earnings gap	0.118			0.146			0.133		

Note: Table shows summary statistics for the connected set for 2007, 2014, and the pooled years 2007–2014. *Source:* RAIS, 2007–2014.

Table A.4. Comparison of summary statistics (selection vs. all)

	Pooled 2007–2014, selection set			Pooled 2007–2014, all			Ratio: selection set to all		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share Nonwhite	0.378	0.409	0.329	0.365	0.397	0.318	103.6	103.0	103.5
Share primary school	0.088	0.117	0.043	0.085	0.115	0.039	103.5	101.7	110.3
Share middle school	0.203	0.245	0.134	0.199	0.245	0.130	102.0	100.0	103.1
Share high school	0.531	0.516	0.554	0.548	0.524	0.584	96.9	98.5	94.9
Share college	0.179	0.122	0.269	0.168	0.115	0.246	106.5	106.1	109.3
Mean years of education	11.1	10.4	12.1	11.0	10.4	12.0	100.9	100.0	100.8
(std. dev.)	(3.3)	(3.3)	(2.9)	(3.2)	(3.3)	(2.8)	103.1	100.0	103.6
Mean years of age	33.6	33.5	33.8	33.0	33.1	33.0	101.8	101.2	102.4
(std. dev.)	(9.4)	(9.4)	(9.4)	(9.4)	(9.4)	(9.4)	100.0	100.0	100.0
Mean employer size	2,775	1,757	4,402	2,406	1,518	3,742	115.3	115.7	117.6
(std. dev.)	(16,298)	(11,448)	(21,828)	(16,264)	(11,559)	(21,410)	100.2	99.0	102.0
Mean gender-employer size	1,743	926	3,048	1,495	761	2,599	116.6	121.7	117.3
(std. dev.)	(10,314)	(4,413)	(15,574)	(10,194)	(4,173)	(15,231)	101.2	105.8	102.3
Mean employer age	28.0	25.7	31.7	23.1	21.6	25.4	121.2	119.0	124.8
(std. dev.)	(22.8)	(21.5)	(24.3)	(21.8)	(20.4)	(23.5)	104.6	105.4	103.4
Mean months employed during year	9.9	9.8	10.0	9.6	9.5	9.7	103.1	103.2	103.1
(std. dev.)	(3.2)	(3.2)	(3.1)	(3.2)	(3.2)	(3.3)	100.0	100.0	93.9
Mean contractual work hours	41.7	42.6	40.3	41.9	42.6	40.7	99.5	100.0	99.0
(std. dev.)	(5.1)	(3.9)	(6.4)	(5.4)	(4.3)	(6.6)	94.4	90.7	97.0
Mean years of tenure	3.9	3.6	4.5	3.2	2.9	3.5	121.9	124.1	128.6
(std. dev.)	(5.6)	(5.2)	(6.1)	(5.0)	(4.7)	(5.4)	112.0	110.6	113.0
Mean log real monthly earnings	7.206	7.259	7.122	7.112	7.166	7.032	101.3	101.3	101.3
(std. dev.)	(0.693)	(0.697)	(0.678)	(0.656)	(0.666)	(0.633)	105.6	104.7	107.1
Number of worker-years	271,400,512	166,964,240	104,436,280	449,390,272	269,897,824	179,492,464	60.4	61.9	58.2
Number of unique workers	56,868,704	33,975,716	22,892,988	77,397,648	44,436,752	32,960,898	73.5	76.5	69.5
Number of unique employers	640,639	420,451	220,188	6,044,593	3,420,973	2,623,620	10.6	12.3	8.4
Share female	0.385			0.399			96.5		
Mean log gender earnings gap	0.137			0.134			102.2		

Note: Table shows summary statistics for the selected sample, all observations, and their ratio for the pooled years 2007–2014. *Source:* RAIS, 2007–2014.

Table A.5. Comparison of summary statistics (connected vs. selection)

	Pooled 2007–2014, connected set			Pooled 2007–2014, selection set			Ratio: connected set to selection set		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share Nonwhite	0.378	0.409	0.327	0.378	0.409	0.329	100.0	100.0	99.4
Share primary school	0.088	0.116	0.042	0.088	0.117	0.043	100.0	99.1	97.7
Share middle school	0.202	0.245	0.133	0.203	0.245	0.134	99.5	100.0	99.3
Share high school	0.531	0.516	0.554	0.531	0.516	0.554	100.0	100.0	100.0
Share college	0.179	0.122	0.272	0.179	0.122	0.269	100.0	100.0	101.1
Mean years of education	11.1	10.4	12.1	11.1	10.4	12.1	100.0	100.0	100.0
(std. dev.)	(3.3)	(3.3)	(2.9)	(3.3)	(3.3)	(2.9)	100.0	100.0	100.0
Mean years of age	33.6	33.5	33.8	33.6	33.5	33.8	100.0	100.0	100.0
(std. dev.)	(9.4)	(9.4)	(9.4)	(9.4)	(9.4)	(9.4)	100.0	100.0	100.0
Mean employer size	2,815	1,774	4,497	2,775	1,757	4,402	101.4	101.0	102.2
(std. dev.)	(16,418)	(11,509)	(22,059)	(16,298)	(11,448)	(21,828)	100.7	100.5	101.1
Mean gender-employer size	1,768	936	3,114	1,743	926	3,048	101.4	101.1	102.2
(std. dev.)	(10,390)	(4,436)	(15,739)	(10,314)	(4,413)	(15,574)	100.7	100.5	101.1
Mean employer age	28.1	25.7	32.0	28.0	25.7	31.7	100.4	100.0	100.9
(std. dev.)	(22.8)	(21.5)	(24.3)	(22.8)	(21.5)	(24.3)	100.0	100.0	100.0
Mean months employed during year	9.9	9.8	10.0	9.9	9.8	10.0	100.0	100.0	100.0
(std. dev.)	(3.2)	(3.2)	(3.1)	(3.2)	(3.2)	(3.1)	100.0	100.0	100.0
Mean contractual work hours	41.7	42.6	40.3	41.7	42.6	40.3	100.0	100.0	100.0
(std. dev.)	(5.1)	(3.9)	(6.4)	(5.1)	(3.9)	(6.4)	100.0	100.0	100.0
Mean years of tenure	3.9	3.6	4.5	3.9	3.6	4.5	100.0	100.0	100.0
(std. dev.)	(5.6)	(5.2)	(6.1)	(5.6)	(5.2)	(6.1)	100.0	100.0	100.0
Mean log real monthly earnings	7.211	7.262	7.129	7.206	7.259	7.122	100.1	100.0	100.1
(std. dev.)	(0.693)	(0.697)	(0.679)	(0.693)	(0.697)	(0.678)	100.0	100.0	100.1
Number of worker-years	267,318,328	165,149,632	102,168,696	271,400,512	166,964,240	104,436,280	98.5	98.9	97.8
Number of unique workers	56,297,308	33,761,656	22,535,652	56,868,704	33,975,716	22,892,988	99.0	99.4	98.4
Number of unique employers	607,029	403,585	203,444	640,639	420,451	220,188	94.8	96.0	92.4
Share female	0.382			0.385			99.2		
Mean log gender earnings gap	0.133			0.137			97.1		

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

B Empirical Gender Pay Gaps and Employer Heterogeneity Appendix

B.1 Detailed AKM Estimation Results

Figure B.1. Predicted AKM contractual work hours FEs, by gender

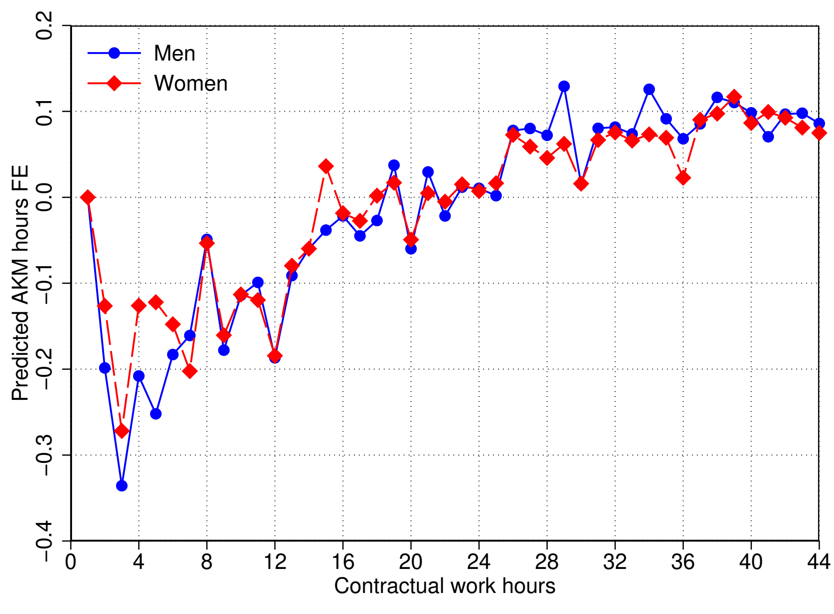

Note: Figure shows predicted AKM contractual work hours FEs separately for men and women based on estimating earnings equation (1). The omitted category is 1 hour, for which the FE value is normalized to 0. *Source:* RAIS, 2007–2014.

Figure B.2. Predicted AKM occupation FEs, by gender

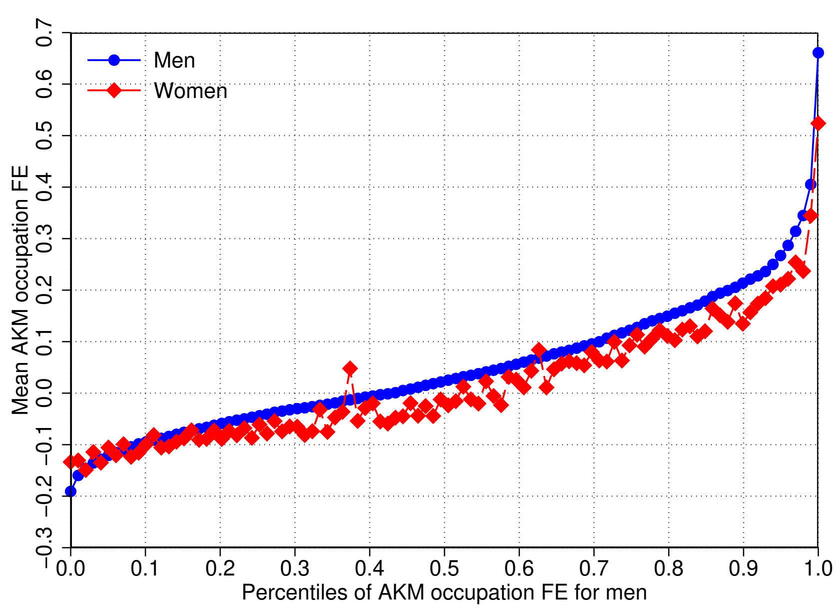

Note: Figure shows predicted AKM occupation FEs separately for men and women based on estimating earnings equation (1). Fixed effects of both genders are sorted by mean FEs of male FE quantiles. *Source:* RAIS, 2007–2014.

Figure B.3. Predicted AKM actual-experience FEs, by gender

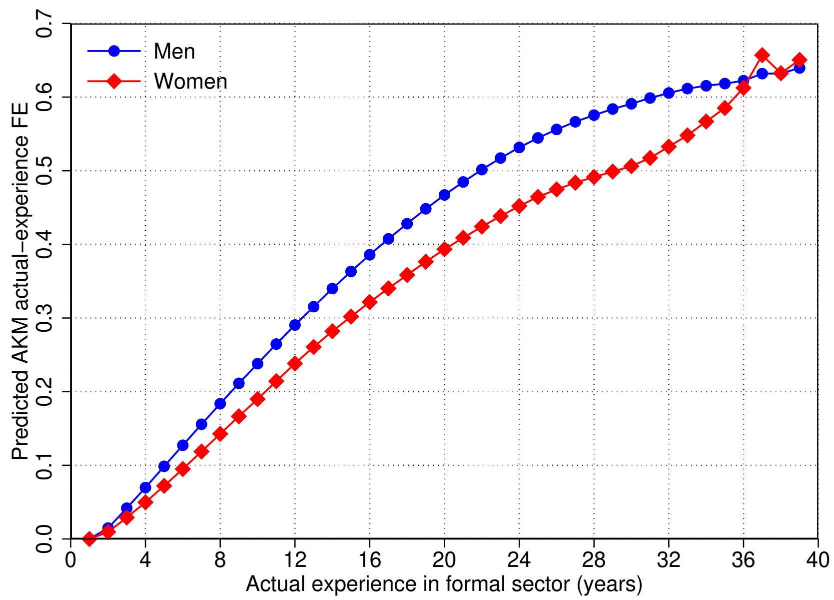

Note: Figure shows predicted AKM actual-experience FEs separately for men and women based on estimating earnings equation (1). *Source:* RAIS, 2007–2014.

Figure B.4. Predicted AKM tenure FEs, by gender

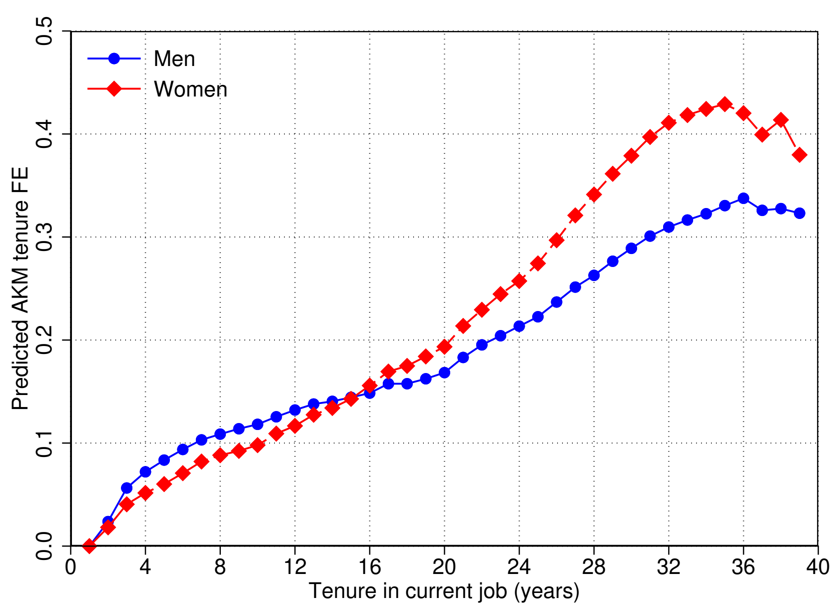

Note: Figure shows predicted AKM tenure FEs separately for men and women based on estimating earnings equation (1). *Source:* RAIS, 2007–2014.

Figure B.5. Predicted AKM education-year FEs, by gender

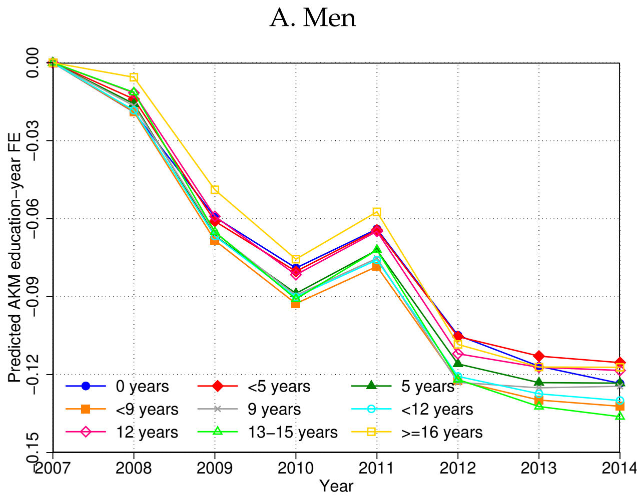

Note: Figure shows predicted AKM education-year FEs separately for men and women based on estimating earnings equation (1). Note that the declining pattern for both genders and all education categories is due to measuring earnings in multiples of the prevailing minimum wage, which increased over this period—see Engbom and Moser (2022) for details. *Source:* RAIS, 2007–2014.

Figure B.6. Predicted AKM education-age FEs, by gender

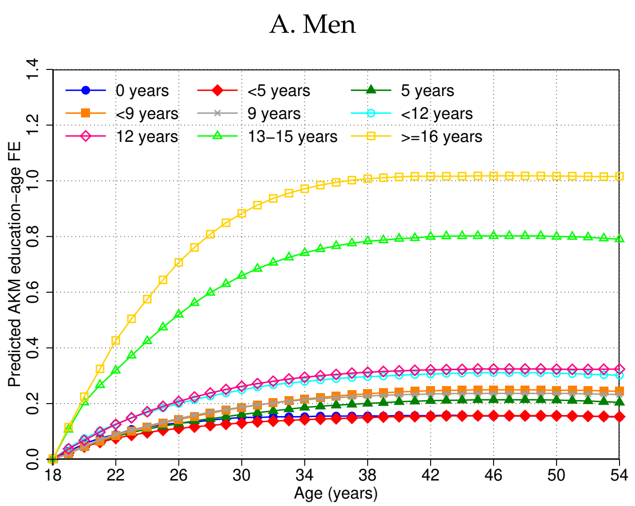

Note: Figure shows predicted AKM education-age FEs separately for men and women based on estimating earnings equation (1). Age-pay profiles for all education groups are constrained to be constant from age 45 to age 49 and unconstrained otherwise. *Source:* RAIS, 2007–2014.

B.2 Further Details on Between vs. Within-Employer Pay Differences

Here, we present two alternative Kitagawa-Oaxaca-Blinder decompositions of the gender gap in pay, only the first of which is shown in the main text. The overall gender gap in employer FEs can be written as

$$\mathbb{E}\left[\psi_{MJ(i,t)}\middle| G(i)=M\right] - \mathbb{E}\left[\psi_{FJ(i,t)}\middle| G(i)=F\right] \tag{B.1}$$

$$= \underbrace{\left(\mathbb{E}\left[\psi_{MJ(i,t)}\middle| G(i)=M\right] - \mathbb{E}\left[\psi_{MJ(i,t)}\middle| G(i)=F\right]\right)}_{\text{between-employer gap}} + \underbrace{\mathbb{E}\left[\psi_{MJ(i,t)} - \psi_{FJ(i,t)}\middle| G(i)=F\right]}_{\text{within-employer gap}} \tag{B.2}$$

$$= \underbrace{\left(\mathbb{E}\left[\psi_{FJ(i,t)}\middle| G(i)=M\right] - \mathbb{E}\left[\psi_{FJ(i,t)}\middle| G(i)=F\right]\right)}_{\text{between-employer gap}} + \underbrace{\mathbb{E}\left[\psi_{MJ(i,t)} - \psi_{FJ(i,t)}\middle| G(i)=M\right]}_{\text{within-employer gap}}. \tag{B.3}$$

Table B.1 shows the results from the two alternative decompositions of the gender log pay gap corresponding to equations (B.2) and (B.3).

Table B.1. Alternative Kitagawa-Oaxaca-Blinder decompositions of the gender log pay gap

		Between-employer gap		Within-employer gap	
	Gender log pay gap	Level	Share (%)	Level	Share (%)
Decomposition 1	0.113	0.089	78.7	0.024	21.3
Decomposition 2	0.113	0.073	64.3	0.040	35.7

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the overall gender log pay gap into a between-employer gap (i.e., pay-policy component) and a within-employer gap (i.e., sorting component). Decomposition 1 corresponds to equation (B.2) and uses men's employer FEs for computing the between-employer component. Decomposition 2 corresponds to equation (B.3) and uses women's employer FEs for computing the between-employer component. *Source:* RAIS, 2007–2014.

To illustrate these two decompositions, Figure B.7 shows the distributions of gender-specific employer FEs underlying the individual terms in equations (B.2) and (B.3).

Figure B.7. Components of Kitagawa-Oaxaca-Blinder decompositions

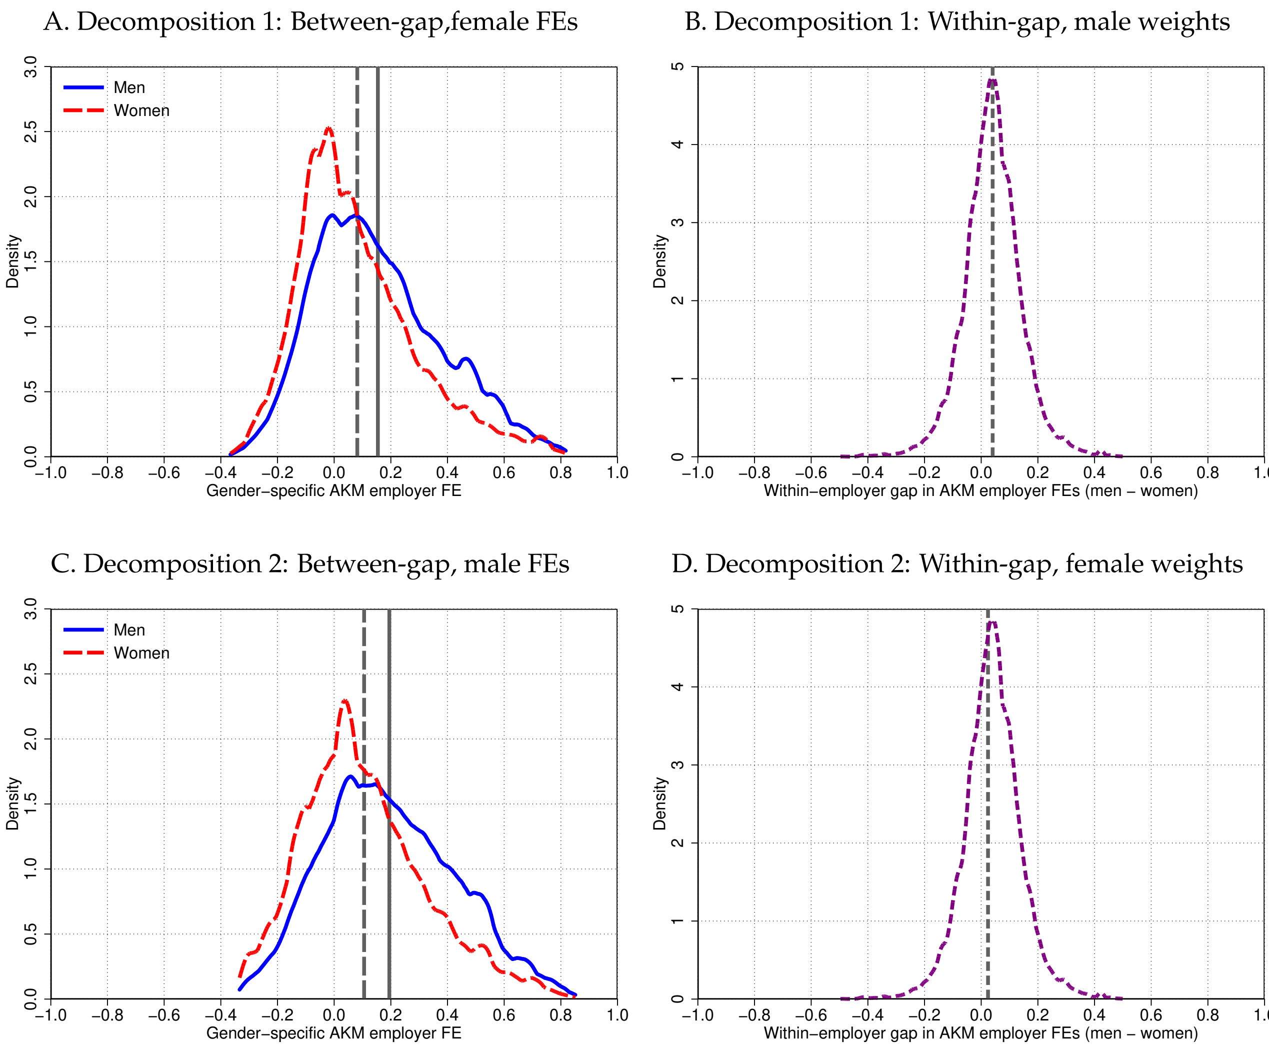

Note: Figure shows pay distributions underlying the Kitagawa-Oaxaca-Blinder decompositions—specifically, the between-gap using female FEs (Panel A) and using male FEs (Panel C). as well as the within-gap using male weights (Panel B) and using female weights (Panel D) Decomposition 1 (Panels A and B) and decomposition 2 (Panels C and D) correspond to equations (B.2) and (B.3) of the main text, respectively. Dashed vertical line shows mean of the distribution. *Source:* RAIS, 2007–2014.

B.3 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. Here, we 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.[58] 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.8 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.

[58] For comparison, the birth rate in Brazil from official birth statistics is 12.9 per 1,000 people (World Bank, 2021a). In our setting, approximately 11.1 workers per 1,000 in the RAIS data go on a registered parental leave during 2014. At face value, this means that our data cover approximately 86.0% of births from vital statistics records, suggesting that there is a high take-up rate of parental leaves among new parents. The difference between the official statistic and that computed on the RAIS data might be due to multiple births per parent (e.g., twins or two separate births within the same calendar year), less than perfect take-up of parental leaves (e.g., continuing office duties in spite of the federally mandated parental leave) or different fertility across workers in formal-sector jobs covered by RAIS (e.g., lower fertility among relatively high-income workers in the formal sector).

Figure B.8. Life-cycle mean gender-specific employer FEs, by gender and parent status

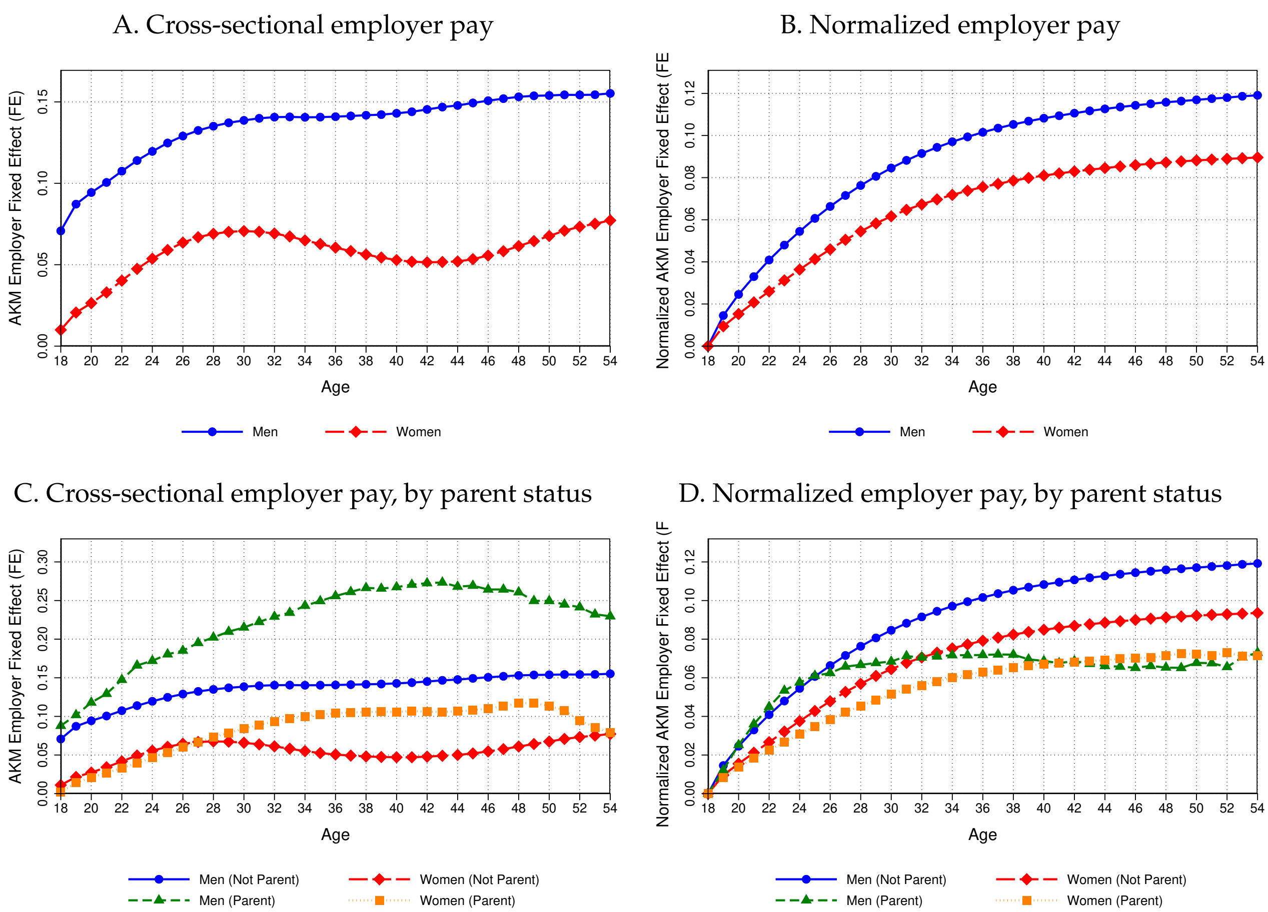

Note: Figure shows cross-sectional (Panels A and C) and normalized (Panels B and D) employer pay across age separately by gender (Panels A and B) and separately by gender and ever-parent status (Panels C and D). Cross-sectional estimates represent binned means. Normalized estimates are binned means of differenced variable, normalized to 0 at age 18. *Source:* RAIS, 2007–2014.

C Equilibrium Model of Employer Pay, Amenities, and Size 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 $\{S_{gz}, W_{gz}\}_{gz}$ and policy functions $\{\phi_{gz}\}_{gz}$; firm value function Π and policy functions $\{w_{gz}, a_{gz}, v_{gz}\}_{gz}$; utility offer distributions $\{F_{gz}(x)\}_{gz}$; measures of nonemployed workers $\{u_{gz}\}_{gz}$, aggregate job searchers $\{U_{gz}\}_{gz}$, aggregate vacancies $\{V_{gz}\}_{gz}$, and labor market tightnesses $\{\theta_{gz}\}_{gz}$; job offer arrival rates $\{\lambda^U_{gz}, \lambda^E_{gz}, \lambda^G_{gz}\}_{gz}$; and firm sizes $\{l_{gz}\}_{gz}$ such that, for all (g, z):*

- *given $F_{gz}(x)$ and $\{\lambda^U_{gz}, \lambda^E_{gz}, \lambda^G_{gz}\}$, workers' value functions S_{gz} and W_{gz} satisfy equations (3) and (4);*
- *nonemployed workers' job acceptance policy follows a threshold rule ϕ_{gz} given by equation (5), and employed workers with flow utility x accept voluntarily any job x' such that $x' > x$;*
- *firm sizes $l_{gz}(\cdot)$ solve equation (16);*
- *given $l_{gz}(\cdot)$, firms' value function $\Pi_{gz}(\cdot)$ satisfies equation (11);*
- *firms choose $\{w_{gz}, a_{gz}, v_{gz}\}$ to maximize their objective given by equation (11);*
- *measures of nonemployed workers are given by equation (6), aggregate job searchers U_{gz} and aggregate vacancies V_{gz} are given by equation (12), and labor market tightness θ_{gz} is given by equation (13);*
- *given θ_{gz}, the job offer arrival rates $\{\lambda^U_{gz}, \lambda^E_{gz}, \lambda^G_{gz}\}$ satisfy equation (14);*
- *given $F_{gz}(x)$, $\{\lambda^U_{gz}, \lambda^E_{gz}\}_{gz}$, λ^G_{gz}, and V_{gz}, firm sizes satisfy equation (16); and*
- *the offer distribution satisfies $F_{gz}(x) = \int_j v_{gz}(j)\mathbf{1}[x_{gz}(j) \le x]\, d\Gamma(j)/V_{gz}$.*

C.2 Proof of Lemma 1 (Optimal Amenities)

Restatement of Lemma 1 (Optimal Amenities). *A firm's optimal amenity policy function $a^*_{gz}(\cdot)$ can be written as $a^*_z(\tilde{c}^{a,0}_g) = (\tilde{c}^{a,0}_g)^{1/(1-\eta^a)} z$. The optimal amenity cost function $c^{a,*}_{gz}(\cdot)$ can be written as $c^{a,*}_z(\tilde{c}^{a,0}_g) = (\tilde{c}^{a,0}_g)^{1/(1-\eta^a)} z/\eta^a$.*

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 (11) as one of choosing in each market a flow utility $x = w + \beta_g a$ and vacancies v that solve the following problem:

$$\max_{x,v} \left\{ \left[(1-\tau_g)pz - c^x_{gz}(x)\right] l_{gz}(x, v) - c^v_{gz}(v) \right\}, \quad \forall (g, z), \tag{C.1}$$

where $c^x_{gz}(x)$ is the solution to the following cost-minimization subproblem in each market (g, z):

$$c^x_{gz}(x) = \min_{w,a} \left\{ w + c^a_{gz}(a) \right\} \quad \text{s.t.} \quad w + \beta_g a = x. \tag{C.2}$$

Once written in this way, it is evident that an interior solution to the firm's cost-minimization problem in equation (C.2) is characterized by the following set of first-order conditions (FOCs):

$$\frac{\partial c^a_{gz}(a^*)}{\partial a} = \beta_g, \tag{C.3}$$

$$w^* = x - \beta_g a^*. \tag{C.4}$$

Equation (C.3) uniquely pins down a firm's optimal amenity choice a^*_{gz} for every market (g, z).

Using the functional form of amenity costs in equation (8), we can rearrange equation (C.3) to get the following expression for the optimal amenity level a^*_{gz} as a function of the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$:

$$a^*_z\left(\tilde{c}^{a,0}_g\right) = \left(\tilde{c}^{a,0}_g\right)^{\frac{1}{1-\eta^a}} z. \tag{C.5}$$

Equation (C.5) shows that a firm's optimal amenity provision is a function of only the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$ and z. Specifically, the optimal amenity policy is linear in worker ability z. Since $\eta^a > 1$, optimal amenities are also decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$. Optimal amenities are invariant to all other parameters. Consequently, a firm with preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$ delivers an amenity valuation of $\beta_g(j)a^*_z(\tilde{c}^{a,0}_g)$ to workers of type (g, z). The optimal wage is then chosen to deliver the remainder of flow utility x according to equation (C.4).[59]

We obtain an explicit expression for the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$ corresponding to a given equilibrium amenity level a^*_{gz} by rearranging equation (C.5) as

$$\tilde{c}^{a,0}_g = \left(\frac{a^*_{gz}}{z}\right)^{1-\eta^a}. \tag{C.6}$$

By plugging equation (C.6) into the amenity cost function from equation (8) and using the expression for optimal amenities in equation (C.5), we can rewrite the equilibrium cost of delivering an amenity valuation $\beta_g a^*$ to the worker as

$$\begin{aligned} c^a_{gz}\left(a^*_{gz}\right) &= c^{a,0}_g \frac{\left(a^*_{gz}/z\right)^{\eta^a}}{\eta^a} z && \text{(C.7)}\\ &= \beta_g \left(\frac{a^*_{gz}}{z}\right)^{1-\eta^a} \frac{\left(a^*_{gz}/z\right)^{\eta^a}}{\eta^a} z && \text{(C.8)}\\ &= \beta_g \frac{a^*_{gz}}{\eta^a} && \text{(C.9)}\\ &= \beta_g \frac{\left(\tilde{c}^{a,0}_g\right)^{\frac{1}{1-\eta^a}} z}{\eta^a}. && \text{(C.10)} \end{aligned}$$

□

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

$$a^{**}_z\left(x, \tilde{c}^{a,0}_g\right) = \begin{cases} x/\beta_g & \text{if } x < \overline{x}\left(\tilde{c}^{a,0}_g\right) \\ a^*_z\left(\tilde{c}^{a,0}_g\right) & \text{if } x \geq \overline{x}\left(\tilde{c}^{a,0}_g\right) \end{cases}, \quad w^{**}_{gz}\left(x, \tilde{c}^{a,0}_g\right) = \begin{cases} 0 & \text{if } x < \overline{x}\left(\tilde{c}^{a,0}_g\right) \\ x - \beta_g a^{**}_z\left(\tilde{c}^{a,0}_g, x\right) & \text{if } x \geq \overline{x}\left(\tilde{c}^{a,0}_g\right) \end{cases},$$

where $\overline{x}\left(\tilde{c}^{a,0}_g\right)$ solves $\partial c^a_{gz}(\overline{x}\left(\tilde{c}^{a,0}_g\right))/\partial a = \beta_g$. Note, however, that in such corner solutions, the optimal wage is $w^{**} = 0$, which is empirically not relevant. Going forward, we focus on the more natural case of an interior solution.

C.3 Proof of Lemma 2 (Optimal Vacancies)

Restatement of Lemma 2 (Optimal Vacancies). *Keeping fixed all other parameters, a firm's optimal vacancy policy $v^*_{gz}(\cdot)$ is strictly increasing in composite productivity $\tilde{p}_{gz}$ and, thus, strictly increasing in productivity p, strictly decreasing in the gender wedge τ for women, and strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$.*

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

$$P_{gz}(\tilde{p}_{gz}, x) = h_{gz}(x) J_{gz}(\tilde{p}_{gz}, x), \tag{C.11}$$

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

$$J_{gz}(\tilde{p}_{gz}, x) = \frac{\tilde{p}_{gz} - x}{\delta_{gz} + \lambda^E_{gz}(1 - F_{gz}(x)) + \lambda^G_{gz}} \tag{C.12}$$

$$= \frac{\left(\tilde{p}_{gz} - x\right) / \left(\delta_{gz} + \lambda^G_{gz}\right)}{1 + \kappa^E_{gz}\left(1 - F_{gz}\left(x\right)\right)}, \tag{C.13}$$

where $\kappa^e_{gz} = \lambda^e_{gz}/(\delta_{gz} + \lambda^G_{gz})$. The acceptance probability for a firm offering x is

$$h_{gz}(x) = \frac{u_{gz} + s^E_{gz}\left(1 - u_{gz}\right) G_{gz}\left(x\right) + s^G_{gz}}{u_{gz} + s^E_{gz}\left(1 - u_{gz}\right) + s^G_{gz}} \tag{C.14}$$

$$= \frac{\delta_{gz} + s^E_{gz}\left(\lambda^U_{gz} + \lambda^G_{gz}\right) G_{gz}\left(x\right) + s^G_{gz}\left(\delta_{gz} + \lambda^U_{gz} + \lambda^G_{gz}\right)}{\delta_{gz} + s^E_{gz}\left(\lambda^U_{gz} + \lambda^G_{gz}\right) + s^G_{gz}\left(\delta_{gz} + \lambda^U_{gz} + \lambda^G_{gz}\right)} \tag{C.15}$$

$$= \frac{1 + s^E_{gz}\kappa^U_{gz} G_{gz}\left(x\right) + s^G_{gz}\left(1 + \kappa^U_{gz}\right)}{1 + s^E_{gz}\kappa^U_{gz} + s^G_{gz}\left(1 + \kappa^U_{gz}\right)} \tag{C.16}$$

$$= \frac{1 + s^E_{gz}\kappa^U_{gz}\left[\frac{F_{gz}(x)}{1+\kappa^E_{gz}\left[1-F_{gz}(x)\right]}\right] + s^G_{gz}\left(1 + \kappa^U_{gz}\right)}{1 + s^E_{gz}\kappa^U_{gz} + s^G_{gz}\left(1 + \kappa^U_{gz}\right)} \tag{C.17}$$

$$= \frac{1 + \kappa^E_{gz}\left[1 - F_{gz}\left(x\right)\right] + s^E_{gz}\kappa^U_{gz} F_{gz}\left(x\right) + s^G_{gz}\left(1 + \kappa^U_{gz}\right)\left[1 + \kappa^E_{gz}\left[1 - F_{gz}\left(x\right)\right]\right]}{\left[1 + s^E_{gz}\kappa^U_{gz} + s^G_{gz}\left(1 + \kappa^U_{gz}\right)\right]\left[1 + \kappa^E_{gz}\left[1 - F_{gz}\left(x\right)\right]\right]}, \tag{C.18}$$

where $\kappa^U_{gz} = (\lambda^U_{gz} + \lambda^G_{gz})/\delta_{gz}$ and u_{gz} is substituted with its expression in equation (6). Combining expressions, expected profits per contacted worker are

$$z\left(\tilde{p}_{gz}, x\right) = h\left(x\right) J\left(\tilde{p}, x\right) \tag{C.19}$$

$$= \frac{\left\{1 + \kappa^E_{gz}\left[1 - F_{gz}\left(x\right)\right] + s^E_{gz}\kappa^U_{gz} F_{gz}\left(x\right) + s^G_{gz}\left(1 + \kappa^U_{gz}\right)\left[1 + \kappa^E_{gz}\left[1 - F_{gz}\left(x\right)\right]\right]\right\}\left(\tilde{p}_{gz} - x\right)}{\left[1 + s^E_{gz}\kappa^U_{gz} + s^G_{gz}\left(1 + \kappa^U_{gz}\right)\right]\left[1 + \kappa^E_{gz}\left(1 - F_{gz}\left(x\right)\right)\right]^2\left(\delta_{gz} + \lambda^G_{gz}\right)}. \tag{C.20}$$

Then, the firm's problem becomes

$$\max_{x,v} \left\{ P_{gz}\left(\tilde{p}_{gz}, x\right) v q_{gz} - c^{v}_{gz}\left(v\right) \right\}, \tag{C.21}$$

where q_{gz} is defined as in equation (14). Therefore, the optimal flow-utility and vacancy policy functions satisfy

$$\begin{aligned} x^{*}_{gz}\left(\tilde{p}_{gz}, \cdot\right) &= \arg\max_{x} P_{gz}\left(\tilde{p}_{gz}, x\right) \\ \frac{\partial c^{v}_{gz}\left(v^{*}\left(\tilde{p}_{gz}, \cdot\right)\right)}{\partial v} &= q_{gz} \max_{x} P_{gz}\left(\tilde{p}_{gz}, x\right). \end{aligned} \tag{C.22}$$

Since the vacancy cost function $c^{v}(\cdot)$ is convex, and $z\left(\tilde{p}_{gz}, x\right)$ in equation (C.20) is strictly increasing in $\tilde{p}_{gz}$, then it follows from an application of the envelope theorem to equation (C.22) that $v^{*}\left(\tilde{p}_{gz}, \cdot\right)$ is strictly increasing in $\tilde{p}_{gz}$. Therefore, $v^{*}_{gz}(\cdot)$ is strictly increasing in productivity p, strictly decreasing (constant) in z_a for women (men). Finally, optimal vacancies are also strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_{g}$ due to the result in Lemma 1. □

C.4 Proof of Lemma 3 (Optimal Flow Utility and Wages)

Restatement of Lemma 3 (Optimal Flow Utility and Wages). *Keeping fixed all other parameters, a firm's optimal flow utility offer $x^{*}_{gz}(\cdot)$ is strictly increasing in composite productivity $\tilde{p}_{gz}$ and, thus, strictly increasing in productivity p for all worker types, strictly decreasing in the gender wedge τ for women, and strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_{g}$. A firm's optimal wage offer $w^{*}_{gz}(\cdot)$ is strictly increasing in productivity p for all worker types and strictly decreasing in the gender wedge τ for women.*

Proof. We proceed in two steps.

Step 1. In the first step, we prove monotonicity of x^{*}_{gz} in components of $\tilde{p}_{gz}$. 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_{gz}:

$$[\partial v_{gz}]: \quad c^{v,0}_{gz} \frac{\partial \tilde{c}^{v}(v_{gz})}{\partial v_{gz}} = T_{gz}(\tilde{p}_{gz} - x_{gz}) \left(\frac{1}{\delta_{gz} + \lambda^{G}_{gz} + \lambda^{E}_{gz}(1 - F_{gz}(x_{gz}))} \right)^{2}, \tag{C.23}$$

$$[\partial x_{gz}]: \quad 1 = (\tilde{p}_{gz} - x_{gz}) \frac{2\lambda^{E}_{gz} f_{gz}(x_{gz})}{\delta_{gz} + \lambda^{G}_{gz} + \lambda^{E}_{gz}(1 - F_{gz}(x_{gz}))}, \tag{C.24}$$

where $T_{gz} = \mu_{gz}[(u_{gz} + s^{G}_{gz})\lambda^{U}_{gz}(\delta_{gz} + \lambda^{G}_{gz} + \lambda^{E}_{gz})]/V_{gz}$. Now consider equation (C.23); because the term on the right-hand side is always positive for $\tilde{p}_{gz} > \phi_{gz}$, it follows that optimal vacancies $v^{*}_{gz}(\tilde{p}_{gz}, c^{v,0}_{gz})$ are always strictly positive.

We now show that the derivative of wages with respect to $\tilde{p}_{gz}$ is always positive. Define $h_{gz}(\tilde{p}_{gz}) =$

$F_{gz}(x^*_{gz}(\tilde{p}_{gz}))$. Thus

$$h_{gz}(\tilde{p}_{gz}) = \frac{\int_{\tilde{p}' \geq \phi_{gz}}^{\tilde{p}_{gz}} v^*_{gz}(\tilde{p}')\gamma_{gz}(\tilde{p}')}{V_{gz}} \, d\,\tilde{p}' \tag{C.25}$$

$$h'_{gz}(\tilde{p}_{gz}) = f_{gz}(x^*_{gz}(\tilde{p}_{gz}))\, {x^*_{gz}}'(\tilde{p}_{gz}) \tag{C.26}$$

$$f_{gz}(x^*_{gz}(\tilde{p}_{gz})) = h'_{gz}(\tilde{p}_{gz}) / {x^*_{gz}}'(\tilde{p}_{gz}), \tag{C.27}$$

where $v^*_{gz}(\tilde{p}_{gz})$ are optimal vacancies conditional on $\tilde{p}_{gz}$, $\gamma_{gz}(\tilde{p}_{gz})$ is the marginal density of composite productivity $\tilde{p}_{gz}$, and $\partial x^*_{gz}(\tilde{p}_{gz})/\partial \tilde{p}_{gz} = {x^*_{gz}}'(\tilde{p}_{gz})$ is the derivative of equilibrium flow utility with respect to $\tilde{p}_{gz}$. Thus, we can rewrite $h'_{gz}(\tilde{p}_{gz}) = v^*_{gz}(\tilde{p}_{gz})/V_{gz}\gamma(\tilde{p}_{gz})$ by differentiating equation (C.27) using Leibniz's integral rule.

Using these identities, we can write $f_{gz}(x^*_{gz}(\tilde{p}_{gz})) = v^*_{gz}(\tilde{p}_{gz})/V_{gz}\gamma_{gz}(\tilde{p}_{gz})\partial\tilde{p}_{gz}/\partial x^*_{gz}(\tilde{p}_{gz})$. Thus, we can rewrite equation (C.24) as

$$\frac{\partial x^*_{gz}(\tilde{p}_{gz})}{\partial \tilde{p}_{gz}} = (\tilde{p}_{gz} - x^*_{gz}) \frac{2\lambda^E_{gz}}{\delta_{gz} + \lambda^G_{gz} + \lambda^E_{gz}(1 - h_{gz}(\tilde{p}_{gz}))} \frac{v^*_{gz}(\tilde{p}_{gz})}{V_{gz}} \gamma_{gz}(\tilde{p}_{gz}). \tag{C.28}$$

Because the right-hand side of this expression is positive for $\tilde{p}_{gz} > \phi_{gz}$, we have $\partial x^*_{gz}(\tilde{p}_{gz})/\partial\tilde{p}_{gz} > 0$, thus proving that equilibrium flow utility is increasing in $\tilde{p}_{gz}$.

Since $\tilde{p}_{gz}$ is strictly increasing in productivity p, strictly decreasing (constant) in the gender wedge τ_g for women (men), and strictly decreasing in the preference-adjusted amenity cost shifter $\tilde{c}^{a,0}_g$ due to Lemma 1, it follows that optimal flow utility is strictly increasing in p, strictly decreasing (constant) in τ_g for women (men), and strictly decreasing in $\tilde{c}^{a,0}_g$.

Step 2. In the second step, we prove the monotonicity of w_{gz} in components of $\tilde{p}_{gz}$. The characterization of $w_{gz} = x_{gz} - \beta_g a_{gz}$ follows from combining Lemmas 1 and 2 with Step 1 above. □

C.5 Alternative Modeling Assumption on Amenity Production

We here present an alternative formulation of firms' amenity production technology. While in the baseline formulation of the model, firms produce gender-specific amenity values for each worker, an alternative formulation has 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 z.

Firms post an amenity vector $\vec{a} = (a_1, a_2, \ldots, a_N) \in \mathbb{R}^{N\times 1}$ subject to cost $c^a(\vec{a})$. Workers derive gender-specific utility from amenities given a preference vector $\vec{\beta}_g = (\beta_{g,1}, \beta_{g,2}, \ldots, \beta_{g,N}) \in \mathbb{R}^{N\times 1}$. A worker of gender g at a firm with amenity vector $\vec{a}$ enjoys amenity utility $\vec{a}'\vec{\beta}_g$.

If $N = 1$, then $\vec{a} = a \in \mathbb{R}$ 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 $sign(\beta_M) = sign(\beta_F)$. Otherwise, if $sign(\beta_M) \neq sign(\beta_F)$, 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:*

$$N \geq 2. \tag{C.29}$$

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

$$\nexists c \in \mathbb{R} \quad s.t. \quad \vec{\beta}_M = c\vec{\beta}_F. \tag{C.30}$$

Under these assumptions, we obtain the following result, 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^a, U_F^a) at a given employer, there exists an amenity vector $\vec{a} = (a_1, a_2, \ldots, a_N)$ such that*

$$\underset{2\times 1}{\begin{bmatrix} U_M^a \\ U_F^a \end{bmatrix}} = \underset{2\times N}{\begin{bmatrix} \vec{\beta}_M' \\ \vec{\beta}_F' \end{bmatrix}} \underset{N\times 1}{\vec{a}}. \tag{C.31}$$

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 $\vec{a}$ in equation (C.31) has full rank. Therefore, the system admits at least one solution. □

If $N = 2$, then Lemma C.1 admits a unique amenity vector $\vec{a} = (a_1, a_2)$ that rationalizes any amenity-utility duplet (U_M^a, U_F^a). If $N > 2$, then there exist multiple amenity vectors $\vec{a}$ that rationalize the same duplet of gender-specific utilities (U_M^a, U_F^a), 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 $\vec{a}$, and the cost of amenity provision is given by*

$$c^a(\vec{a}, l_M, l_F) = \int_{j=0}^{l_M+l_F} \sum_{i=1}^{N} c_i^{a,0} \tilde{c}^a(a_i)\, dj, \tag{C.32}$$

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

Next, we show that there exist (unique) values of productivity p, the gender wedge τ_g; there also exist amenity cost shifters $c_i^{a,0}$ for $i \in \{1, \ldots, N\}$ that rationalize a given vector of amenities $\vec{a}$ 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 $\vec{a}$, there exists a firm-specific amenity cost function $c^a(\vec{a})$ such that $\vec{a}$ solves*

$$(\vec{a}, w_M, w_F, v_M, v_F) = \underset{\tilde{\vec{a}}, \tilde{w}_M, \tilde{w}_F, \tilde{v}_M', \tilde{v}_F}{\arg\max} \left\{ \sum_{g=M,F} \left[(1-\tau_g)p - \tilde{w}_g - c^a(\tilde{\vec{a}})\right] l_g(\tilde{\vec{a}}, \tilde{w}_g, \tilde{v}_g) - \sum_{g=M,F} c^v(\tilde{v}_g) \right\} \tag{C.33}$$

for some levels of firm productivity p and gender wedge τ_g.

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

$$[\partial a_i]: \quad c_i^{a,0}\frac{\partial \tilde{c}^a(a_i)}{\partial a_i} = \frac{[p - w_M - c^a(\vec{a})]\frac{\partial l_M(\vec{a},w_M,v_M)}{\partial a_i} + [(1-\tau)p - w_F - c^a(\vec{a})]\frac{\partial l_F(\vec{a},w_F,v_F)}{\partial a_i}}{l_M(\vec{a},w_g,v_g) + l_F(\vec{a},w_F,v_F)} \tag{C.34}$$

$$[\partial w_M]: \quad 1 = [p - w_M - c^a(\vec{a})]\frac{\frac{\partial l_M(\vec{a},w_M,v_M)}{\partial w_M}}{l_M(\vec{a},w_M,v_M)} \tag{C.35}$$

$$[\partial w_F]: \quad 1 = [(1-\tau)p - w_F - c^a(\vec{a})]\frac{\frac{\partial l_F(\vec{a},w_F,v_F)}{\partial w_F}}{l_F(\vec{a},w_F,v_F)} \tag{C.36}$$

$$[\partial v_M]: \quad \frac{\partial c_M^v(v_M)}{v_M} = [p - w_M - c^a(\vec{a})]\frac{\partial l_M(\vec{a},w_M,v_M)}{\partial v_M} \tag{C.37}$$

$$[\partial v_F]: \quad \frac{\partial c_F^v(v_F)}{v_F} = [p - w_F - c^a(\vec{a}) - z]\frac{\partial l_F(\vec{a},w_F,v_F)}{\partial v_F}. \tag{C.38}$$

Note that the only FOC containing the term $\partial\tilde{c}^a(a_i)/\partial a_i$ is equation (C.34). All other FOCs depend on the level of the amenity cost $c^a(\vec{a})$ but not its derivative. Hence, we can scale the amenity cost function $c^a(\vec{a})$, the productivity level for both genders, or the gender wedge for women to satisfy the wage FOCs in equations (C.35)–(C.36) and vacancy posting in equations (C.37)–(C.38). By Assumption C.3, there are N equations (i.e., FOCs with respect to a_i for $i = 1, 2, \ldots, N$) with N free parameters (i.e., amenity cost shifters $c_i^{a,0}$ for $i = 1, 2, \ldots, N$), so there exists N amenity cost shifters $c_i^{a,0}$ that satisfy firm optimality and rationalize $\vec{a}$. □

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

$$[\partial a_i]: \quad c_i^{a,0} \times \frac{\partial\tilde{c}^a(a_i)}{\partial a_i}\sum_{g=M,F} l_g(x_g, v_g) = \beta_{M,i} l_M(x_M, v_M) + \beta_{F,i} l_F(x_F, v_F), \quad \forall i. \tag{C.39}$$

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 men and women, as demonstrated in Section 3.2. For example, in the data, women comprise 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 $\{\vec{a}_j\}_j$ for each worker j and the cost of amenity provision given by*

$$c^a\left(\{\vec{a}_j\}_j\right) = \int_{j=0}^{l_M+l_F}\sum_{i=1}^{N} c_i^{a,0}\tilde{c}^a(a_{i,j})\,dj, \tag{C.40}$$

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

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 $c_i^{a,0}$ for $i \in \{1, \ldots, N\}$ 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 $a_{i,j}^*(\cdot)$ for all (i, j) is strictly decreasing in its amenity cost shifter $c_i^{a,0}$, increasing in gender utility weights $\beta_{g,i}$ for $g \in \{M, F\}$, 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 $x_g = w_g + \vec{a}'_g \vec{\beta}_g$ and vacancies v_g that solve the following problem:

$$\max_{\{x_g, v_g\}_{g=M,F}} \left\{ \sum_{g=M,F} [(1-\tau_g)p]\, l_g(x_g, v_g) - c^x(x_M, x_F) - \sum_{g=M,F} c_g^v(v_g) \right\}, \quad \forall g, \tag{C.41}$$

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

$$c^x(x_M, x_F) = \min_{w_M, w_F, \{\vec{a}_j\}_j} \left\{ w_M l_M(x_M, v_M) + w_F l_F(x_F, v_F) + \int_{j=0}^{l_M + l_F} \sum_{i=1}^{N} c^a(a_{i,j})\, dj \right\} \tag{C.42}$$
$$\text{s.t.} \quad w_g + \vec{a}'_{g(j)} \vec{\beta}_g = x_g \quad \forall j$$

$$= \min_{\{\vec{a}_g\}_g} \Bigg\{ (x_M - \vec{a}'_M \beta_M)\, l_M(x_M, v_M) + (x_F - \vec{a}'_F \beta_F)\, l_F(x_F, v_F) \tag{C.43}$$
$$+ \sum_{i=1}^{N} [l_M(x_M, v_M)\, c^a(a_{M,i}) + l_F(x_F, v_F)\, c^a(a_{F,i})] \Bigg\},$$

where we move from equation (C.42) to equation (C.43) by using Assumption C.4. Note that the cost of amenity production $c^a(\cdot)$ 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 $\vec{a}_g$ 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:

$$[\partial a_{g,i}] : \quad c_i^{a,0} \times \frac{\partial \tilde{c}^a(a_{g,i})}{\partial a_{g,i}} = \beta_{g,i}, \quad \forall (g, i). \tag{C.44}$$

Equation (C.44) pins down a firm's optimal amenity choice $a_{g,i}^*(c_i^{a,0}, \beta_{g,i})$ as a function of the heterogeneous amenity cost shifter $c_i^{a,0}$ as well as the set of gender-specific amenity-utility weights $\beta_{g,i}$. The optimal wage is then chosen to deliver the remainder of flow utility x_g to workers of gender $g = M, F$. □

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 8. Specifically, it follows from Lemma C.3 that the gender-specific amenity vector $\vec{a}_g^*$ is independent of productivity p, the gender wedge τ, 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.

Choice of Model. 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 $\vec{a}$ 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 a significant share 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 $\vec{a}$. 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 parametrizations of the cost function.

C.6 Further Discussion of Model Assumptions

Output is Linear within Worker Types. 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 by pooling all workers of the same gender. 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.

Further supporting this assumption, Fukui et al. (2023) find a small crowd-out between women and men in the U.S. Therefore, a natural starting point treats men and women as interchangeable inputs in production.

Labor Market Segmentation. Here, we argue that our assumption of labor market segmentation by worker types is reasonable, given the empirical patterns we observe. 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 et al., 2024) and accept jobs subject to deadlines (Cortés et al., 2023). 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. Our model provides a natural way to rationalize these differences.

C.7 Alternative Modeling Assumptions on Vacancy Posting

C.7.1 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_{Mz} of vacancies targeted at male workers and v_{Fz} vacancies targeted at women. The total cost of posting vacancies (v_{Mz}, v_{Fz}) for men and women is given by $c_z^v(v_{Mz} + v_{Fz})$, where the function c_z^v retains the properties laid out in the main text: $c_z^v(0) = 0$, $\partial c_z^v(\cdot)/\partial v > 0$, $\partial^2 c_z^v(\cdot)/\partial v^2 > 0$.

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

$$\max_{x_{Mz}, x_{Fz}, v_{Mz}, v_{Fz}} \left\{ \sum_{g=M,F} (\tilde{p}_{gz} - x_{gz}) l_{gz}(x_{gz}, v_{gz}) - c_z^v(v_{Mz} + v_{Fz}). \right\} \tag{C.45}$$

The FOCs with respect to vacancy posting now read

$$[\partial v_{Mz}]: \quad c_z^{v\prime}(v_{Mz} + v_{Fz}) = T_{Mz}(\tilde{p}_{Mz} - x_{Mz}) \left(\frac{1}{\delta_{Mz} + \lambda_{Mz}^G + \lambda_{Mz}^E(1 - F_{Mz}(x_{Mz}))} \right)^2, \tag{C.46}$$

$$[\partial v_{Fz}]: \quad c_z^{v\prime}(v_{Mz} + v_{Fz}) = T_{Fz}(\tilde{p}_{Fz} - x_{Fz}) \left(\frac{1}{\delta_{Fz} + \lambda_{Fz}^G + \lambda_{Fz}^E(1 - F_{Fz}(x_{Fz}))} \right)^2. \tag{C.47}$$

Putting equations (C.46) and (C.47) 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 FOCs, which do not depend on the amount of vacancies posted by that firm.

Since the right-hand sides in equations (C.46) and (C.47) 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.7.2 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_z of gender-neutral vacancies for workers of each ability level z at cost $c_z^v(v_z)$. In such a model, a firm's problem can be written as

$$\max_{x_{Mz},x_{Fz},v_a} \left\{ \sum_{g=M,F} (\tilde{p}_{gz} - x_{gz}) l_{gz}(x_{gz}, v_a) - c_z^v(v_z) \right\}. \tag{C.48}$$

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 $x_{gz} < \phi_{gz}$ such that no worker of gender g will accept it. Consequently, while a total of V_z vacancies are posted in each submarket in the aggregate, only $V_{gz} \leq V_z = \int v_z(\tilde{p}_{gz})\, d\Gamma_{gz}(\tilde{p}_{gz})$ vacancies are accepted in equilibrium by workers of type (g, z). This implies that the number of matches produced in the labor market is given by

$$m_{gz} = \chi_{gz}[\mu_{gz}(u_{gz} + s_{gz}^E(1 - u_{gz}) + s_{gz}^G)]^\alpha V_z^{1-\alpha} \frac{V_{gz}}{V_z}, \tag{C.49}$$

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, $f_{gz}/q_{gz} = V_z/[u_{gz} + s_{gz}^E(1 - u_{gz}) + s_{gz}^G]$, where $f_{gz} = m_{gz}/[u_{gz} + s_{gz}^E(1 - u_{gz}) + s_{gz}^G]$ is the job-finding rate per effective job searcher and $q_{gz} = m_{gz}/V_z$ is the vacancy yield rate.

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

$$\dot{l}_{gz}(x, v) = -\delta_{gz} l_{gz}(x, v) - s_{gz}\lambda_{gz}^E(1 - F_{gz}(x)) l_{gz}(x, v) + \tag{C.50}$$

$$v q_{gz} \left[\frac{u_{gz} + s_{gz}^G}{u_{gz} + s_{gz}^G + (1 - u_{gz}) s_{gz}^e} + \frac{(1 - u_{gz}) s_{gz}^e}{u_{gz} + s_{gz}^G + (1 - u_{gz}) s_{gz})} G_{gz}(x) \right]. \tag{C.51}$$

Solving for the stationary solution,

$$l_{gz}(x_{gz}, v_z) = \left(\frac{1}{\delta_{gz} + \lambda_{gz}^G + \lambda_{gz}^E(1 - F_{gz}(x_{gz}))} \right)^2 \frac{v_z}{V_z} \mu_{gz}(u_{gz} + s_{gz}^G)\lambda_{gz}^U(\delta_{gz} + \lambda_{gz}^G + \lambda_{gz}^E). \tag{C.52}$$

To find the firm's policy functions, define $T_{gz} = \mu_{gz}[u_{gz}\lambda_{gz}^U(\delta_{gz} + s_{gz}\lambda_{gz}^U)]/V_z$and composite productivity $\tilde{p}_{gz} = (1 - \tau_g)zp + a_{gz} - c_{gz}^a(a_{gz}$. we rewrite the firm's problem as a function of the steady state mass of employed workers as follows:

$$\max_{x_{Mz},x_{Fz},v_z} \left\{ T_{Mz} v_z (\tilde{p}_{Mz} - x_{Mz}) \left(\frac{1}{\delta_{Mz} + \lambda_{Mz}^G + \lambda_{Mz}^E(1 - F_{Mz}(x_{Mz}))} \right)^2 \right. \tag{C.53}$$

$$\left. + \; T_{Fz} v_z (\tilde{p}_{Fz} - x_{Fz}) \left(\frac{1}{\delta_{Fz} + \lambda_{Fz}^G + \lambda_{Fz}^E(1 - F_{Fz}(x_{Fz}))} \right)^2 - c_z(v_z) \right\}. \tag{C.54}$$

The associated FOCs read

$$c'(v_z) = T_{Mz}(\tilde{p}_{Mz} - x_{Mz})\left(\frac{1}{\delta_{Mz} + \lambda^G_{gz} + \lambda^E_{Mz}(1 - F_{Mz}(x_{Mz}))}\right)^2 \tag{C.55}$$

$$+ T_{Fz}(\tilde{p}_{Fz} - x_{Fz})\left(\frac{1}{\delta_{Fz} + \lambda^G_{gz} + \lambda^E_{Fz}(1 - F_{Fz}(x_{Fz}))}\right)^2 \tag{C.56}$$

$$1 = (\tilde{p}_{Mz} - x_{Mz})\frac{2\lambda^E_{Mz} f_{Mz}(x_{Mz})}{\delta_{Mz} + \lambda^G_{Mz} + \lambda^E_{Mz}(1 - F_{Mz}(x_{Mz}))} \tag{C.57}$$

$$1 = (\tilde{p}_{Fz} - x_{Fz})\frac{2\lambda^E_{Fz} f_{Fz}(x_{Fz})}{\delta_{Fz} + \lambda^G_{Fz} + \lambda^E_{Fz}(1 - F_{Fz}(x_{Fz}))}. \tag{C.58}$$

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.52), we can write the female share of a firm as

$$s_f = \frac{l_{Fz}(x_{Fz}, v_z)}{l_{Fz}(x_{Fz}, v_z) + l_{Mz}(x_{Mz}, v_z)} \tag{C.59}$$

$$= \frac{1}{1 + \frac{\left(\frac{1}{\delta_{Mz}+\lambda^G_{Mz}+\lambda^E_{Mz}(1-F_{Mz}(x_{Mz}))}\right)^2 (u_{Mz}+s^G_{Mz})\lambda^U_{Mz}(\delta_{Mz}+\lambda^G_{Mz}+\lambda^E_{Mz})}{\left(\frac{1}{\delta_{Fz}+\lambda^G_{Fz}+\lambda^U_{Fz}(1-F_{Fz}(x_{Fz}))}\right)^2 (u_{Fz}+s^G_{Fz})\lambda^U_{Fz}(\delta_{Fz}+\lambda^G_{Fz}+\lambda^E_{Fz})}}. \tag{C.60}$$

On the right-hand side, we can substitute our empirical estimates of the U-E transition rates λ^u_{gz}, the E-U transition rates δ_{gz}, the compulsory offer arrival rate λ^G_{gz}, the voluntary offer arrival rate λ^e_{gz}, and the compulsory offer search units s^G_{gz} from Table 4. Thus, we obtain:

$$(u_{Mz} + s^G_{Mz})\lambda^U_{Mz}(\delta_{Mz} + \lambda^G_{Mz} + \lambda^E_{Mz}) = (0.236 + 0.101) \times 0.104 \times (0.036 + 0.010 + 0.009) \approx 0.0019 \tag{C.61}$$

$$(u_{Fz} + s^G_{Fz})\lambda^U_{Fz}(\delta_{Fz} + \lambda^G_{Fz} + \lambda^E_{Fz}) = (0.219 + 0.083) \times 0.091 \times (0.028 + 0.007 + 0.007) \approx 0.0012. \tag{C.62}$$

Thus, this ratio is approximately equal to 0.0019/0.0012 = 1.583, and the expression simplifies to

$$s_f = \frac{1}{1 + 1.583 \times \left(\frac{\delta_{Fz}+\lambda^G_{Fz}+\lambda^E_{Fz}(1-F_{Fz}(x_{Fz}))}{\delta_{Mz}+\lambda^G_{Mz}+\lambda^E_{Mz}(1-F_{Mz}(x_{Mz}))}\right)^2} = \frac{1}{1 + 1.583 \times \left(\frac{0.036+0.007\times(1-F_{Fz}(x_{Fz}))}{0.046+0.009\times(1-F_{Mz}(x_{Mz}))}\right)^2}. \tag{C.63}$$

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 $\overline{s}_f$ 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 $F_{Fz} = 1$ and $F_{Mz} = 0$. Conversely, the firm with the lowest female employment share has $F_{Fz} = 0$ and $F_{Mz} = 1$.

Thus, we find that the minimum (maximum) female employment share in the model is ≈ 0.419 (0.596), which is inconsistent with the minimum (maximum) female employment share being close to 0 (1) in the data.

C.8 Comparison with Sorkin (2018)

Table C.1. Comparison with Sorkin (2018)

	Sorkin (2018)	Current paper
Equilibrium concept	Partial equilibrium: does not model underlying sources of firm heterogeneity	General equilibrium: firm wages, amenities, and vacancies are determined in equilibrium
Worker heterogeneity	None	Ability, gender
Firm heterogeneity	Amenity cost parameter, separation shock rate, relocation shock rate	Productivity, gender wedge, gender-specific amenity cost parameters
Idiosyncratic heterogeneity	Type-I extreme value distribution with scale parameter 1	None
Utility function	$V = \omega_1[\ln(w) + \ln(a)]$	$V = \omega_0 + \omega_1[w + a]$
Determination of utility	Exogenous firm values	Endogenous firm values
Firm optimization	Choose pay and amenities to maximize profit s.t. exogenous firm value	Choose pay, amenities, and vacancies s.t. profit maximization
Steady state?	No	Yes
Endogenous job destruction?	Yes	No
Time	discrete	Continuous
Unemployed reject offers?	Yes	No
Reasons to move to lower pay	Amenities, idiosyncratic utility shock, relocation shock	Amenities, relocation shock
Reasons to move to lower rank	Idiosyncratic utility shock, relocation shock	Relocation shock
Compensating differentials	Only "Rosen" motive, cannot identify "Mortensen" motive	Entire joint distribution of (w, a)
Variance of amenities	Identify "Rosen" component but not "Mortensen" component of $Var(\ln(a_j))$	Identify both "Rosen" and "Mortensen" components of $Var(a_j)$
Variance of pay	Decompose Var($\ln w_j$) into rents and compensating differentials	Decompose $Var(\ln w_j)$ into $Var(\ln x_j)$, $Var(\tilde{a}_j)$, and $Cov(x_j, \tilde{a}_j)$
Country	U.S.	Brazil
Data source	Linked employer-employee data (LEHD)	Linked employer-employee data (RAIS)
Data coverage	Employees in 27 U.S. states	All formal-sector employees
Start date	2000Q4	January 2007
End date	2008Q1	December 2014
Period length	29 quarters	84 months
Data frequency	Quarterly	Monthly

Note: This table compares the theoretical framework in Sorkin (2018) with that in the current paper.

C.9 Model with Heterogeneous Separation Rates to Nonemployment

In this subsection, we describe how our model would change if we allowed separation rates to nonemployment to be heterogeneous across firms. In particular, we study a case in which separation rates are a non-decreasing function of composite productivity $\tilde{p}$. We otherwise retain all other assumptions on demographics, market structure, matching function, utility of workers, production, firm heterogeneity, cost functions and utility-posting behaviour.

In this case, we are still able to rank firms by $\tilde{p}$ and therefore we can still write the offer distribution as $F(x) = F(x^*(\tilde{p})$. Thus, with a slight abuse of notation, denote $\delta(x) = \delta(\tilde{p} : x^*(\tilde{p}) = x)$ as the separation rate of a firm with composite productivity $\tilde{p}$ optimally offering flow utility x. [60] There are a few key changes to the equilibrium structure of our model. First of all, The steady-state nonemployment rate for each worker type is now calculated as

$$u_{gz} = \frac{\bar{\bar{\delta}}_{gz}}{\bar{\bar{\delta}}_{gz} + \lambda^U_{gz} + \lambda^G_{gz}}. \tag{C.64}$$

where $\bar{\bar{\delta}}_g = \int_x \delta_g(x)\, dG(x)$ is the average employment-weighted separation rate across firms and $G(x)$ is the cross-sectional distribution of flow utilities of employed workers. Second, the cross-sectional distribution of flow utilities for each worker type is now written as

$$G_{gz}(x) = \frac{F_{gz}(x)}{1 + \kappa^E_{gz}(x)\left[1 - F_{gz}(x)\right]}, \tag{C.65}$$

where $\kappa^E_{gz}(x) \equiv \lambda^E_{gz} / (\int_{x'<x} \delta_g(x')\, dG(x')/G(x) + \lambda^G_{gz})$ governs the effective speed of climbing the firm ladder in utility space, which depends on the average separation rate faced by workers earning utility below x. Third, solving for the stationary employment distribution in each market (g, z), firm sizes are now given by

$$l_{gz}(x, v) = \left(\frac{1}{\delta_g(x) + \lambda^G_{gz} + \lambda^E_{gz}\left[1 - F_{gz}(x)\right]}\right) \frac{v}{V_{gz}} \mu_{gz} \lambda^U_{gz} \left[u_{gz} + (1 - u_{gz}) s^E_g G_{gz}(x) + s^G_g\right] \tag{C.66}$$

where this time we could not simplify $G_{gz}(x)$ due to it including integral terms that are x-dependent. Nevertheless, we can still conclude that firm size is going to be increasing in x due to fewer workers quitting to other firms, and a larger pool of effective searchers moving in from other firms. Also, firm size is decreasing in $\delta_g(x)$.

[60] In a more general model where δ is distributed arbitrarily across firms, workers would rank firms in two dimensions: flow utility x and job security $(1 - \delta)$, leading to a substantially more complicated ladder model which at present we are not able to solve in general equilibrium.

First Order Conditions To solve the firm's problem, we still take FOCs with respect to utility x and vacancies posted v, bearing in mind that $\delta_g(\tilde{p}_{gz})$ is not a choice variable:

$$\frac{d\Pi_{gz}(x,v)}{dv}: \quad c_g^{v,0}\frac{dc_g^v(v)}{dv} = (\tilde{p}_{gz} - x)\frac{1}{V_{gz}}\left(\frac{u_{gz} + (1-u_{gz})s_g^E G_{gz}(x) + s_g^G}{\delta_g(\tilde{p}_{gz}) + \lambda_{gz}^G + \lambda_{gz}^E\left[1 - F_{gz}(x)\right]}\right)\mu_{gz}\lambda_{gz}^U \tag{C.67}$$

$$\begin{aligned}\frac{d\Pi_{gz}(x,v)}{dx}: \quad l(x,v) = {}& (\tilde{p}_{gz} - x)\lambda_g^E f_{gz}(x)\frac{v}{V_{gz}}\left(\frac{u_{gz} + (1-u_{gz})s_g^E G_{gz}(x) + s_g^G}{\left[\delta_g(\tilde{p}_{gz}) + \lambda_{gz}^G + \lambda_{gz}^E\left[1 - F_{gz}(x)\right]\right]^2}\right)\mu_{gz}\lambda_{gz}^U \\ & + (\tilde{p}_{gz} - x)\frac{v}{V_{gz}}\left(\frac{(1-u_{gz})s_g^E g_{gz}(x)}{\delta_g(\tilde{p}_{gz}) + \lambda_{gz}^G + \lambda_{gz}^E\left[1 - F_{gz}(x)\right]}\right)\mu_{gz}\lambda_{gz}^U \end{aligned} \tag{C.68}$$

Equations (C.67) and (C.68) are more general versions of equations (C.23) and (C.23), where we allow $\delta_g(\tilde{p}_{gz}$ to be a decreasing function of $\tilde{p}_{gz}$. Equation (C.68) can be further simplified as:

$$1 = (\tilde{p}_{gz} - x)\left(\frac{\lambda_g^E f_g(x)}{\delta_g(x) + \lambda_{gz}^G + \lambda_{gz}^E\left[1 - F_{gz}(x)\right]} + \frac{(1-u_{gz})s_g^E g_{gz}(x)}{u_{gz} + (1-u_{gz})s_g^E G_{gz}(x) + s_g^G}\right) \tag{C.69}$$

D Identification Appendix

D.1 Overview of Identification Procedure

Figure D.1. Overview of Identification Procedure

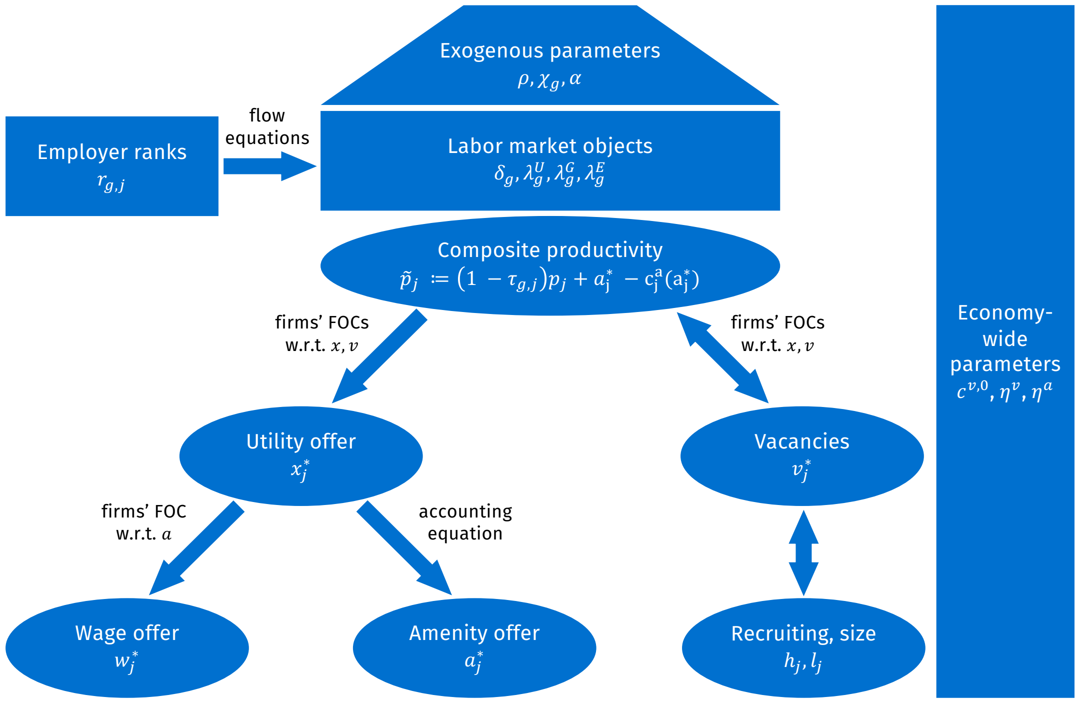

Note: This figure provides a schematic overview of the identification procedure described in Section 5.

D.2 Proof of Proposition 1 (Equilibrium Wage Equation)

Restatement of Proposition 1 (Equilibrium Wage Equation). *The equilibrium wage of a worker of gender g and ability z at a firm with composite productivity $\tilde{p}_g$, preference weight β_g and preference-adjusted amenity cost shifter $\tilde{c}_g^{a,0}$ is*

$$\ln w_{gz}\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right) = \underbrace{\alpha_z}_{\text{“worker wage FE”}} + \underbrace{\psi_g^w\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right)}_{\text{“gender-firm wage FE”}}, \tag{D.1}$$

where

$$\alpha_z = \ln z, \tag{D.2}$$

$$\psi_g^w\left(\tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right) = \ln\left(\tilde{p}_g - \beta_g a_g^*\left(\tilde{c}_g^{a,0}\right) - \int_{\tilde{p}' \geq \phi_g}^{\tilde{p}_g} \left[\frac{1 + \kappa_g^E\left[1 - F_g\left(x_g^*\left(\tilde{p}_g\right)\right)\right]}{1 + \kappa_g^E\left[1 - F_g\left(x_g^*\left(\tilde{p}'\right)\right)\right]}\right]^2 d\tilde{p}'\right). \tag{D.3}$$

Proof. We proceed in two steps. First, we prove the proposition under exogenous firm-level vacancies that are constant within but may differ across genders. Second, we prove that the same result applies under endogenous vacancy posting.

Step 1. Suppose firms differ in their exogenous number of vacancies for each gender, $\{v_g\}_g$. Define $T_{gz} = \mu_{gz}[(u_{gz} + s^G_{gz})\lambda^U_{gz}(\delta_{gz} + \lambda^G_{gz} + \lambda^E_{gz})]/V_{gz}$. First of all, we guess (and later verify) that $\lambda^U_{gz} = \lambda^U_g$ for all ability types z, and therefore also $\lambda^G_{gz} = \lambda^G_g$ and $\lambda^E_{gz} = \lambda^E_g$. Thus, our assumptions also imply that $T_{gz} = T_g$ for all z. Second, under exogenous vacancies, a firm's type is defined by its composite productivity $\tilde{p}_{gz}$ and its exogenous vacancies v_g, which are constant across ability markets. As a consequence, $V_{gz} = V_g$ in all z-markets. Using equation (16), the firm's problem can therefore be written as

$$x^*_{gz}(\tilde{p}_{gz}) = \arg\max_x (\tilde{p}_{gz} - x)\left(\frac{1}{\delta_g + \lambda^G_g + \lambda^E_g(1 - F_{gz}(x))}\right)^2 v_g T_g. \tag{D.4}$$

Thus, given fixed vacancies, equilibrium firm profits in equation (18) can be written as

$$\Pi_{gz}(\tilde{p}_{gz}, v_g) = (\tilde{p}_{gz} - x^*(\tilde{p}_{gz}))\left(\frac{1}{\delta_g + \lambda^G_g + \lambda^E_g(1 - F_{gz}(x^*(\tilde{p}_{gz}))}\right)^2 v_g T_g. \tag{D.5}$$

We can write the offer distribution as

$$F_{gz}(x^*(\tilde{p}_{gz})) = h_{gz}(\tilde{p}_{gz}) = \frac{1}{V_{gz}} \int_{p' > \phi_{gz}}^{\tilde{p}_{gz}} \int v_g \gamma(p', v')\, dp'\, dv', \tag{D.6}$$

where $\gamma_{gz}(p', v')$ is the joint density function of $\tilde{p}_{gz}$ and v_g, and $h_{gz}(\tilde{p}_{gz})$ is the CDF of the marginal distribution of $\tilde{p}_{gz}$, in which values are weighted by vacancies posted by firms with each particular $\tilde{p}_{gz}$. The expression for h_{gz} is equivalent to equation (C.27) in Lemma 3, except that here we are integrating over exogenous rather than endogenous vacancies. Applying the Envelope Theorem yields

$$\frac{\partial \Pi_{gz}(\tilde{p}_{gz}, v_g)}{\partial \tilde{p}_{gz}} = \left(\frac{1}{\delta_g + \lambda^G_g + \lambda^E_g(1 - h_{gz}(\tilde{p})}\right)^2 v_g T_{gz}. \tag{D.7}$$

When $\tilde{p}_{gz} = \phi_{gz}$, $\Pi(\phi_{gz}, v_g) = 0$ for all v_g, which gives us a boundary condition to solve the differential equation for profits. Rearranging (D.5) and integrating equation (D.7) yields

$$x_{gz}\left(\tilde{p}_{gz}\right) = \tilde{p}_{gz} - \int_{y \geq \phi_{gz}}^{\tilde{p}_{gz}} \left[\frac{1 + \kappa^E_g(1 - h_{gz}(\tilde{p}_{gz}))}{1 + \kappa^E_g(1 - h_{gz}(y))}\right]^2 dy, \tag{D.8}$$

where $\kappa^E_g = \lambda^E_g/(\delta_g + \lambda^G_g)$. This equation parallels equation (47) in Burdett and Mortensen (1998), where composite productivity $\tilde{p}_{gz}$ in the current model plays the role of job productivity differentials in their model.

Lemma 1 already proves that amenities are proportional to ability z, therefore we can express them as $a_{gz} = a_g z$, and the cost of producing amenities as $\tilde{c}^a_{gz}(a_{gz}) = z\tilde{c}^a_g(a_g)$.

Summing up, it follows that $a_{gz} = za_g$ and $\tilde{c}^a_{gz}(a_{gz}) = z\tilde{c}^a_g(a_g)$. Therefore, composite productivity $\tilde{p}_{gz} = (1 - \tau_g)pz + \beta_g a_{gz} - \tilde{c}^a_{gz}(a_{gz})$ is proportional to z, and we can write $\tilde{p}_{gz} = z\tilde{p}_g$, where $\tilde{p}_g = (1 - \tau_g)p + \beta_g a_g - \tilde{c}^a_g(a_g)$ is distributed according to $h_g(\tilde{p}_g)$. By definition, $h_{gz}(\tilde{p}_{gz}) = h_{gz}(z\tilde{p}_g) = h_g(\tilde{p}_g)$. Thus, with a change of variables and using that vacancies of each firm are constant across ability markets, we can rewrite equation (D.8) as

$$x_g\left(z, \tilde{p}_g\right) = \tilde{p}_g z - \int_{y \geq \phi_{gz}}^{\tilde{p}_g} z \left[\frac{1 + \kappa^E_g(1 - h_g(\tilde{p}_g)}{1 + \kappa^E_g(1 - h_g(y))}\right]^2 dy. \tag{D.9}$$

We still need to prove that ϕ_{gz} is also proportional to z under the assumption that $b_{gz} = zb_g$. We use a guess-and-verify approach: we guess that the case in which ϕ_{gz} and equilibrium flow utility $x\left(\tilde{p}_g, v_g\right)$ are proportional to z is an equilibrium of the model and we verify it below. From equation (5), we have

$$\phi_{gz} = zb_g + (\lambda_g^U - \lambda_g^E) \int_{x' \geq \phi_{gz}} \frac{1 - F_{gz}(x')}{\rho + \delta_g + \lambda_g^G + \lambda_g^E(1 - F_{gz}(x'))} \, dx' \, . \tag{D.10}$$

We proceed to show that if $\phi_{gz} = z\phi_g$, then $x\left(\tilde{p}_g\right)$ is also proportional to z. The proof follows trivially from equation (D.9) since $\phi_{gz} = z\phi_g$ implies that

$$x_g\left(z, \tilde{p}_g\right) = z\tilde{p}_g - z\int_{y \geq \phi_g}^{\tilde{p}_g} \left[\frac{1 + \kappa_g^E(1 - h_g(\tilde{p}_g))}{1 + \kappa_g^E(1 - h_g(y))}\right]^2 dy \, . \tag{D.11}$$

Next we show that if $x\left(z, \tilde{p}_g\right)$ is proportional to z, then ϕ_{gz} must be proportional to z. Consider the bijective mapping $\tilde{p}_g(x, z) = [x^*(z, \tilde{p}_g)]^{-1}$. We can rewrite the outside option as

$$\phi_{gz} = zb_g + (\lambda_g^U - \lambda_g^E) \int_{x' \geq \phi_{gz}} \frac{1 - h_{gz}([x'(z, \tilde{p}_g)]^{-1})}{\rho + \delta_g + \lambda_g^G + \lambda_g^E(1 - h_{gz}([x'(z, \tilde{p}_g)]^{-1}))} \, dx' \tag{D.12}$$

$$= zb_g + z(\lambda_g^U - \lambda_g^E) \int_{x' \geq \phi_{gz}} \frac{1 - h_g([x'(1, \tilde{p}_g)]^{-1})}{\rho + \delta_g + \lambda_g^G + \lambda_g^E(1 - h_{gz}([x'(1, \tilde{p}_g)]^{-1}))} \, dx' \, , \tag{D.13}$$

which implies that the only solution to this equation satisfies $\phi_{gz} = z\phi_g$.

Finally, recalling that $\tilde{p}_g = (1 - \tau)p + a_g - c_g^a(a_g)$ and that $w = x - \beta a$, we can write monetary wages as

$$w\left(z, \tilde{p}_g, \beta_g, \tilde{c}_g^{a,0}\right) = z\left[\tilde{p}_g - \beta_g a_g(\tilde{c}_g^{a,0}) - \int_{\tilde{p}' \geq \phi_g}^{\tilde{p}_g} \left[\frac{1 + \kappa_g^E(1 - h_g(\tilde{p}_g))}{1 + \kappa_g^E(1 - h_g(\tilde{p}'))}\right]^2 d\tilde{p}'\right] , \tag{D.14}$$

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 $v_{gz}^* = v_g^*$ for all z, so that the offer distribution h_{gz} is the same across all ability markets. We follow a guess-and-verify approach. Suppose that $x_{gz}^*(\tilde{p}_g)$ is proportional to ability z. Using that $F_{gz}(x_{gz}^*(\tilde{p}_{gz})) = h_{gz}(\tilde{p}_{gz})$, we can write the FOC for vacancy creation in equation (C.23) as

$$zc_g^{v,0} \frac{\partial \tilde{c}_g^v(v_{gz})}{\partial v_{gz}} = T_{gz}(\tilde{p}_{gz} - x_{gz}^*(\tilde{p}_{gz})) \left(\frac{1}{\delta_{gz} + \lambda_{gz}^G + \lambda_{gz}^E(1 - h_{gz}(\tilde{p}_{gz}))}\right)^2 \tag{D.15}$$

$$c_g^{v,0} \frac{\partial \tilde{c}_g^v(v_{gz})}{\partial v_{gz}} = T_g(\tilde{p}_g - x_g^*(\tilde{p}_g)) \left(\frac{1}{\delta_g + \lambda_g^G + \lambda_g^E(1 - h_g(\tilde{p}_g))}\right)^2 , \tag{D.16}$$

immediately proving that $v_{gz} = v_g$ for all z. Equation (12) thus implies that aggregate vacancies satisfy $V_{gz} = V_g$, which also implies that in equilibrium, $\lambda_{gz}^U = \lambda_g^U$ (which we previously guessed) and $u_{gz} = u_g$. As a consequence, all terms in the wage equation (D.14) scale linearly in ability. Therefore, log wages take the form of the desired equilibrium wage equation. □

D.3 Equilibrium Amenity Equation

An analogous result shows that equilibrium amenities in this environment also have a log-additive structure, akin to the treatment of wages by Card et al.'s (2016) variant of the AKM framework.

Corollary 1 (Equilibrium Amenity Equation)**.** *The equilibrium amenity of a worker of gender g and ability z at a firm with preference weight β_g and preference-adjusted amenity cost shifter $\tilde{c}_g^{a,0}$ is*

$$\ln \beta_g a_{gz}\left(\tilde{c}_g^{a,0}\right) = \underbrace{\alpha_z}_{\text{"worker amenity FE"}} + \underbrace{\psi^a\left(\beta_g, \tilde{c}_g^{a,0}\right)}_{\text{"gender-firm amenity FE"}}, \tag{D.17}$$

where

$$\alpha_z = \ln z, \tag{D.18}$$

$$\tilde{c}_g^{a,0} = \frac{c_g^{a,0}}{\beta_g} \tag{D.19}$$

$$\psi^a\left(\beta_g, \tilde{c}_g^{a,0}\right) = \ln \beta_g + \frac{1}{1-\eta^a} \ln \tilde{c}_g^{a,0}. \tag{D.20}$$

Proof. This result follows directly from equation (C.5) in the proof of Lemma 1 in Section C.2. □

Corollary 1 shows that amenities, like wages, follow a specification that is log-additive between worker heterogeneity (α_z) and gender-specific firm heterogeneity (ψ^a). The worker amenity FE α_z is a strictly monotonic transformation of worker ability. The gender-firm amenity FE $\psi^a(\beta_g, \tilde{c}_g^{a,0})$ depends only on a firm's preference weight β_g and amenity cost shifter $c_g^{a,0}$ for each gender, scaled by a function of the economy-wide amenity cost elasticity η^a. However, an explicit treatment of amenities and compensating differentials was missing from the analysis in Card et al. (2016). Our framework fills this gap by explicitly modeling firms' equilibrium wage and amenity choices.

D.4 Normalization of Gender-Specific Firm Pay

General Model with Endogenous Amenities. Here, we discuss the normalization of gender-specific firm pay in our baseline model with endogenous amenities. Recall that we estimate gender-specific firm pay components by applying a two-way fixed effect model à la AKM separately for each gender:

$$\ln(w_{ijt}) = \alpha_i + \psi_{G(i)j} + X_{it}\beta_{G(i)} + \varepsilon_{ijt}. \tag{D.21}$$

From here on, with a slight abuse of notation and to be consistent with the notation in our structural model, we write $w_{gr} \equiv \exp(\psi_{gJ_g(r)})$ to denote the level pay of firm $j = J_g(r)$ at rank $r \in [0,1]$ for gender g. As explained in Abowd et al. (2002) and Card et al. (2016), equation (D.21) identifies the gender-specific firm fixed effects up to a constant within each connected set—i.e., one normalization must be imposed on a reference firm for each connected set. Because the connected sets for men and women are disconnected by construction, comparing firm pay across genders requires an additional normalization on the level of firm fixed effects for men compared to women. The baseline assumption in Card et al. (2016) is that firms in a lower range of the value added per worker distribution have firm fixed effects equal to zero for both genders. Card et al. (2016) show that similar results are obtained when imposing an alternative normalization that sets mean firm fixed effects equal to zero in the hotel and restaurant sector for both genders.

Through the lens of our equilibrium model, any normalization of gender-specific firm fixed effects must take into account the fact that amenities differ between firms for a given gender as well as within a firm across genders, in addition to the existence of firm-specific gender wedges. Thus, we derive a model-consistent normalization to the gender-specific estimates of fixed effects in equation (D.21).

Building on the equilibrium wage equation (19), we can derive a model-consistent normalization of gender-firm FEs that allows us to compare firm pay between men and women. Intuitively, our model suggests that such a normalization requires finding a set of firms j with the same $\tilde{p}_g$, β_g and $\tilde{c}_g^{a,0}$ across genders $g \in \{M, F\}$ so that $\psi_M(\tilde{p}_M, \beta_M, \tilde{c}_M^{a,0}) = \psi_F(\tilde{p}_F, \beta_F, c_F^{a,0})$. An advantage of our model is that it provides us guidance on how to find such a set of firms. To this end, let $\mathcal{B}_g \equiv [0, \hat{r}_g] \subseteq [0,1]$ for some strictly positive but small $\hat{r}_g \approx 0$ be a set of firms with utility ranks near the bottom for gender g. Let $\mathcal{D}_g \subseteq [0,1]$ be a set of firms with $\tau_{Fr} \approx 0$. Let $\mathcal{A}_g \subseteq [0,1]$ be a set of dual-gender firms with $\beta_{gr} a_{gr} \approx \beta_{-gR_{-g}(J(r))} a_{-gR_{-g}(J(r))}$. Then the following result provides us a model-consistent normalization of gender-firm FEs:

Proposition 6 (Firm Pay Normalization)**.** *Firm pay ψ_{gj} can be equated across genders $g \in \{M, F\}$ for a set of firms j with rank $R_g(j) \in \mathcal{B}_g \cap \mathcal{D}_g \cap \mathcal{A}_g$ simultaneously for both genders g:*

$$\mathbb{E}_j[\psi_{Mj}] = \mathbb{E}_j[\psi_{Fj}] \quad \forall j: \ R_g(j) \in \mathcal{B}_g \cap \mathcal{D}_g \cap \mathcal{A}_g, \forall g \in \{M, F\}. \tag{D.22}$$

Proof. Recall the definition of composite productivity of a firm at rank r for gender g,

$$\tilde{p}_{gr} \equiv (1 - \tau_{gr}) p_r + \beta_{gr} a_{gr} - \tilde{c}_{gr}^a(a_{gr}). \tag{D.23}$$

Our model-consistent normalization applies to the intersection of three sets of firms.

First, we consider a set of close-to-zero profit firms. Let $\mathcal{B}_g \equiv [0, \hat{r}_g] \subseteq [0,1]$ for some strictly positive but small $\hat{r}_g \approx 0$ be a set of firms with utility ranks near the bottom for gender g. Bottom-ranked firms with $r \in \mathcal{B}_g$ provide worker utility $x_{gr} \equiv w_{gr} + \beta_{gr} a_{gr}$ approximately equal to their composite productivity $\tilde{p}_{gr}$ under the assumption that the lower bound of the support of (p, τ_g) extends low enough, which can always be guaranteed:

$$\tilde{p}_{gr} \approx x_{gr} \qquad \forall r \in \mathcal{B}_g. \tag{D.24}$$

Second, we consider a set of firms that treat workers of both genders interchangeably in production. Let $\mathcal{D}_g \subseteq [0,1]$ be a set of firms with $\tau_{Fr} \approx 0$. Also, recall that $\tau_{Mr} = 0$ by assumption. For those firms, we have

$$\tilde{p}_{gr} \approx p_r + \beta_{gr} a_{gr} - \tilde{c}_g^a(a_{gr}) \qquad \forall r \in \mathcal{D}_g, \tag{D.25}$$

so the only gender-specific aspect of firms $r \in \mathcal{D}_g$ is $\beta_{gr} a_{gr} - \tilde{c}_g^a(a_{gr})$. Combining equations (D.24) and (D.25), we have $p_r + \beta_{gr} a_{gr} - c_g^a(a_{gr}) \approx w_{gr} + \beta_{gr} a_{gr}$ for $r \in \mathcal{B}_g \cap \mathcal{D}$ for both genders.

Third, we consider a set of firms that provide the same level of amenities to both genders. Let $\mathcal{A}_g \subseteq [0,1]$ be a set of dual-gender firms with $\beta_{gr} a_{gr} \approx \beta_{-gR_{-g}(J(r))} a_{-gR_{-g}(J(r))}$. In words, the amenities enjoyed by gender g are enjoyed to the same level by gender $-g$ (e.g., both gender-specific amenities equal zero), where $-g \equiv \{M, F\} \setminus \{g\}$.

Due to our model result that a firm's optimal amenity choice a^* is strictly decreasing in the preference-adjusted amenity cost parameter $\tilde{c}^{a,0}$, a firm optimally provides the same level of amenities to both genders if and only if it faces the same cost function / preference combination for both genders. As a result,

$$\beta_{gr} a_{gr} - \tilde{c}_g^a(a_{gr}) = \beta_{-gR_{-g}(J(r))} a_{-gR_{-g}(J(r))} - \tilde{c}_{-g}^a(a_{-gR_{-g}(J(r))}) \qquad \forall r \in \mathcal{A}_g, \tag{D.26}$$

so the only gender-specific aspect of firms $r \in \mathcal{A}_g$ is τ_{gr}.

Combining the above insights, the definition of $\tilde{p}_{gr}$ in equation (D.23) yields that $w_{MR_M(j)} = w_{FR_F(j)}$

for any firm j with rank $R_g(j) \in \mathcal{B}_g \cap \mathcal{D}_g \cap \mathcal{A}_g$ for both genders g at the same time. In words, for a subset of firms that simultaneously (i) are located near the bottom of both genders' utility ranks, (ii) treat men and women as perfect substitutes in production, and (iii) provide the same amenity level to men and women, we can equalize the gender-specific AKM firm fixed effects in equation (D.21) across men and women according to equation (D.22). □

In words, Proposition 6 states that we can equalize the gender-specific AKM firm fixed effects in equation (19) across men and women for a subset of firms that simultaneously (i) are located near the bottom of both genders' utility ranks, (ii) treat men and women as perfect substitutes in production, and (iii) provide a fixed amenity level to men and women.

The normalization in equation (D.22) embeds two untestable assumptions, namely that (i) men and women are treated as perfect substitutes (i.e., $\tau_{Fr} \approx 0$) at firms with rank $r \in \mathcal{D}_F$ for women, and (ii) the same amenity value is provided to both men and women (i.e., $\beta_{MR_M(j)} a_{MR_M(j)} \approx \beta_{FR_F(j)} a_{FR_F(j)}$) at firms j such that $R_g(j) \in \mathcal{A}_g$ for both genders g. Note that the fact that we are looking for firms near the bottom of both genders' utility ranks (i.e. $R_g(j) \in \mathcal{B}_g$ for both genders g) is a verifiable condition given our revealed-preference measures of gender-specific firm ranks.

To summarize, while the original normalization of gender-firm FEs based on value added per worker proposed by Card et al. (2016) is not valid in our environment, equation (D.22) provides a normalization that extends the argument in Card et al. (2016) to our environment with gender-specific amenities and compensating differentials.

Special Case of the Model with Exogenous Amenities. Now suppose that gender-specific firm amenities a_{gr} are exogenous. Then $c_g^a(a_{gr}) = 0$ without loss of generality. In this case, equation (D.24) reduces to

$$(1 - \tau_{gr}) p_r - w_{gr} \approx 0 \qquad \forall r \in \mathcal{B}. \tag{D.27}$$

Therefore, if we can find a subset of firms in $\mathcal{B} \cap \mathcal{D}$ that simultaneously (i) are located near the bottom of the utility ranks for both genders (i.e., $r \in \mathcal{B}$) and (ii) treat men and women as perfect substitutes in production net of employers' taste for gender (i.e., $r \in \mathcal{D}$), then equalizing firm pay across genders for this set of firms constitutes a normalization of the gender-specific AKM firm fixed effects in equation (D.21) that is consistent with our model featuring exogenous amenities:

$$\mathbb{E}_r[\psi_{Fr}] = \mathbb{E}_r[\psi_{Mr}] \qquad \forall r \in \mathcal{B} \cap \mathcal{D}. \tag{D.28}$$

The normalization in equation (D.28) embeds only one untestable assumption, namely that men and women are treated as perfect substitutes (i.e., $\tau_r \approx 0$) at firms with rank $r \in \mathcal{D}$. Note that the same normalization as imposed in the case with endogenous amenities, discussed in Section D.4 is still valid in this case.

D.5 Proof of Proposition 2 (Employer Ranks)

Restatement of Proposition 2 (Employer Ranks) *All workers of the same gender share a common employer ranking $r \in [0, 1]$, which is identified given employer sizes $l(r)$.*

Proof. First, note that Propositions 1 and Corollary 1 together imply that the flow consumption that a worker of gender g and ability z receives at firm j can be written as

$$x_{gz}(j) = w_{gz}(j) + \beta_g(j)a_{gz}(j) \tag{D.29}$$

$$= \exp\left(\alpha_z + \psi_g^w(j)\right) + \exp\left(\alpha_z + \psi^a(j)\right) \tag{D.30}$$

$$= \exp(\alpha_z)\exp\left(\psi_g^w(j)\right) + \exp(\alpha_z)\exp\left(\psi^a(j)\right) \tag{D.31}$$

$$= \exp(\alpha_z)\left[\exp\left(\psi_g^w(j)\right) + \exp\left(\psi^a(j)\right)\right] \tag{D.32}$$

$$= z\left[\exp\left(\psi_g^w(j)\right) + \exp\left(\psi^a(j)\right)\right] \tag{D.33}$$

Fixing gender g, equation (D.33) is linear in worker ability z. In particular, the firm-specific term, $\exp(\psi_g^w(j)) + \exp(\psi^a(j))$, is independent of worker ability z. Therefore, in the eyes of a worker of type (g, z), a firm j is higher-ranked than another firm j' if and only if it provides higher flow utility $x_{gz}(j) > x_{gz}(j')$ if and only if $\exp(\psi_g^w(j)) + \exp(\psi^a(j)) > \exp(\psi_g^w(j')) + \exp(\psi^a(j'))$. This proves that all workers of the same gender share a common ranking of all firms in the economy. Therefore, for the remainder of the proof we drop the ability subscript z, and we focus on one gender at a time, thus dropping the gender subscript g too for readability.

As we have proved in Lemma 3 that utility $x(\tilde{p})$ is increasing in composite productivity $\tilde{p}$, and that higher utility x means that an employer is higher-ranked, it follows that employers with higher $\tilde{p}$ are higher-ranked. Therefore, higher-ranked employers:

1. Post higher utility: $x(r') > x(r)$ whenever $r' > r$, from Lemma 3;
2. Retain more workers: by definition $F(x(r')) > F(x(r))$ whenever $x(r') > x(r)$;
3. Post more vacancies: from Lemma 2, $v(r') > v(r)$ whenever $r' > r$;
4. Have larger size: as a result of posting more vacancies and offering higher utility, $l(r') > l(r)$ whenever $r' > r$.

These prove trivially that firm sizes $l(r)$ are monotonically increasing in a firm's rank, and therefore we can recover firm ranks r by ordering firms by size, proving the proposition.

□

D.6 Proof of Proposition 3 (Labor Market Objects)

Restatement of Proposition 3 (Labor Market Objects) *Gender-firm-specific recruiting intensities $f(r)$ and vacancies $v(r)$ as well as gender-specific separation hazards δ, job offer hazards from nonemployment λ^U, involuntary job offer hazards λ^G, voluntary on-the-job offer hazards λ^E, and aggregate vacancies V are identified given employer ranks and data on worker flows between employment states.*

Proof. We proceed in steps, with each step linking one empirical object to one model object (i.e., either a model parameter or an equilibrium outcome of interest).

<u>Empirical Firm Nonemployment Hiring Shares $\leftrightarrow$ Model Recruiting Intensities.</u> We obtain the empirical gender-specific distribution of recruiting intensities $f(r)$ by inverting the gender-specific distribution of employment G_g across ranks by solving equation (7) for the offer distribution F_g. We then obtain the hiring distribution as the change in F_g across ranks.

This empirical object directly corresponds to firms' recruiting intensities v_r/V in the model, where v_r denotes the vacancies posted by firm r and $V \equiv \int_r v_r dr$ denotes the total number of vacancies in the economy.[61]

Empirical Rate of Moving into Nonemployment ↔ Model Rate of Exogenous Separations. We identify δ off rates of workers i moving into nonemployment:

$$\hat{\delta} = \mathbb{E}_i \mathbf{1}\left[\text{nonemployed}_{i,t+1} \middle| \text{employed}_{i,t}\right]. \tag{D.34}$$

Empirical Job Finding Rate ↔ Model Rate of Job Offers from Nonemployment. A simple log-hazard model of worker E-N-E transitions, where E denotes employment at some firm and N denotes nonemployment, can be used to recover $\widehat{\lambda^U + \lambda^G}$ separately by gender.[62]

Empirical Share of Moves Down the Firm Ranks ↔ Model Rate of Involuntary Job Offer Shocks. Two insights allow us to use information on worker transitions between employers to identify λ^G. First, we focus on transitions in rank space, not pay space. Second, the share of rank-increasing transitions due to involuntary on-the-job offers declines in F_r. Formally, the total number of job-to-job transitions from employer rank r is

$$J2J_r = l_r[\lambda^E(1 - F^r) + \lambda^G], \tag{D.35}$$

where l_r is the gender-specific number of workers at firm r. Rearranging and averaging across all firms, we have

$$\hat{\lambda}^G = \mathbb{E}_r\left[\frac{J2J_{r\downarrow}}{l_r F_r}\right], \tag{D.36}$$

where $J2J_{r\downarrow} = J2J_r - l_r(\lambda^E + \lambda^G)(1 - F_r)$ is the number of job-to-job transitions to lower ranks. Based on this, we derive the parameter estimate $\hat{s}^G \equiv \hat{\lambda}^G / \hat{\lambda}^U$.

Empirical Share of Moves up the Firm Ranks ↔ Model Rate of Voluntary On-the-Job Offers. On-the-job offers not associated with involuntary transitions must have been voluntary. Hence, once we know $\hat{\lambda}^G$, we can use equation (D.35) to estimate λ^E as

$$\hat{\lambda}^E = \frac{J2J_r/n_r - \hat{\lambda}^G}{1 - F_r}. \tag{D.37}$$

Notice that all of these parameters are over-identified, as in principle we could use just a fraction of the firms and of the job-to-job moves to identify them. We choose to use the overall sample average of these two moments. Based on this, we derive the parameter estimate $\hat{s}^E \equiv \hat{\lambda}^E / \hat{\lambda}^U$.

Aggregate Vacancies. Given estimates of δ, λ^U, λ^E, and λ^G, in addition to the relative mass of workers of a given gender, μ, we are equipped to deduce aggregate vacancies V. To this end, recall that a worker's job-finding probability due to the aggregate matching function is

$$\frac{m}{u} = \lambda^U + \lambda^G = \chi\left[\mu\left(u + s^E(1-u) + s^G\right)\right]^{\alpha} V^{1-\alpha}, \tag{D.38}$$

where m is the number of matches and $u = \delta/(\delta + \lambda^U + \lambda^G)$ is the nonemployment rate. Given the

[61] Since $f_r = v_r/V$ refers to shares (i.e., not levels), the mass of aggregate vacancies V does not matter for its computation. At the end of this section, we spell out how to deduce the mass of aggregate vacancies V, which will be useful later on.

[62] Since in our model, the job-finding rate from nonemployment is $\lambda^U + \lambda^G$, we use our estimate $\hat{\lambda}^G$ below to obtain $\hat{\lambda}^U$.

normalization $\chi = 1$, we can solve equation (D.38) for aggregate vacancies V:[63]

$$V = \left(\frac{\lambda^U + \lambda^G}{\left[\mu \left(u + s^E(1-u) + s^G\right)\right]^{\alpha}} \right)^{\frac{1}{1-\alpha}}. \tag{D.39}$$

□

D.7 Proof of Proposition 4 (Firm-Level Parameters)

Restatement of Proposition 4 (Firm-Level Parameters) *The following gender-firm-specific parameters as functions of r are point identified given employer ranks and labor market objects: productivity $p(r)$, the gender wedge $\tau(r)$, the amenity valuation $\beta(r)a^*(r)$, and amenity costs $c^a(a^*(r))$.*

Proof. Here, we present an identification result based on continuous firm types, as in the model developed in Section 4 and further discussed in Section 5. In Appendix D.8, we extend this argument to the case of discrete firm types, as in the data.

First, by the results of Proposition 1 and Corollary 1, we abstract from heterogeneity in ability and replace $z = 1$ in all functional forms without loss of generality. Therefore, the vacancy cost function can be written as $c^v(v) = c^{v,0} v^{\eta^v} / \eta^v$ and the amenity cost function can be written as $c^a(a) = c^{a,0} a^{\eta^a} / \eta^a$.

Recall that the composite productivity of firm r is given by $\tilde{p}(r) = (1 - \tau(r))p(r) + \beta(r)a(r) - c^a(a(r); r)$. Let

$$T \equiv \frac{\mu[(u + s^G)\lambda^u(\delta + \lambda^G + \lambda^E)]}{V} \tag{D.40}$$

be a constant which only depends on previously estimated labor market rates and aggregate vacancies. We start from the firm's FOCs for optimality in equations (C.23) and (C.24). We substitute the functional form of $c^v(v)$ in equation (C.23) and express the FOCs as a coupled pair of differential equations:

$$h'(\tilde{p}(r)) = \frac{1}{V} \left[\frac{T(\tilde{p}(r) - x(\tilde{p}(r)))}{c^{v,0}\left[\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}(r)))\right]^2} \right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}(r)), \tag{D.41}$$

$$x'(\tilde{p}(r)) = \frac{1}{V} \frac{2\lambda^E(\tilde{p}(r) - x(\tilde{p}(r)))}{\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}(r)))} \left[\frac{T(\tilde{p}(r) - x(\tilde{p}(r)))}{c^{v,0}\left[\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}(r)))\right]^2} \right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}(r)). \tag{D.42}$$

We first derived the differential equation (D.42) as equation (C.28) in Appendix C.4, and here we substituted therein the explicit solution for vacancies in equation (D.41). In equations (D.41)–(D.42) above, $h(\tilde{p}(r)) \equiv F(x(\tilde{p}(r))) = F(r)$ is the CDF of *composite productivity offers* (i.e., weighted by firms' vacancies) and $h'(\tilde{p}(r)) = \partial h(\tilde{p}(r))/\partial \tilde{p}$ is the derivative thereof, while $\Gamma(\tilde{p}(r))$ and $\gamma(\tilde{p}(r)) \equiv \Gamma'(\tilde{p}(r))$ are the CDF of *composite productivity* and its derivative. Note that we can obtain $h(\tilde{p}(r))$ by integrating the hiring density $f(r)$ over observed ranks r, since $F(r) = \int_{r'=0}^{r} f(r')\, dr'$.

Next, we perform a change of variables using the fact that, by definition,

$$\Gamma(\tilde{p}(r)) = \int_{\tilde{p}'=\underline{\tilde{p}}}^{\tilde{p}(r)} \gamma(\tilde{p}')\, d\tilde{p}' = \int_{r'=0}^{r} 1\, dr' = r, \tag{D.43}$$

That is, the share of firms with composite productivity up to $\tilde{p}(r)$ equals the share of firms with rank

[63] As discussed in the main text, this normalization is inconsequential for our purposes.

up to r. Therefore, $\gamma(\tilde{p}(r))\, d\tilde{p}/\, dr = 1\, \forall r$. Equivalently, $h'(\tilde{p}(r))\, d\tilde{p}/\, dr = f(r)$. Using the definition that $F(r) = h(\tilde{p}(r))$, we can rewrite equation (D.41) to obtain an expression for the firm's *flow payoff per matched worker*,

$$P(r) \equiv \tilde{p}(r) - x(\tilde{p}(r)) = [f(r)V]^{\eta^v - 1} \frac{c^{v,0}}{T} \left[\delta + \lambda^G + \lambda^E(1 - F(r))\right]^2, \tag{D.44}$$

In equation (D.44), $f(r)$, $c^{v,0}$, and $F(r)$ are all known quantities identified by applying Proposition 3 to the data. Thus, equation (D.44) can be rearranged to yield an explicit expression for $P(r)$ as a function of known objects. Finally, we perform a similar change of variables to write the derivative of utility $x'(r)$ as a function of ranks:

$$x'(r) = \frac{1}{V} \frac{2\lambda^E P(r)}{\delta + \lambda^G + \lambda^E(1 - F(r))} \left[\frac{TP(r)}{c^{v,0}\left[\delta + \lambda^G + \lambda^E(1 - F(r))\right]^2}\right]^{\frac{1}{\eta^v - 1}}. \tag{D.45}$$

Plugging $P(r)$ from equation (D.44) into this expression allows us to identify flow utility across ranks, $x(r) = K + \int_{r'=0}^{r} x'(r')\, dr'$, up to a constant of integration K. Intuitively, the model helps us pin down differences in utilities between rungs of the job ladder. We choose a value for the constant of integration K such that the equilibrium distribution of amenities—derived in equation (D.46) below—attains a lower bound strictly above but arbitrarily close to zero.[64]

This allows us to identify gender-firm-specific amenity valuations, $\beta(r)a(r)$, by simply taking the difference between utility and wages at each firm, and to identify productivity net of the gender wedge, $(1 - \tau(r))p(r)$, by adding wages and the cost of amenities to the firm's flow payoff per matched worker:

$$\beta(r)a(r) = x(r) - w(r), \tag{D.46}$$

$$(1 - \tau(r))p(r) = P(r) + w(r) + c^a(a(r); r). \tag{D.47}$$

On the right-hand side of equation (D.46), $x(r)$ is known from integrating equation (D.45) and $w(r)$ is known from the data, so we can infer $\beta(r)a(r)$. On the right-hand side of equation (D.47), $P(r)$ is known from equation (D.44), $w(r)$ is known from the data, and $c^a(a(r); r) = \beta(r)a(r)/\eta^a$ is known from equation (C.9) of Lemma 1, with the known value for $\beta(r)a(r)$ substituted from equation (D.46).

If, in addition, we fix the preference weights $\beta(r)$ at a known value, say $\beta(r) = 1$, then we can obtain gender- and firm-specific preference-adjusted amenity cost shifters by inverting the firm's FOC for optimal amenity creation—see Lemma 1 and the functional form in equation (8)—which yields

$$\tilde{c}^{a,0}(r) = [a(r)]^{1-\eta^a}. \tag{D.48}$$

Regardless of whether or not the preference weights $\beta(r)$ are known, after having identified productivity $p(r)$ for men, under the normalization that $\tau(r) = 0$ for them, and $(1 - \tau(r))p(r)$ for women from equation (D.47), we recover the gender wedge $\tau(r)$ based on the ratio of estimated productivities net of the gender wedge at dual-gender firms. It is worth noting that, for this last step, the exact functional form by which $\tau(r)$ enters firms' payoff function (e.g., multiplicatively or additively) is

[64] Our choice of amenities starting strictly above but arbitrarily close to zero seems natural given the interpretation of amenities being endogenously produced by firms. It can be rationalized through an assumption that some employers put a vanishingly small weight on the (unmeasured by the tax authorities) amenity-related welfare of their workforce. At the same time, this choice minimizes the share of amenity valuations in total compensation across firms and maximizes the remaining gender gap in log total compensation. In contrast, if we modeled amenities as exogenous firm characteristics, then the choice of the constant of integration K would be inconsequential for our analysis as the absolute level of amenities would not be pinned down in that case.

immaterial for identification and estimation. □

D.8 Identification of Firm-Level Parameters with Discrete Firm Types

Here, we adapt the identification proof in Section D.7 to discrete data where we observe a finite number of firms N. We observe their ranks, defined as $r \in \{1/N, 2/N, \ldots, 1\}$, their empirical hiring intensities f_r and their wage w_r.[65] The FOCs read:

$$h'(\tilde{p}_r) = \frac{1}{V}\left[\frac{T(\tilde{p}_r - x(\tilde{p}_r))}{c^{v,0}}\left(\frac{1}{\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}_r))}\right)^2\right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}_r), \tag{D.49}$$

$$x'(\tilde{p}_r) = \frac{1}{V}\frac{2\lambda^E(\tilde{p}_r - x(\tilde{p}_r))}{\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}_r))}\left[\frac{T(\tilde{p}_r - x(\tilde{p}_r))}{c^{v,0}}\left(\frac{1}{\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}_r))}\right)^2\right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}_r). \tag{D.50}$$

How do we move from the continuous representation to the discrete case? In what follows, we want to move from functions of model objects (e.g., $\tilde{p}_r$) to functions of ranks, r. For instance, the change in the CDF of recruiting intensities between two (discrete) ranks $r - 1$ and r is

$$\Delta h(\tilde{p}_r) \equiv \int_{\tilde{p}_{r-1}}^{\tilde{p}_r} \frac{v(\tilde{p})}{V}\gamma(\tilde{p})\, d\tilde{p} \tag{D.51}$$

$$\approx \frac{v(\tilde{p}_{r-1})}{V}\gamma(\tilde{p}_{r-1}) \times [\tilde{p}_r - \tilde{p}_{r-1}], \tag{D.52}$$

where Δ is an operator that takes differences between the current and the previous (discrete) rank. We write a discretized version of equation (D.43) as follows:

$$\Gamma(\tilde{p}_r) = \int_{\tilde{p}_1}^{\tilde{p}_r} \gamma(\tilde{p})\, d\tilde{p} \tag{D.53}$$

and interpreting the empirical (discrete) distribution of firms as representative of the theoretical (continuous) distribution of firms, the CDF of composite productivity $\tilde{p}$ of the first n firms in the ranking is simply[66]

$$\Gamma(\tilde{p}_r) = r \in \left\{\frac{1}{N}, \frac{2}{N}, \ldots, 1\right\}. \tag{D.54}$$

Therefore, a good approximation for the change in the CDF of composite productivity is simply

$$\gamma(\tilde{p}_{r-1})(\tilde{p}_r - \tilde{p}_{r-1}) \approx 1/N \tag{D.55}$$

[65] Instead of functional notation, here we denote firm-specific objects by subscript r to highlight that there is a discrete number of firms in the data.

[66] Recall that $r \in \{0, 1/(N-1), \ldots, 1\}$, where N is the total number of firms, so $n \equiv rN \in \{1, 2, \ldots, N\}$ is the rescaled rank of a firm, with workers of a given gender preferring firms with higher values of n.

at all rungs of the ladder.[67] Now we rewrite the change in the CDF of composite productivity offers using a finite-difference approximation and equation (D.49):

$$\Delta F(x(\tilde{p}_r)) = h'(\tilde{p}_{r-1})(\tilde{p}_r - \tilde{p}_{r-1}) \tag{D.56}$$

$$= \frac{1}{V}\left[\frac{T(\tilde{p}_{r-1} - x(\tilde{p}_{r-1}))}{c^{v,0}}\left(\frac{1}{\delta + \lambda^G + \lambda^E(1 - h(\tilde{p}_{r-1}))}\right)^2\right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}_{r-1})(\tilde{p}_r - \tilde{p}_{r-1}) \tag{D.57}$$

Now replace $\Delta F(x(\tilde{p}_r))$ by $\hat{f}_r$, replace $\gamma(\tilde{p}_{r-1})(\tilde{p}_r - \tilde{p}_{r-1})$ by $1/N$, and replace $h(\tilde{p}_{r-1})$ by F_{r-1}. It then follows that

$$\hat{f}_r = \left[\frac{T(\tilde{p}_{r-1} - x(\tilde{p}_{r-1}))}{c^{v,0}}\left(\frac{1}{\delta + \lambda^G + \lambda^E(1 - F_{r-1})}\right)^2\right]^{\frac{1}{\eta^v - 1}} \frac{1}{VN} \tag{D.58}$$

All of the elements of equation (D.58) are known from the data, except for the firm's flow payoff per matched worker, $P_{r-1} \equiv (\tilde{p}_{r-1} - x(\tilde{p}_{r-1}))$. Thus, we have one equation (D.58) in one unknown (P_r) for each firm rank r, which can be rewritten as

$$P_{r-1} \equiv (\tilde{p}_{r-1} - x_{r-1}(\tilde{p}_{r-1})) = \left(\hat{f}_r VN\right)^{\eta^v - 1} \frac{Vc^{v,0}}{T}\left(\delta + \lambda^G + \lambda^E(1 - F_{r-1})\right)^2. \tag{D.59}$$

Intuitively, a firm posts more vacancies if it receives a greater flow payoff per matched worker.

Going back to equation (D.50), we apply a similar finite-difference approximation:

$$\Delta x_r = x'(\tilde{p}_{r-1})(\tilde{p}_r - \tilde{p}_{r-1}) \tag{D.60}$$

$$= \frac{2\lambda^E(P_{r-1})}{\delta + \lambda^G + \lambda^E(1 - F_{r-1})}\left[\frac{TP_{r-1}}{c^{v,0}}\left(\frac{1}{\delta + \lambda^G + \lambda^E(1 - F_{r-1})}\right)^2\right]^{\frac{1}{\eta^v - 1}} \frac{1}{VN}. \tag{D.61}$$

Recall that equation (D.59) above already identifies P_r. Therefore, we can iteratively apply equation (D.61) through ranks $r > 0$ to deduce x_r as follows:

$$x_r = x_0 + \sum_{i=1}^{r} \Delta x_i. \tag{D.62}$$

For a given initial condition x_0, based on equation (D.62), we deduce

$$\beta_r a_r = x_r - w_r, \tag{D.63}$$

where $\beta_r a_r$ is the amenity valuation to be pinned down, x_r is the total compensation recursively defined in equation (D.62), and w_r is the known firm component of wages at firm rank r. Similarly to the continuous case, we set x_0 such that $\min\{\beta_r a_r\}_r \approx 0$. This allows us to pin down the distribution of firm-level amenity valuations, $\beta_r a_r$, based on equation (D.63). From this, we deduce the amenity cost, $c^a(a_r)$, by plugging back $\beta_r a_r$ into the expression for the optimal amenity cost in equation (C.9). Finally, as we have already identified $P_r = \tilde{p}_r - x_r = p_r - w_r - c^a(a_r)$, we can deduce p_r firm by firm as

$$p_r = P_r + w_r + c^a(a_r), \tag{D.64}$$

[67] Note that our data covers the universe (i.e., a very large number) of firms in Brazil. If we considered small sample of firms randomly drawn from the population, then for the first few firms in the left tail of the distribution of ranks our approximation may be relatively poor. However, as we sum over increasingly higher ranks, our approximation becomes increasingly precise since errors vanish rather than accumulate.

proceeding exactly as in the continuous case in Appendix D.7.

Having identified p_{Mj} for men and p_{Fj} for women at the same firm j, we can recover the gender wedge τ_j as

$$\tau_j = 1 - \frac{p_{Fj}}{p_{Mj}}. \tag{D.65}$$

In summary, equations (D.63), (D.64), and (D.65) jointly identify the unobserved multidimensional type $(p_{g,r}, \beta_{g,r} a_{g,r}, \tau_{g,r})$ of firms at any rank r and for both genders g. Intuitively, what our identification proof exploits is the fact that firm surplus and utility offers—consisting of productivity, the gender wedge, and the amenity value net of amenity creation costs—are identified by firms' labor demand as revealed through hires from nonemployment, while information on wage offers together with the already-revealed utility offers identifies firm amenity values. Finally, comparing outcomes across men and women at the same firm allows us to deduce the gender wedge, which stands in for differences in firm surplus of employing otherwise identical male and female workers at the same wage rate and amenity value net of amenity creation costs.

D.9 Proof of Proposition 5 (Economy-Wide Parameters)

Restatement of Proposition 5 (Economy-Wide Parameters). *(i) The vacancy cost shifter $c^{v,0}$ is identified based on the aggregate labor share; (ii) the elasticity of the vacancy cost function η^v is identified based on the firm pay-profit gradient; (iii) the elasticity of the amenity cost function η^a is identified based on the aggregate amenity cost share in the data.*

Proof. We proceed in three parts.

Part (i): vacancy cost shifter. First, we find an expression for firm profits $\rho\Pi(r)$. Recall the definition of profits per matched worker,

$$p(r) - w(r) - c^a(a_r) = [f(r)V]^{\eta^v - 1} \frac{c^{v,0}}{T} \left[\delta + \lambda^G + \lambda^E(1 - F(r))\right]^2. \tag{D.66}$$

From this expression, we multiply by size $l(r) = f(r)VT/[\delta + \lambda^G + \lambda^E(1 - F(r)]^2$ and subtract vacancy posting costs to obtain flow profits $\rho\Pi(r)$,

$$\rho\Pi(r) = (p(r) - w(r) - c^a(a_r))\, l(r) - c^v(v(r)) \tag{D.67}$$

$$= [f(r)V]^{\eta^v} c^{v,0} - c^v(v(r)). \tag{D.68}$$

We now substitute the functional form of the cost of posting vacancies $c^v(v(r)) = c^{v,0} v(r)^{\eta^v} / \eta^v$ and $f(r)V = v(r)$ to obtain

$$\rho\Pi(r) = [v(r)]^{\eta^v} c^{v,0} - c^{v,0} \frac{[v(r)]^{\eta^v}}{\eta^v} \tag{D.69}$$

$$= \left(1 - \frac{1}{\eta^v}\right) [v(r)]^{\eta^v} c^{v,0}. \tag{D.70}$$

From this expression it is immediately clear that flow profits $\rho\Pi(r)$ are proportional to $c^{v,0}$. Specifically, profits are always positive but they can be arbitrarily large as $c^{v,0} \to \infty$ and arbitrarily small as $c^{v,0} \to 0^+$, and therefore the profit share of the economy can range from 0 to 1 depending on the value of $c^{v,0}$.

Substituting back $v(r) = f(r)V$ to make clear the dependency on observed hiring intensities $f(r)$ we obtain

$$\rho\Pi(r) = \left(1 - \frac{1}{\eta^v}\right) [f(r)V]^{\eta^v} c^{v,0} \tag{D.71}$$

In equation (D.71), η^v is treated as unknown, however its value does not matter for any of our argument concerning identification of the vacancy cost shifter $c^{v,0}$. We can write the labor share as

$$\mathcal{L} \equiv \frac{\int w(r)l(r)\,dr}{\int [p(r)l(r) - c^a(a(r))l(r) - c^v(v(r))]\,dr} \tag{D.72}$$

$$= \frac{\int w(r)l(r)\,dr}{\int [w(r)l(r) + \rho\Pi(r)]\,dr} \tag{D.73}$$

$$= 1 - \frac{\int \rho\Pi(r)\,dr}{\int [w(r)l(r) + \rho\Pi(r)]\,dr}. \tag{D.74}$$

Since flow profits $\rho\Pi(r)$ in equation (D.71) are strictly increasing in the vacancy cost shifter $c^{v,0}$ and the labor share in equation (D.74) is strictly decreasing in profits $\rho\Pi(r)$, it follows that the labor share is strictly decreasing in the vacancy cost shifter $c^{v,0}$. Furthermore, the labor share in equation (D.74) can take on any value in $(0, 1)$ by choosing an appropriate value of the vacancy cost shifter $c^{v,0}$. As a result, the vacancy cost shifter $c^{v,0}$ is identified based on the aggregate labor share in the data. This proves part (i).

Part (ii): elasticity of the vacancy cost function. Next, we show that the variance of log profits is monotonically increasing in η^v. Applying natural logarithms to both sides of (D.71) yields

$$\ln(\rho\Pi(r)) = \ln\left(1 - \frac{1}{\eta^v}\right) + \eta^v \ln[f(r)V] + \ln(c^{v,0}). \tag{D.75}$$

Taking variances on both sides of equation (D.75), we get

$$Var\left[\ln(\rho\Pi(r))\right] = Var\left[(\eta^v \times \ln f(r)) + \eta^v \times \ln V + \ln\left(c^{v,0}\right)\right] \tag{D.76}$$

$$= (\eta^v)^2 \times Var\left[\ln f(r)\right] \tag{D.77}$$

Except for $f(r)$, all terms in equation (D.76) are constant across firms, so they drop out of the calculation of the variance. As a result, equation (D.77) shows that the variance of log profits is proportional to $(\eta^v)^2$ and thus monotonically increasing in η^v. Now consider a regression of log firm pay on log profits,

$$\ln w(r) = \alpha + \beta \ln \Pi(r) + \varepsilon(r), \tag{D.78}$$

where the regression coefficient β in equation (D.78) captures the elasticity of firm pay with respect to firm profits. The regression coefficient

$$\beta = \frac{Cov[\ln w(r), \ln(\rho\Pi(r))]}{Var[\ln(\rho\Pi(r))]}, \tag{D.79}$$

is inversely proportional to the variance of log profits, which scales in $(\eta^v)^2$, and proportional to the covariance between log profits and observed log pay, which scales in η^v, we know that the regression coefficient is strictly decreasing in η^v. Thus, if an empirical regression coefficient β is attained for some parameter value η^v, then this is the unique value of η^v that rationalizes this empirical β. As a result,

η^v is identified based on the elasticity of firm pay to firm profits. This proves part (ii).

Part (iii): amenity cost elasticity. Finally, in equation (C.9) in Lemma 1 we demonstrated that amenity costs are inversely proportional to η^a. Therefore, we can write the economy-wide cost share of amenities as

$$\mathcal{A} \equiv \frac{\int c^a(a^*(r))l(r)\,dr}{\int [p(r)l(r) - c^a(a^*(r))l(r) - c^v(v(r))]\,dr}. \tag{D.80}$$

Obviously, the aggregate amenity cost share in equation (C.9) is monotonically decreasing in η^a, as the numerator is monotonically decreasing in η^a and the denominator is monotonically increasing in η^a. Furthermore, as $\eta^a \to \infty$, the amenities cost share monotonically tends to zero. Thus, if an empirical value of the aggregate amenity cost share $\mathcal{A}$ is attained for some parameter value η^a, then this is the unique value of η^a that rationalizes this empirical $\mathcal{A}$. This proves that the elasticity of the amenity cost function η^a is identified based on the aggregate amenity cost share in the data. □

D.10 Recovering Parameter Values 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. As our proof shows that all parameters are point-identified, this exercise is not strictly necessary, but we view it as a proof of concept and as further validation of our strategy.

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 $(1-\tau)p$ if the algorithm is run on women's data. We start by drawing 100,000 firms, characterized by productivity p and amenities a, jointly normally distributed with correlation $\rho(p,a)$. We then transform productivity to have Pareto marginal distribution.

We then feed the nonparametric joint distribution of $\{p,a\}$ to our discrete numerical solution algorithm, detailed in Appendix G.1, to solve our model. The outputs of the simulation are firm-level wages, amenities, ranks and hiring intensities. We use this data to construct our estimates of rank r, hiring intensities f_r, and offer CDF F_r as explained in Proposition 3. 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, hiring intensities f_r, and CDF F_r to estimate amenities and productivity at the firm level.

Our results are summarized in Table D.1, which shows moments of the distribution of recovered estimates under different parametrizations of the underlying amenities distribution. Under different parametrizations 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 *identical* to the true values up to machine precision. In all experiments, we keep the marginal distribution of productivity fixed, but we alter the economy-wide parameters η^v and η^a, as well as the underlying variance of amenities $Var(a)$ and the correlation of amenities and productivity $\rho(p,a)$ in the initial joint normal distribution. η^v and η^a can be easily read in Panel D of Table D.1. In Column (1), $Var(a) = 0.1$ and $\rho(p,a) = 0$. In Column (2), $Var(a) = 0.15$ and $\rho(p,a) = 0$. In Column (3), $Var(a) = 0.125$ and $\rho(p,a) = 0$. In Column (4), $Var(a) = 0.2$ and $\rho(p,a) = -0.5$. In Column (5), $Var(a) = 0.15$ and $\rho(p,a) = 0.5$

To evaluate the goodness of fit of our procedure, Panel E shows the correlations and mean squared error (MSE) between our amenity estimates and true amenity values, and those between our productivity estimates and true productivity values. The correlation for all estimates is equal to 1 approximated at the seventh decimal digit. The MSE for amenities is close to machine zero.

Figure D.2 visualizes the fit of our estimation routine with respect to the main objects of interest—namely, the amenity values. Across all five simulations, estimates lie on the 45 degrees line, showing how our procedure is robust to different joint distributions of amenities and productivity, and different values of economy-wide parameters.

Table D.1. Monte Carlo simulation and estimation on 100,000 Firms

	(1)	(2)	(3)	(4)	(5)
Panel A. Properties of true wages w					
Correlation between w and r	0.318	0.205	0.248	0.292	-0.012
Panel B. Properties of true amenity values a					
Standard deviation of a	0.317	0.386	0.352	0.447	0.387
Correlation between a and w	-0.589	-0.640	-0.632	-0.750	-0.558
Correlation between a and r	0.555	0.611	0.585	0.404	0.817
Panel C. Properties of true productivity p					
Standard deviation of p	0.747	0.747	0.742	0.719	0.758
Correlation between p and w	0.701	0.590	0.637	0.822	0.376
Correlation between p and r	0.701	0.724	0.702	0.579	0.739
Correlation between p and a	0.051	0.131	0.070	-0.361	0.461
Panel D. Properties of estimated amenity values $\hat{a}$ and productivity $\hat{p}$					
Standard deviation of $\hat{a}$	0.317	0.386	0.352	0.447	0.387
Correlation between $\hat{a}$ and w	-0.589	-0.640	-0.632	-0.750	-0.558
Correlation between $\hat{a}$ and r	0.555	0.611	0.585	0.404	0.817
Standard deviation of $\hat{p}$	0.747	0.747	0.742	0.719	0.758
Correlation between $\hat{p}$ and w	0.701	0.590	0.637	0.822	0.376
Correlation between $\hat{p}$ and r	0.701	0.724	0.702	0.579	0.739
Correlation between $\hat{p}$ and $\hat{a}$	0.051	0.131	0.070	-0.361	0.461
Panel E. Goodness of fit					
Correlation between $\hat{a}$ and a	1.000	1.000	1.000	1.000	1.000
Mean Squared Error	0.000	0.000	0.000	0.000	0.000
Correlation between $\hat{p}$ and p	1.000	1.000	1.000	1.000	1.000
Mean Squared Error	0.000	0.000	0.000	0.000	0.000

Note: Table reports estimation results using simulated data from 100,000 firms under different parametrizations of the underlying distribution of firm heterogeneity, including amenity values a, productivity p, wages w, and employer ranks r. MSE denotes the mean squared error. Columns (1)–(5) show separate simulations under different parametrizations of the data-generating process described in the text of Appendix D.10. *Source:* Model simulations.

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

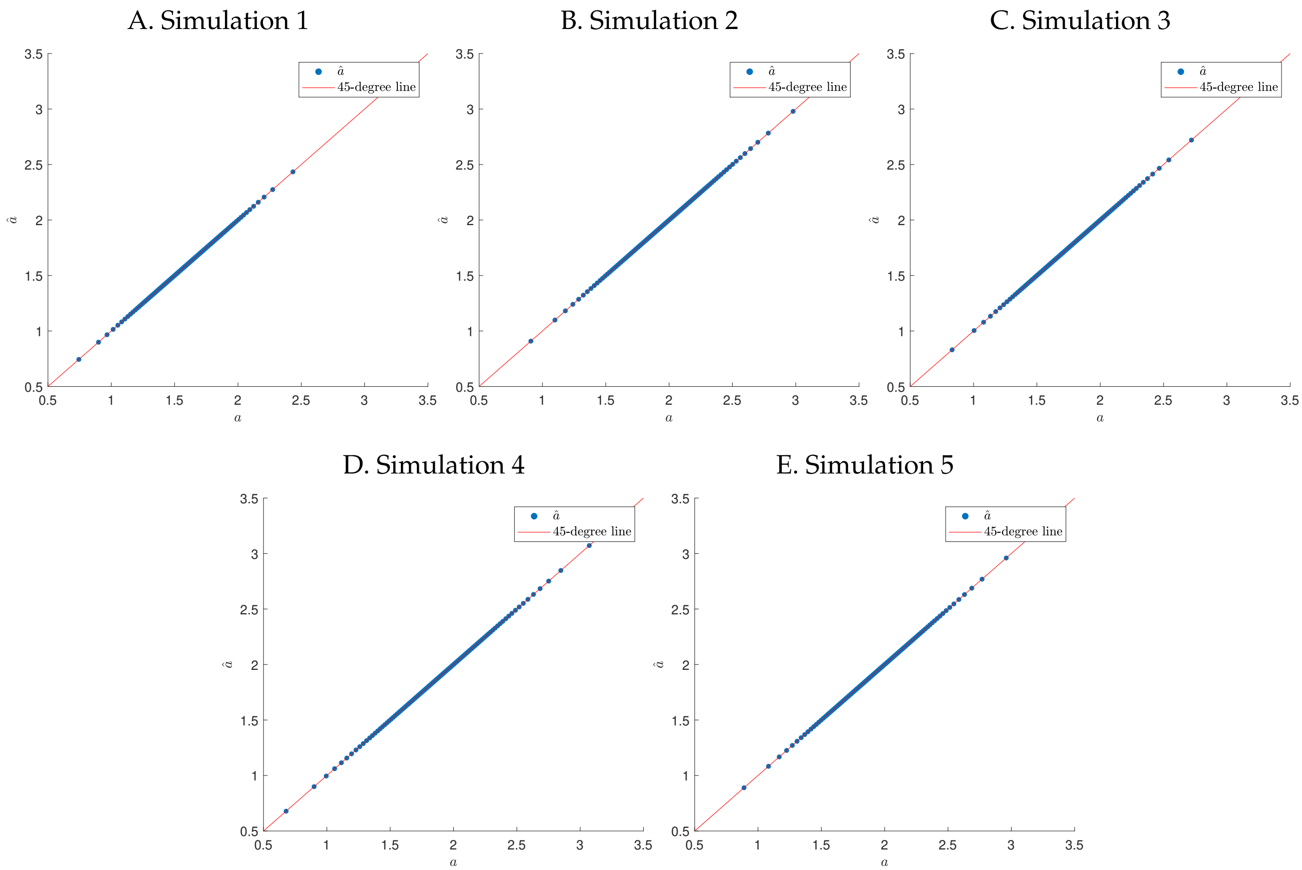

Note: Figure plots average amenity estimates against percentile bins of true amenity values. The simulations in Panels A–E correspond to columns (1)–(5) in Table D.1. *Source:* Model simulations.

D.11 Robustness to Model Misspecification and Subgroup Heterogeneity

To assess the robustness of our estimation results to model misspecification, we use a series of Monte Carlo simulations using a more flexible DGP that allows for heterogeneity in idiosyncratic worker preferences and firm attributes across population subgroups, even within gender. In what follows, we describe the details of how we estimate our (misspecified) model based on simulated data from such a DGP, and the results from this analysis. The bottom line is that we find our model to be remarkably robust to a wide range of misspecifications in the form of idiosyncratic worker and firm attributes.

To address the concern that preferences may be heterogeneous in the worker population and that, therefore, firm rankings may be heterogeneous, we run the following experiment. We assume that there are two groups of workers labeled A and B. Each firm is characterized by six objects: a shared productivity fundamental, p, a shared amenity cost shifter fundamental, $c^{a,0}$, two group-specific productivity components, $\{\epsilon_i\}_{i\in\{A,B\}}$, and two group-specific amenity cost components $\{\zeta_i\}_{i\in\{A,B\}}$. We define the group-specific productivity and amenity cost shifter, respectively, for population group $i \in \{A, B\}$ at firm j as:

$$p_{ij} = p_j + \epsilon_i,$$
$$c^{a,0}_{ij} = c^{a,0}_j + \zeta_i.$$

Figure D.3. Productivity estimates against true productivities in Monte Carlo simulations

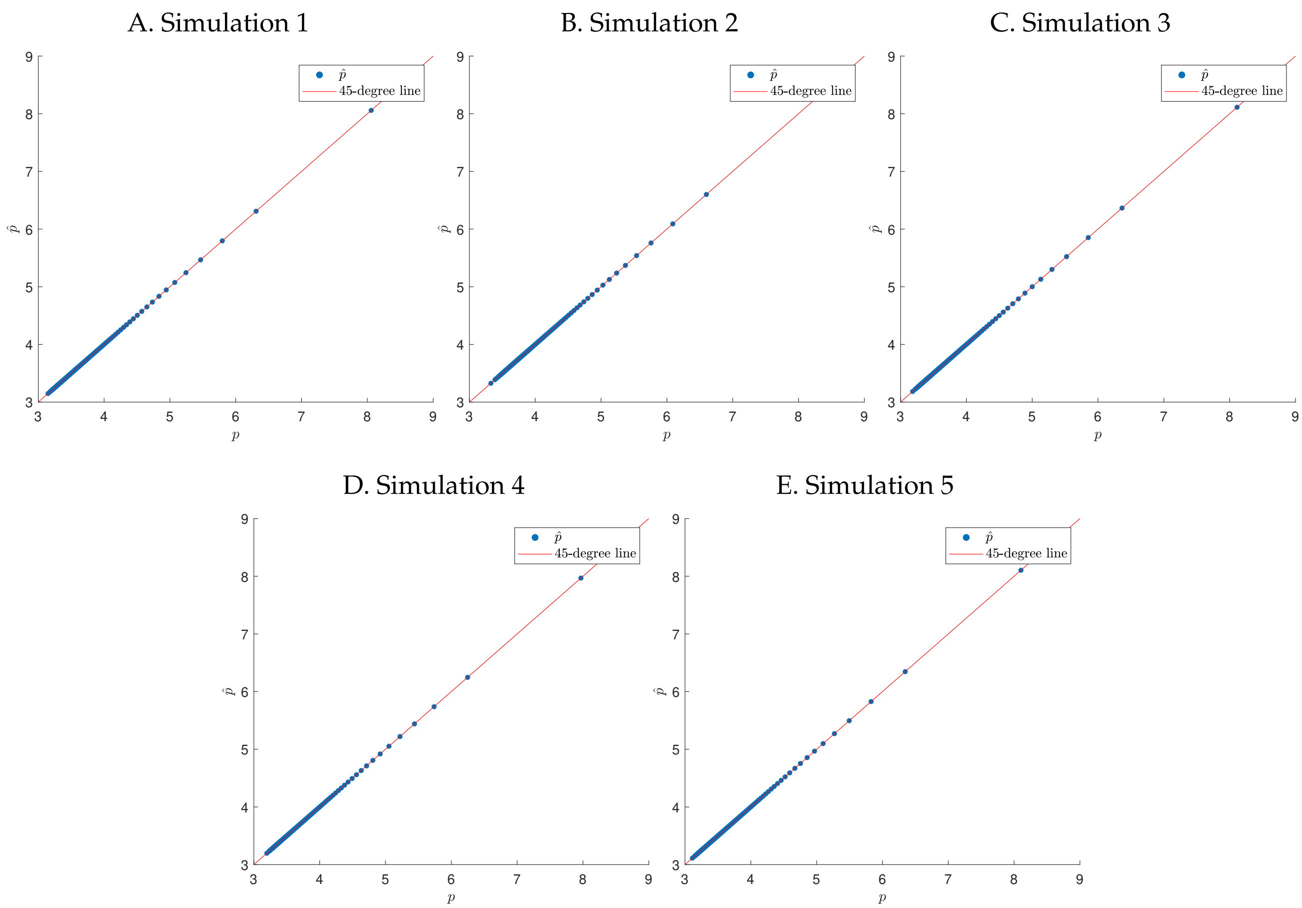

Note: Figure plots average productivity estimates against percentile bins of true productivity values. The simulations in Panels A–E correspond to columns (1)–(5) in Table D.1. *Source:* Model simulations.

We refer to p_{ij} and $c_{ij}^{a,0}$—and thus the implied equilibrium amenities a_{ij}—as the *group-specific fundamentals*, in contrast to the to fundamentals p_j and $c_j^{a,0}$ that are shared between worker groups at the same firm j. Starting from these two sets of group-specific fundamentals, we fix all labor market parameters $(\lambda^U, \lambda^E, \lambda^G, \delta)$ across groups and across simulations, and we fix the economy-wide parameters η^v and η^a across simulations. We then solve the general equilibrium across firms j in our model, separately for each group $i \in \{A, B\}$, as if these were two separate economies populated by workers who climb separate job ladders. This is the sense in which this simulation exercise features idiosyncratic worker preferences and firm rankings. Importantly, we can link workers within the same firm j by use of the simulated firm identifier.

Next, we treat the group identity $i \in \{A, B\}$ as unobserved by the econometrician. Instead, the econometrician observes pooled population averages as we measure them in our paper. To this end, we construct firm-level total employment g_j as the sum of employment of workers of type i, g_{ij}, and firm-level mean wages w_j as the weighted average of wages of workers of type i, w_{ij}, for each firm j by aggregating group-level quantities across model simulations for worker groups $i \in \{A, B\}$:

$$\begin{aligned} g_j =& g_{Aj} + g_{Bj}, \\ w_j =& \frac{g_{Aj}}{g_j} w_{Aj} + \frac{g_{Bj}}{g_j} w_{Bj}. \end{aligned}$$

We then use data on firm-level employment g_j and firm-level wages w_j as inputs to our estimation procedure, again assuming that the econometrician cannot separately observe outcomes of individual worker groups $i \in \{A, B\}$. Specifically, we use firm-level average employment g_j to infer firm ranks, we back out firm-level hires following the steps described in the paper, and we assign each firm their average wages as computed above. Importantly, our model with shared worker preferences across firms is generally misspecified vis-à-vis the DGP featuring two separate firm rankings of worker groups $i \in \{A, B\}$.

We construct three sets of experiments. In the first, the group-specific components ϵ_i and ζ_i are uncorrelated across groups. In the second, the group-specific components are perfectly negatively correlated across groups. In the last, the group-specific components are perfectly positively correlated. For each set of experiments, we simulate six different scenarios in which we let the variance of the group-specific components become an increasingly large multiple of the variance of the shared fundamental. In our most extreme case study, group-specific components of p_{ij} and $c_{ij}^{a,0}$ feature 200% of the dispersion as that of the shared fundamentals, including firm productivity and amenities. We then evaluate the performance of our algorithm by comparing the estimated values of productivity $\hat{p}$ and estimated amenities $\hat{a}$ with the true employment-weighted average productivity at the same firm $\bar{p}$ and the true employment-weighted average amenity valuation at the same firm $\bar{a}$, which we define as the employment-weighted average of group-specific fundamentals.

The simulation results from our first experiment, in which group-specific fundamentals are uncorrelated across groups, is summarized in Table D.2. We find that our estimation algorithm recovers well the *average* firm-level fundamentals in all exercises. As we increase the importance of group-specific fundamentals, the performance of our estimation procedure worsens, but it remains strong throughout, where the correlation between estimates and true averages never drops below 0.97. For comparison, we also display the correlation of the firm-specific average fundamentals that we estimate and the group-specific fundamentals. Obviously, as we increase the importance of group-specific drivers, the firm-level average becomes a worse predictor of the group-specific fundamental. However, it remains highly reliable at estimating the average fundamental.

For our second experiment, we assume that group-specific components are perfectly negatively correlated. Intuitively, it is more difficult for our estimation algorithm to accurately approximate the

Table D.2. Estimation under uncorrelated group-specific components

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Properties of group-specific components						
Relative variance of productivity component, ϵ_i	0.000	0.100	0.200	0.500	1.000	2.000
Relative variance of amenity component, ζ_i	0.000	0.100	0.200	0.500	1.000	2.000
Panel B. Goodness of fit of pooled estimates vs. population average						
Correlation between $\hat{p}$ and $\bar{p}$	1.000	0.999	0.996	0.988	0.980	0.975
Correlation between $\hat{a}$ and $\bar{a}$	1.000	0.999	0.998	0.995	0.988	0.980
Mean squared error of $\hat{p} - \bar{p}$	0.000	0.005	0.015	0.050	0.095	0.138
Mean squared error of $\hat{a} - \bar{a}$	0.000	0.000	0.003	0.019	0.035	0.068
Panel C. Goodness of fit of pooled estimates vs. group-specific fundamentals						
Correlation between $\hat{p}$ and p_A	1.000	0.974	0.949	0.886	0.821	0.775
Correlation between $\hat{p}$ and p_B	1.000	0.971	0.946	0.896	0.848	0.787
Correlation between $\hat{a}$ and a_A	1.000	0.974	0.952	0.907	0.858	0.791
Correlation between $\hat{a}$ and a_B	1.000	0.979	0.958	0.904	0.844	0.793
Mean squared error of $\hat{p} - p_A$	0.000	0.085	0.154	0.316	0.465	0.553
Mean squared error of $\hat{p} - p_B$	0.000	0.097	0.169	0.301	0.418	0.550
Mean squared error of $\hat{a} - a_A$	0.000	0.016	0.014	0.044	0.081	0.106
Mean squared error of $\hat{a} - a_B$	0.000	0.011	0.022	0.037	0.056	0.107

Note: This table reports results from estimation in the case of a misspecified model in which two groups of workers have different, uncorrelated group-specific components. Column 1 is the baseline estimation with no group-specific components. Between columns 2 and 6, the relative variances of group-specific components progressively increase from 10% to 200% of those of the shared components.

objects of interest when group-specific components are perfectly negatively correlated.[68] The simulation results from this second experiment are summarized in Table D.3. While the performance of our estimation algorithm slightly worsens, it is still surprisingly effective at recovering the firm-specific average productivities and amenity valuations, even when the relative variance of group-specific components is large. For productivity, the correlation between the estimated value and the true average remains at or above 0.972 through all exercises. For amenities, it decreases to 0.924, highlighting the challenge of identifying group-specific determinants of utility and employment when they are negatively correlated. Obviously, the relationship between our estimated average fundamentals and the group-specific fundamentals worsens.

Table D.3. Estimation under perfectly negatively correlated group-specific components

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Properties of group-specific components						
Relative variance of productivity component, ϵ_i	0.000	0.100	0.200	0.500	1.000	2.000
Relative variance of amenity component, ζ_i	0.000	0.100	0.200	0.500	1.000	2.000
Panel B. Goodness of fit of pooled estimates vs. population average						
Correlation between $\hat{p}$ and $\bar{p}$	1.000	0.996	0.988	0.975	0.975	0.981
Correlation between $\hat{a}$ and $\bar{a}$	1.000	0.997	0.994	0.980	0.957	0.915
Mean squared error of $\hat{p} - \bar{p}$	0.000	0.016	0.047	0.128	0.198	0.250
Mean squared error of $\hat{a} - \bar{a}$	0.000	0.001	0.003	0.012	0.076	0.274
Panel C. Goodness of fit of pooled estimates vs. group-specific fundamentals						
Correlation between $\hat{p}$ and p_A	1.000	0.946	0.888	0.780	0.704	0.684
Correlation between $\hat{p}$ and p_B	1.000	0.941	0.906	0.793	0.672	0.545
Correlation between $\hat{a}$ and a_A	1.000	0.944	0.912	0.795	0.646	0.447
Correlation between $\hat{a}$ and a_B	1.000	0.958	0.897	0.785	0.665	0.546
Mean squared error of $\hat{p} - p_A$	0.000	0.169	0.330	0.556	0.691	0.742
Mean squared error of $\hat{p} - p_B$	0.000	0.203	0.287	0.531	0.738	0.951
Mean squared error of $\hat{a} - a_A$	0.000	0.038	0.026	0.062	0.128	0.254
Mean squared error of $\hat{a} - a_B$	0.000	0.021	0.042	0.060	0.146	0.420

Note: This table reports results from estimation in the case of a misspecified model in which two groups of workers have different, perfectly negatively correlated group-specific components. Column 1 is the baseline estimation with no group-specific components. Between columns 2 and 6, the relative variances of group-specific components progressively increase from 10% to 200% of those of the shared components.

As a third experiment and sanity check, we let the group-specific components be perfectly positively correlated across groups. In this case, there is, of course, a common ranking for both worker groups $i \in \{A, B\}$. As a result of this shared ranking among all workers in the population, our estimation should deliver the exact values of productivity and amenities for any variance of the different components. The simulation results from this third experiment are summarized in Table D.4. Reassuringly, in this case, our estimation procedure successfully recovers the exact fundamentals—see also the additional Monte Carlo simulation results in Appendix D.10 of the revised paper.

[68] To see why a negative correlation causes difficulties, consider the case in which the shared fundamental is irrelevant so that group-specific components are the *only* determinant of a firm's rank, wages, and employment—i.e., $p_{ij} = \epsilon_i$ and $c_{ij}^{a,0} = \zeta_i$. If the group-specific components are perfectly negatively correlated across groups, then average employment at each firm is constant in expectation once aggregated across worker groups. This is despite the existence of clear firm rankings and meaningful firm size distributions within each worker group.

Table D.4. Estimation under perfectly positively correlated group-specific components

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Properties of group-specific components						
Relative variance of productivity component, ϵ_i	0.000	0.100	0.200	0.500	1.000	2.000
Relative variance of amenity component, ζ_i	0.000	0.100	0.200	0.500	1.000	2.000
Panel B. Goodness of fit of pooled estimates vs. population average						
Correlation between $\hat{p}$ and $\bar{p}$	1.000	1.000	1.000	1.000	1.000	1.000
Correlation between $\hat{a}$ and $\bar{a}$	1.000	1.000	1.000	1.000	1.000	1.000
Mean squared error of $\hat{p} - \bar{p}$	0.000	0.000	0.000	0.000	0.000	0.000
Mean squared error of $\hat{a} - \bar{a}$	0.000	0.000	0.000	0.000	0.000	0.000
Panel C. Goodness of fit of pooled estimates vs. group-specific fundamentals						
Correlation between $\hat{p}$ and p_A	1.000	1.000	1.000	1.000	1.000	1.000
Correlation between $\hat{p}$ and p_B	1.000	1.000	1.000	1.000	1.000	1.000
Correlation between $\hat{a}$ and a_A	1.000	1.000	1.000	1.000	1.000	1.000
Correlation between $\hat{a}$ and a_B	1.000	1.000	1.000	1.000	1.000	1.000
Mean squared error of $\hat{p} - p_A$	0.000	0.000	0.000	0.000	0.000	0.000
Mean squared error of $\hat{p} - p_B$	0.000	0.000	0.000	0.000	0.000	0.000
Mean squared error of $\hat{a} - a_A$	0.000	0.000	0.000	0.000	0.000	0.000
Mean squared error of $\hat{a} - a_B$	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports results from estimation in the case of a misspecified model in which two groups of workers have different, perfectly positively correlated group-specific components. Column 1 is the baseline estimation with no group-specific components. Between columns 2 and 6, the relative variances of group-specific components progressively increase from 10% to 200% of those of the shared components.

Our simulation-based findings that individual heterogeneity and other sources of noise approximately wash out in our model estimation can be rationalized based on an interesting property of our model. To see this, we start by revisiting equation (D.44) from the identification proof of Proposition 4 in Appendix D.7:

$$P_{r-1} \equiv (\tilde{p}_{r-1} - x_{r-1}(\tilde{p}_{r-1})) = \left(\hat{f}_r VN\right)^{\eta^v - 1} \frac{Vc^{v,0}}{T} \left(\delta + \lambda^G + \lambda^E(1 - F_{r-1})\right)\right)^2,$$

where P_r is the flow payoff per matched worker, $\tilde{p}_r$ is the composite productivity, $x_r(\tilde{p}_r)$ is the flow utility delivered to workers, f_r is the recruiting intensity, and F_r is the ladder rank, all for a firm at rank r, and V is a constant reflecting aggregate vacancies, N is the number of firms, η^v is the elasticity of the vacancy cost function, $c^{v,0}$ is the economy-wide vacancy cost intercept, T is a constant defined in equation (D.40), and $(\delta, \lambda^G, \lambda^E)$ are labor market transition rates.

What is key is that there are two idiosyncratic terms in the expression for P_r above: the recruiting intensity f_r and the firm rank F_r. Hence, both f_r and F_r are potential sources for estimation error based on idiosyncratic variation due to preference heterogeneity or other sources of noise. Let us first inspect the term reflecting a firm's estimated rank, F_r. Importantly, we see that F_r is premultiplied by $\lambda^E \approx 1\%$ and the whole term containing F_r is of magnitude $\approx 5\%$, which means that the squared term becomes of magnitude $\approx (5\%)^2 \ll 1\%$. Therefore, idiosyncratic variation in the estimate of a firm's rank, F_r, has limited influence on our estimates of P_r and thus firm productivity p and amenities a.

This leaves us with only one source of idiosyncratic variation in the estimation, namely firms' recruiting intensity f_r. In general, the impact of noise in f_r can be substantial in our model through the term $(\hat{f}_r VN)^{\eta^v - 1}$ on the right-hand side of the equation above. However, a result of our estimation procedure is that our estimated elasticity of the vacancy cost function is $\eta^v = 2.063 \approx 2$, which implies that $(\hat{f}_r VN)^{\eta^v - 1} \approx \hat{f}_r VN$. Therefore, P_r is approximately linear in $\hat{f}_r$, which implies that any idiosyncratic variation in $\hat{f}_r$ washes out in expectation when estimating P_r.[69] As a result, our estimation procedure recovers estimates of P_r that are unbiased. As a result of having an unbiased estimate of the firm's flow payoff from a matched worker, P_r, we obtain the change in utility offered across rungs of the firm ladder, Δx, defined in equation (D.61). Furthermore, when computing the levels of flow utility $x_r = x_0 + \sum_{i=1}^{r} \Delta x_i$ offered at each firm by use of equation (D.62), idiosyncratic noise in individual firms' P_r washes out because we are summing over many such terms across rungs of the ladder.[70] Therefore, our estimation procedure yields unbiased and stable estimates of flow utility x_r, which is robust to idiosyncratic variation or noise in f_r. As a result of plugging x_r back into equation (D.63), we also see that our amenity estimates are unbiased and stable for a given observation of the empirical firm pay w_r. Plugging all of these objects back into the definition of P_r above, we see that the model object that should pick up most of the idiosyncratic variation in f_r should be firm productivity p_r.

Finally, it is interesting to assess how much the data clearly violate the assumption of homogeneous employer preferences across population subgroups within each gender. To this end, we have split our sample by three worker characteristics, one at a time. These worker characteristics are: low education (i.e., at most high school) vs. high education (i.e., at least high school), ever parents (i.e., registered at least one parental leave between 2007 and 2018) vs. never parents (i.e., did not register any parental leaves between 2007 and 2018), and young cohorts (i.e., year of birth starting in 1968) vs. old cohorts (i.e., year of birth up to 1967). We then compute firm ranks as employment at a given firm, separately for each population subgroup. With this in hand, we are interested in how related the two

[69] The same logic also predicts that idiosyncratic variation in f_r should have more substantial effects on the accuracy of our model estimates, which we confirm in Monte Carlo simulations as described above.

[70] Again, the same logic also predicts that we should have more accurate estimates in the center of the distribution compared to the tails, which we confirm in additional Monte Carlo simulations.

population subgroups' firm ranks are within each gender. For example, this allows us to ask: To what extent does the firm ranking of women with kids differ from that of women without kids?

Figure D.4 provides a first look at the results from this exercise by showing binscatter plots of empirical employer ranks across population subgroups, separately by gender. The key take-away is that there is a strong and close-to-linear relationship between employer ranks of different subpopulations within each gender. In particular, the ranks are close to monotonic in each other. Of course, the linearity and monotonicity of the plotted relationships are not perfect. For example, bottom-ranked firms for population subgroup 1 (on the horizontal axis of each panel) tend to be higher-ranked for members of population subgroup 2 (on the vertical axis of each panel). These imperfections could either reflect preference heterogeneity or measurement error, both of which might be particularly pronounced in the tails of the distribution of firm ranks. Note also that the parent vs. nonparent comparison for men shows a relatively worse fit than for the other population subgroups and genders. This may be partly due to misclassification in what we group into the "parent" category, which consists of all workers who took a registered parental leave from work between 2007 and 2018. Since women are orders of magnitude more likely to go on paid parental leave than men, our sample of men who on parental leave might be a special population subgroup and also subject to measurement error. Overall, though, we take these results to be reassuring that firm ranks strongly tend to be shared across population subgroups within workers of a given gender.

D.12 Identification Under Heterogeneous Separation Rates

To discuss identification of our model when separation rates are heterogeneous, as sketched in Appendix C.9, we start as in Appendix D.7, expressing the first-order conditions (C.67) and (C.68) as a coupled pair of differential equations. We use the same changes of variables as in Appendix C.4, but we apply them to this more complex case in which $\delta_g(\tilde{p}_{gz})$ is a function of composite productivity $\tilde{p}_{gz}$. From now on, we drop the productivity and gender subscripts for readability. Define:

$$h(\tilde{p}) = F(x^*(\tilde{p})) = \frac{\int_{\tilde{p}' \geq \phi}^{\tilde{p}} v^*(\tilde{p}')\gamma(\tilde{p}')\, d\tilde{p}'}{V} \tag{D.81}$$

$$j(\tilde{p}) = G(x^*(\tilde{p})) = \frac{\int_{\tilde{p}' \geq \phi}^{\tilde{p}} l^*(\tilde{p}')\gamma(\tilde{p}')\, d\tilde{p}'}{\mu(1-u)} \tag{D.82}$$

$$h'(\tilde{p}) = f(x^*(\tilde{p}))x^{*\prime}(\tilde{p}) \tag{D.83}$$

$$j'(\tilde{p}) = g(x^*(\tilde{p}))x^{*\prime}(\tilde{p}) \tag{D.84}$$

$$f(x^*(\tilde{p})) = h'(\tilde{p})/x^{*\prime}(\tilde{p}) \tag{D.85}$$

$$g(x^*(\tilde{p})) = j'(\tilde{p})/x^{*\prime}(\tilde{p}) \tag{D.86}$$

Thus, by differentiating equations (D.81) and (D.82) using Leibniz's integral rule, we can write $f(x^*(\tilde{p})) = (v^*(\tilde{p'}/V)\gamma(\tilde{p})\, d\tilde{p}/\, dx^*(\tilde{p})$, and $g(x^*(\tilde{p})) = (l^*(\tilde{p})/(\mu(1-u))\gamma(\tilde{p})\, d\tilde{p}/\, dx^*(\tilde{p})$. Therefore, we can rewrite the first-order conditions as:

$$h'(\tilde{p}) = \frac{1}{V}\left[(\tilde{p}-x)\frac{1}{c^{v,0}V}\left(\frac{u+(1-u)s^E j(\tilde{p})+s^G}{\delta(\tilde{p})+\lambda^G+\lambda^E\left[1-h(\tilde{p})\right]}\right)\mu\lambda^U\right]^{\frac{1}{\eta^v-1}}\gamma(\tilde{p}) \tag{D.87}$$

$$x'(\tilde{p}) = (\tilde{p}-x)\left(\frac{\lambda^E h'(\tilde{p})}{\delta(x)+\lambda^G+\lambda^E\left[1-h(x)\right]}+\frac{(1-u)s^E j'(\tilde{p})}{u+(1-u)s^E j(\tilde{p})+s_g^G}\right). \tag{D.88}$$

Once again, equations (D.87) and (D.88) are more general versions of equations (D.49) and (D.50) where we allow separation rates to be a decreasing function of $\tilde{p}$. Using similar changes of variables

Figure D.4. Binscatter plots of employer ranks across population subgroups, by gender

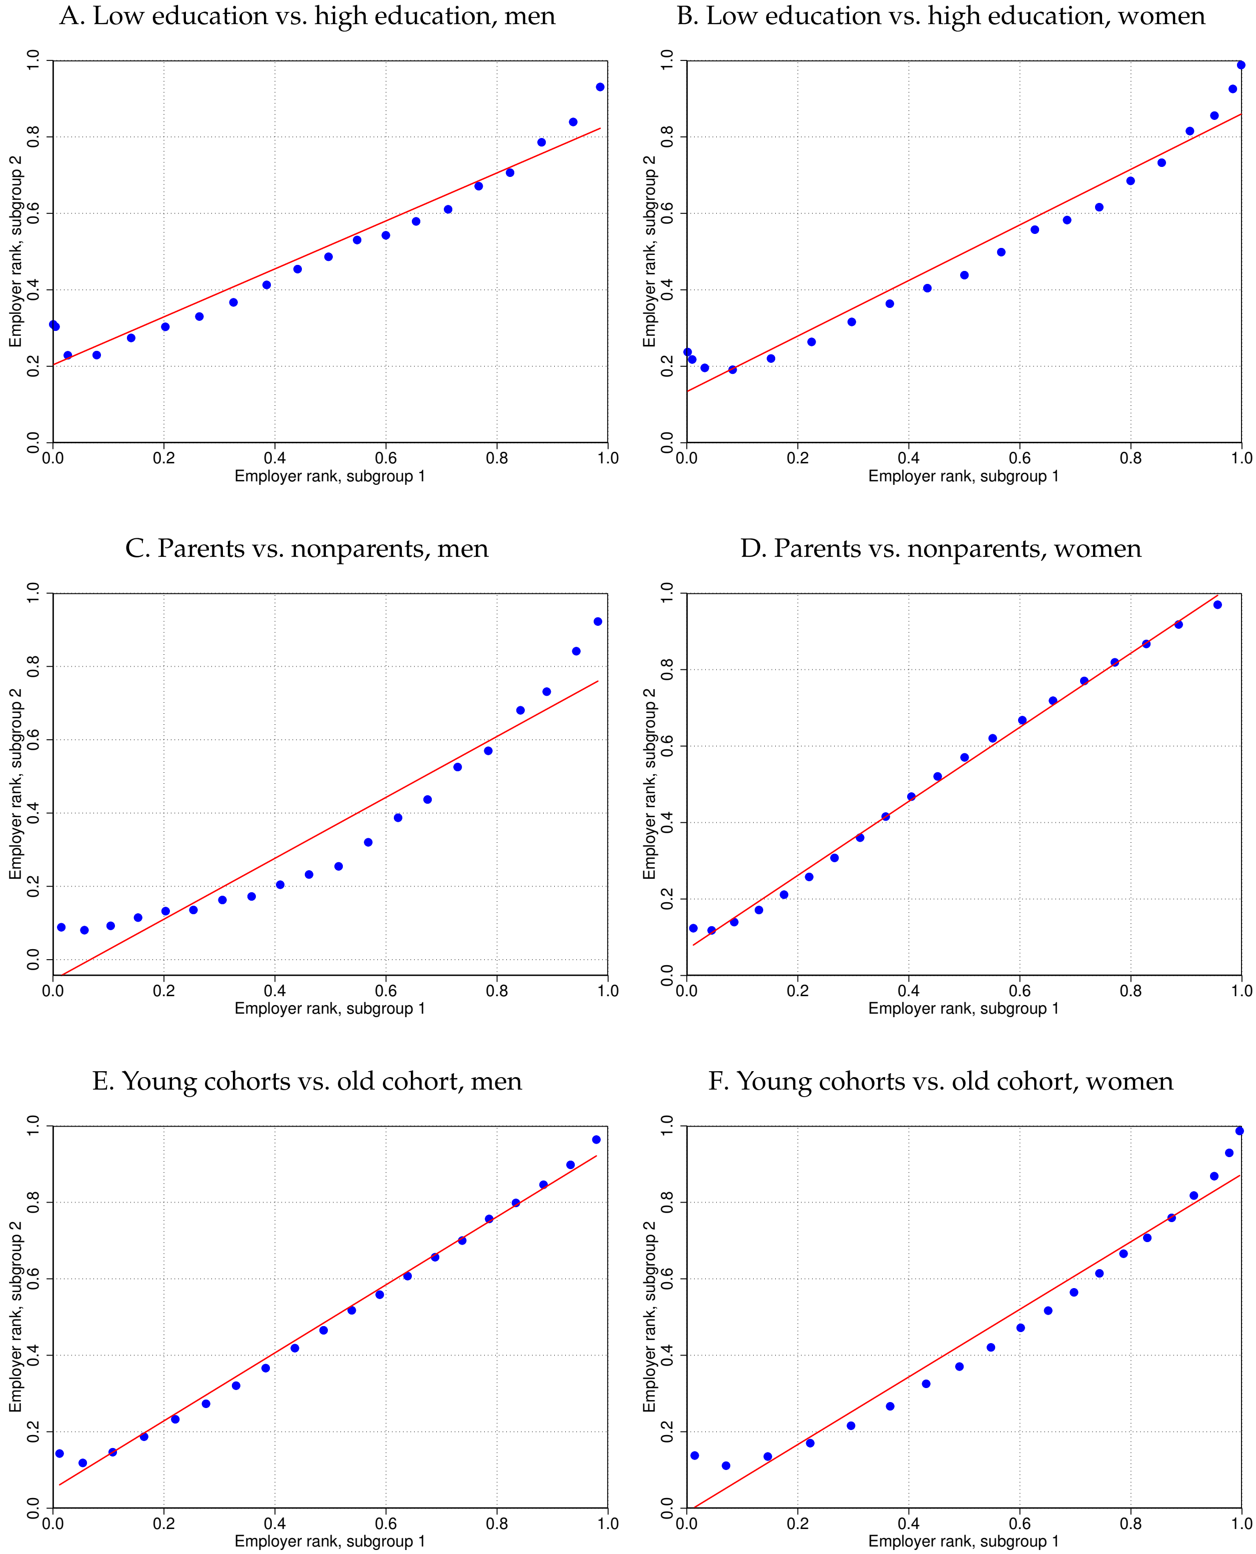

Note: This figure shows binscatter plots of empirical employer ranks across population subgroups, separately by gender. In each panel, the vertical axis shows the mean employer rank for population subgroup 2 corresponding to a set of firms in a given quantile of the distribution of employer ranks for population subgroup 1. Here, population subgroups 1 and 2 refer to the two population subgroups referenced in the relevant panel's caption. *Source:* RAIS, 2007–2014.

as we did in Appendix D.7, we can write an expression for profits per worker depending on ranks rearranging equation (D.87):

$$P(r) = \tilde{p}(r) - x(\tilde{p}(r)) = [f(r)V]^{\eta^v - 1} \frac{c^{v,0}V\left(\delta(r) + \lambda^G + \lambda^E[1 - h(r)]\right)}{\mu\lambda^U\left(u + (1-u)s^E j(r) + s^G\right)}, \tag{D.89}$$

which allows us to extract profits per worker from the data using a modified version of our original equation (D.44).

We now turn to the identification of x. Further substituting $j'(\tilde{p}) = l^*(\tilde{p})\gamma(\tilde{p})/(\mu(1-u))$, then $l^*(\tilde{p})$ with its solution from equation (C.66) and $h'(\tilde{p})$ from equation (D.87) yields:

$$\begin{aligned} x'(\tilde{p}) =& (\tilde{p} - x)\Bigg(\frac{\lambda^E}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \frac{1}{V} \left[(\tilde{p} - x)\frac{1}{c^{v,0}V}\left(\frac{u + (1-u)s^E j(\tilde{p}) + s^G}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \right)\mu\lambda^U \right]^{\frac{1}{\eta^v - 1}} \\ &+ \frac{s^E}{u + (1-u)s^E j(\tilde{p}) + s^G_g}\left(\frac{1}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \right) \frac{v^*(\tilde{p})}{V}\frac{\mu}{\mu}\lambda^U\left[u + (1-u)s^E_g j(\tilde{p}) + s^G \right] \Bigg)\gamma(\tilde{p}) \end{aligned} \tag{D.90}$$

which can be further simplified as:

$$\begin{aligned} x'(\tilde{p}) =& (\tilde{p} - x)\Bigg(\frac{\lambda^E}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \frac{1}{V} \left[(\tilde{p} - x)\frac{1}{c^{v,0}V}\left(\frac{u + (1-u)s^E j(\tilde{p}) + s^G}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \right)\mu\lambda^U \right]^{\frac{1}{\eta^v - 1}} \\ &+ \left(\frac{s^E\lambda^U}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \right)\frac{v^*(\tilde{p})}{V} \Bigg)\gamma(\tilde{p}). \end{aligned} \tag{D.91}$$

This further yields

$$x'(\tilde{p}) = \frac{2\lambda^E(\tilde{p} - x)}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \frac{1}{V} \left[(\tilde{p} - x)\frac{1}{c^{v,0}V}\left(\frac{u + (1-u)s^E j(\tilde{p}) + s^G}{\delta(\tilde{p}) + \lambda^G + \lambda^E\left[1 - h(\tilde{p})\right]} \right)\mu\lambda^U \right]^{\frac{1}{\eta^v - 1}} \gamma(\tilde{p}) \tag{D.92}$$

which is identical to our expression in (D.45), except for a change in the way firm sizes (and therefore vacancies) are computed. Interestingly, the optimality condition for utility turns out to be unaffected by heterogeneity in $\delta(\tilde{p})$: the reason is that, while the distribution of workers across firms in the economy is affected by how the separation rate varies with composite productivity, at the margin the employment distribution only varies with $f(x^*(\tilde{p})$ and $\delta(x(\tilde{p}))$, the local competition of a firm offering utility x. Thus, with a final change of variables and calling $\delta(r) = \delta(\tilde{p}(r))$, we can write:

$$\Delta x(r) = \frac{2\lambda^E P(r)}{\delta(r) + \lambda^G + \lambda^E\left[1 - F(r)\right]} \frac{1}{V} \left[P(r)\frac{1}{c^{v,0}V}\left(\frac{u + (1-u)s^E G(r) + s^G}{\delta(r) + \lambda^G + \lambda^E\left[1 - h(x)\right]} \right)\mu\lambda^U \right]^{\frac{1}{\eta^v - 1}} \tag{D.93}$$

Therefore, we proceed as in Appendix D.8, adapting our equations to discrete data. We first identify profits per worker $P(r)$ by applying equation D.89, plugging in the data on the employment distribution $G(r)$, the implied hiring intensities $f(r)$ and the associated empirical offer distribution $F(r)$ and . We then identify utility offers $x(r)$ by applying equation (D.93) iteratively, plugging in previously obtained profits per worker $P(r)$, the empirical offer distribution $F(r)$ and the empirical employment distribution $G(r)$. Finally, we proceed exactly as in Appendix D.8 to identify amenities, productivity and gender wedges.

E Estimation Results Appendix

E.1 Constructing Aggregate Statistics

We use three aggregate statistics for identification in our model: the aggregate labor share, the elasticity of firm pay with respect to firm value added per worker, and the aggregate amenity cost share.

Aggregate Labor Share. Part (i) of our identification result in Proposition 5 of Section 5 links the vacancy cost shifter $c^{v,0}$ to the aggregate labor share in the data. To bridge the model with the data, we define the labor share as the share of value added accruing to workers in pay,

$$\mathcal{L} \equiv \frac{\sum_g \sum_j w_{gj} l_{gj}}{\sum_g \sum_j \left[p_j l_{gj} - c_g^a(a_{gj}) l_{gj} - c_g^v(v_{gj}) \right]} = 48\%, \tag{E.1}$$

where the numerator contains the wage bill, $\sum_g \sum_j w_{gj} l_{gj}$, while the denominator contains value added, $\sum_g \sum_j [p_j l_{gj} - c_g^a(a_{gj}) l_{gj} - c_g^v(v_{gj})]$. The labor share value of 0.48 is the 2007 value of the share of labor compensation of employees (i.e., excluding the self-employed, who are outside of our model) for Brazil based on Feenstra et al. (2015) as retrieved via FRED.

Elasticity of Firm Pay with Respect to Firm Value Added per Worker. Part (ii) of our identification result in Proposition 5 of Section 5 links the elasticity of the vacancy cost function η^v to the elasticity of firm pay with respect to firm profits in the data. While firm profits are readily measured in firm financial data, we instead rely on existing estimates of the elasticity of firm pay with respect to value added per worker,

$$\tilde{\beta} = \frac{Cov(\ln w(r), \ln \left(\frac{\Pi(r) + w(r) l(r)}{l(r)} \right))}{Var(\ln \left(\frac{\Pi(r) + w(r) l(r)}{l(r)} \right))}, \tag{E.2}$$

To put a number on the elasticity $\tilde{\beta}$ in equation (E.2), we take the coefficient from a regression of firm fixed effects in wages on log value added per worker at the firm level in Brazil from Alvarez et al. (2018). The estimated coefficient is 0.179 in a balanced panel of Brazilian manufacturing firms from 2004–2008 (see Table E1 of Alvarez et al., 2018), which matches the beginning of our sample period.

Aggregate Amenity Cost Share. Part (iii) of our identification result in Proposition 5 of Section 5 links the elasticity of the amenity cost function η^a to the aggregate amenity cost share in the data. A challenge we face is that for the share of amenity costs in value added no precise estimate exists in the literature, much less so for the Brazilian context. Therefore, any assumed value necessarily comes with significant uncertainty. With this caveat in mind, related estimates on the value of local amenities (Bieri et al., 2023) suggest that the approximate cost share of amenities in value added is

$$\mathcal{A} \equiv \frac{\sum_g \sum_j c_g^a(a_{gj}) l_{gj}}{\sum_g \sum_j \left[p_j l_{gj} - c_g^a(a_{gj}) l_{gj} - c_g^v(v_{gj}) \right]} = 8\%. \tag{E.3}$$

Because there is significant uncertainty around this estimate, we use this number as a baseline and conduct robustness checks around it.

E.2 Details on Covariates Related to Productivity Estimates

To obtain productivity estimates for a subset of the firms in the RAIS linked employer-employee data, we merge in firm financial data from Bureau van Dijk's Orbis Historical database. The Orbis Historical data is the largest cross-country firm-level database containing information on public and private firms' balance sheets and income statements—see Kalemli-Özcan et al. (2024) for a detailed description. We merge the Orbis Historical data into RAIS using Brazilian firms' tax identifiers (*Cadastro Nacional de Pessoas Jurídicas*, or *CNPJ*). We then use as our empirical productivity proxy revenue per worker in the Orbis Historical data, which we relate to our model notion of firm productivity p.

E.3 Details on Covariates Related to Gender Wedge Estimates

We include as covariates in Z_j in equation (25) the following six variables that we construct using the RAIS data, in addition to two sets of FEs.

Female Manager: An indicator for whether the employer has a woman in the highest-paid position.

Nonroutine Manual Task Intensity: The mean z-score for nonroutine manual task intensity measured by linking 5-digit occupation codes from the Brazilian *Classificação Brasileira de Ocupações (CBO)* 1994 occupation classification to United States 1990 Census occupation codes based on the occupational crosswalk of de Souza (2022) extending previous work of Autor and Dorn (2009) and Acemoglu and Autor (2011).

Nonroutine Interpersonal Task Intensity: The mean z-score for nonroutine interpersonal task intensity measured by linking 5-digit occupation codes from the Brazilian *Classificação Brasileira de Ocupações (CBO)* 1994 occupation classification to United States 1990 Census occupation codes based on the occupational crosswalk of de Souza (2022) extending previous work of Autor and Dorn (2009) and Acemoglu and Autor (2011).

Mean Working Hours: The mean log number of contractual work hours.

No Major Financial Stakeholders: An indicator for whether an employer has no major financial stakeholder, as proxied by their participation in the small-business tax regime *Simples Nacional*.[71]

Employer Size: Total employer size measured as the log number of full-time equivalent employees during a year.

Municipality FEs: Dummies for 4,733 municipalities represented in our sample.

Sector FEs: Dummies for 661 five-digit sectors represented in our sample.

E.4 Details on Covariates Related to Amenity Estimates

We include in our analysis in Section 3.2 and also as covariates in Z_{gj} in equation (26) of Section 6.3 the following eight variables that we construct using the RAIS data, in addition to two sets of FEs.

[71] Eligibility 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.

Part-Time Work Incidence: The share of workers with contractual work hours below 40.

Working Hours Flexibility: The share of workers who change contractual work hours between any two consecutive years.

Parental Leave Generosity: An indicator for whether workers at an employer have parental leave duration above the national median.

Income Fluctuations: The share of workers who change mean earnings between any two consecutive years.

Workplace Hazards: The share of workers who report absence from work due to work-related illness, multiplied by 100.

Incidence of Unjust Firings: The share of workers who report ending their job due to an employer-induced firing for no officially recognized cause.

Incidence of Workplace Deaths: The share of workers who report ending their job due to death at the workplace.

Employer Size: Total employer size measured as the log number of full-time equivalent employees during a year.

Municipality FEs: Dummies for 4,733 municipalities represented in our sample.

Sector FEs: Dummies for 661 five-digit sectors represented in our sample.

E.5 Detailed Results from the Estimation of Labor Market Objects

Figure E.1. Employment-weighted density of employer ranks, by gender

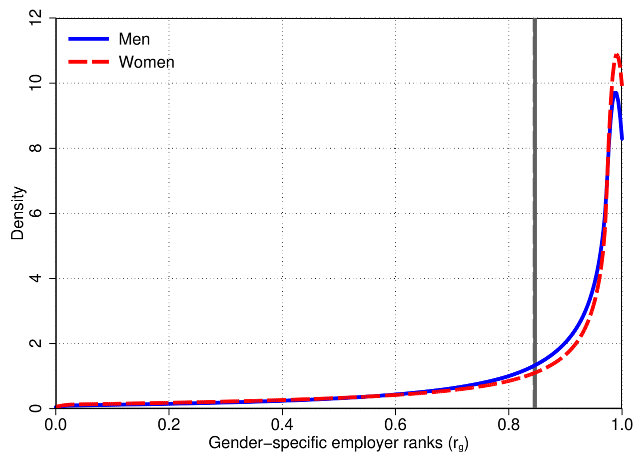

Note: This figure shows the distributions over employer ranks r_g using gender-specific employment weights separately for men (blue solid line) and women (red dashed line). Employer ranks r_g are defined to be uniformly distributed across firms. Grey vertical patterned lines represent mean values for workers of a given gender. *Source:* Model estimates based on RAIS, 2007–2014.

Figure E.2. Density estimates of recruiting intensities, by gender

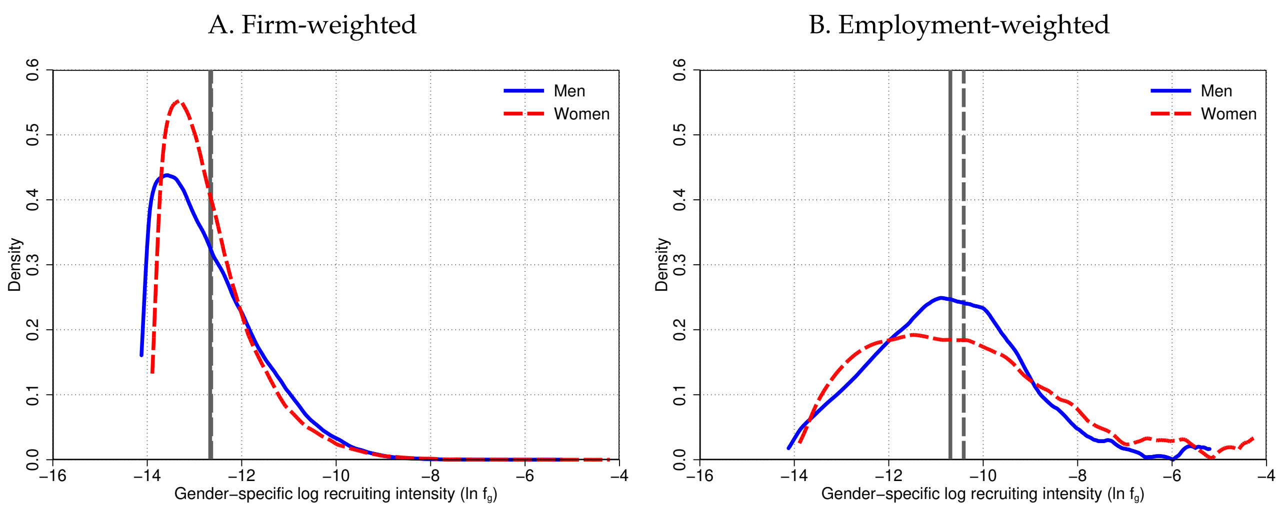

Note: This figure shows density estimates of logarithmic firm recruiting intensities $\ln f_g = \ln(v_g/V_g)$ separately for men (blue solid line) and women (red dashed line). Panel A is firm-weighted, while Panel B uses gender-specific employment distributions. Grey vertical patterned lines represent mean values for workers of a given gender. *Source:* Model estimates based on RAIS, 2007–2014.

E.6 Detailed Results from the Estimation of on Gender-Specific Firm Types

Figure E.3. Employment-weighted densities of log estimated productivity, by gender

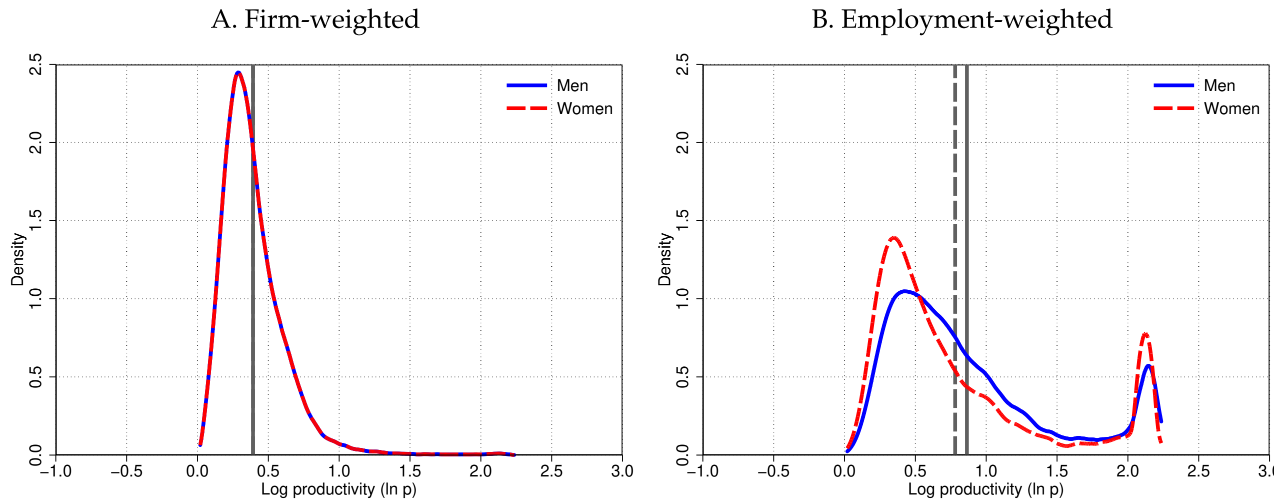

Note: This figure shows employment-weighted densities of log estimated productivity (p) separately for men (blue solid line) and women (red dashed line). Panel A is firm-weighted, while Panel B uses gender-specific employment distributions. Grey vertical patterned lines represent mean values for workers of a given gender. *Source:* Model estimates based on RAIS, 2007–2014.

Figure E.4. Employment-weighted densities of estimated gender wedges, by gender

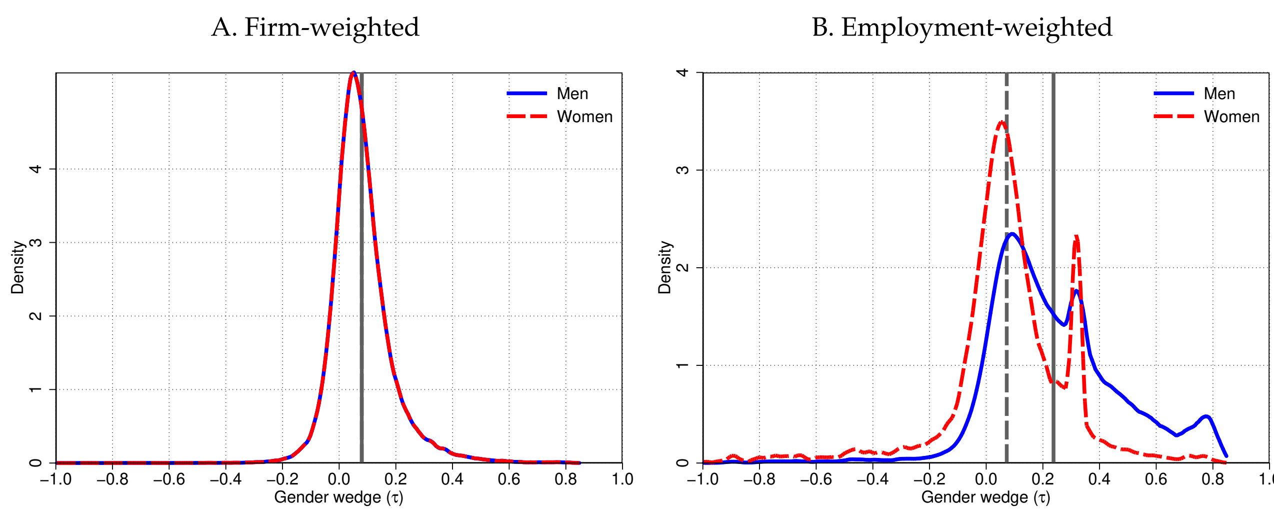

Note: This figure shows employment-weighted densities of women's gender wedges (τ) separately for men (blue solid line) and women (red dashed line). Panel A is firm-weighted, while Panel B uses gender-specific employment distributions. Grey vertical patterned lines represent mean values for workers of a given gender. *Source:* Model estimates based on RAIS, 2007–2014.

E.7 Additional Model-Fit Statistics

Figure E.6 shows the model versus data predictions for firm size and pay across 2-digit CNAE industry codes. These results suggest that the model captures well the distribution of pay and firm size

Figure E.5. Employment-weighted densities of log amenity cost shifters, by gender

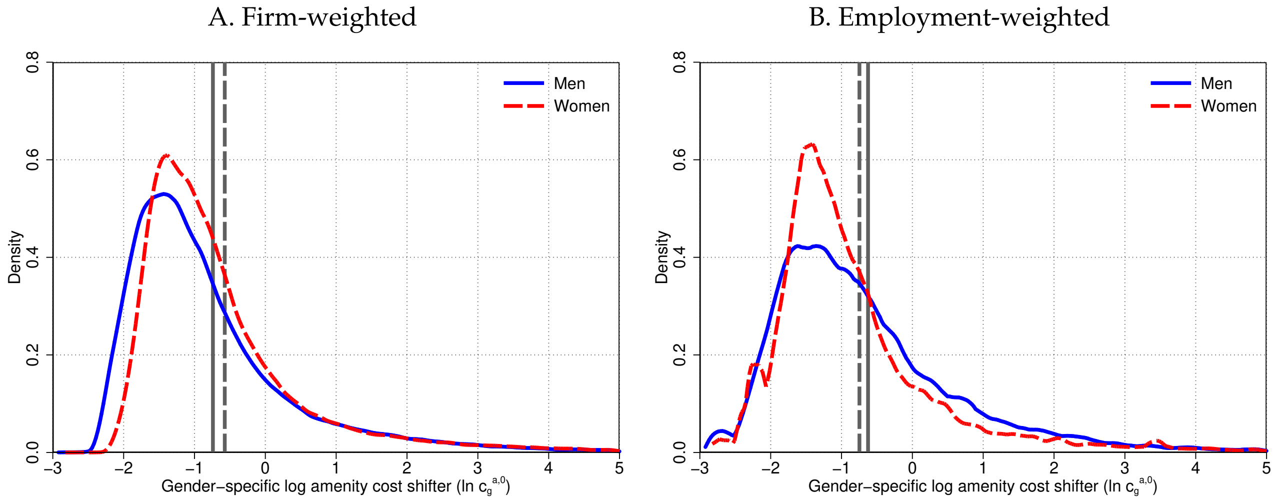

Note: This figure shows employment-weighted densities of log amenity cost shifters ($c_g^{a,0}$) separately for men (blue solid line) and women (red dashed line). Panel A is firm-weighted, while Panel B uses gender-specific employment distributions. Grey vertical patterned lines represent mean values for workers of a given gender. *Source:* Model estimates based on RAIS, 2007–2014.

Table E.1. Correlation table for estimated employer parameters

A. Men

	w_M	a_M	x_M	p	l_M	r_M
w_M	1.000					
a_M	−0.914	1.000				
x_M	0.246	0.168	1.000			
p	0.342	0.064	0.985	1.000		
l_M	0.097	0.133	0.552	0.504	1.000	
r_M	0.225	−0.025	0.486	0.456	0.160	1.000

B. Women

	w_F	a_F	x_F	$(1-\tau)p$	l_F	r_F
w_F	1.000					
a_F	−0.937	1.000				
x_F	0.020	0.331	1.000			
$(1-\tau)p$	0.162	0.187	0.970	1.000		
l_F	−0.085	0.282	0.578	0.476	1.000	
r_F	0.009	0.134	0.408	0.424	0.161	1.000

C. Cross-gender correlations

	w_g	a_g	x_g	$(1-\tau_g)p_g$	l_g	r_g
Cross-gender correlation	0.909	0.884	0.806	0.776	0.891	0.576

Note: This table reports employment-weighted pairwise correlations across employers between gender-specific pay (w_g), gender-specific amenities (a_g), gender-specific flow utilities (x_g), productivity net of the gender wedge ($(1-\tau_g)p$), gender-specific employment (l_g), and gender-specific employer ranks (r_g) separately for workers of gender $g \in \{M, F\}$. Panel A shows these correlations within the set of employers for men, while Panel B shows the same correlations for women. Panel C shows cross-gender correlations within the same employers. *Source:* Model estimates based on RAIS, 2007–2014.

across industries for both genders.

Figure E.6. Model fit in terms of mean log pay against log industry size, by gender

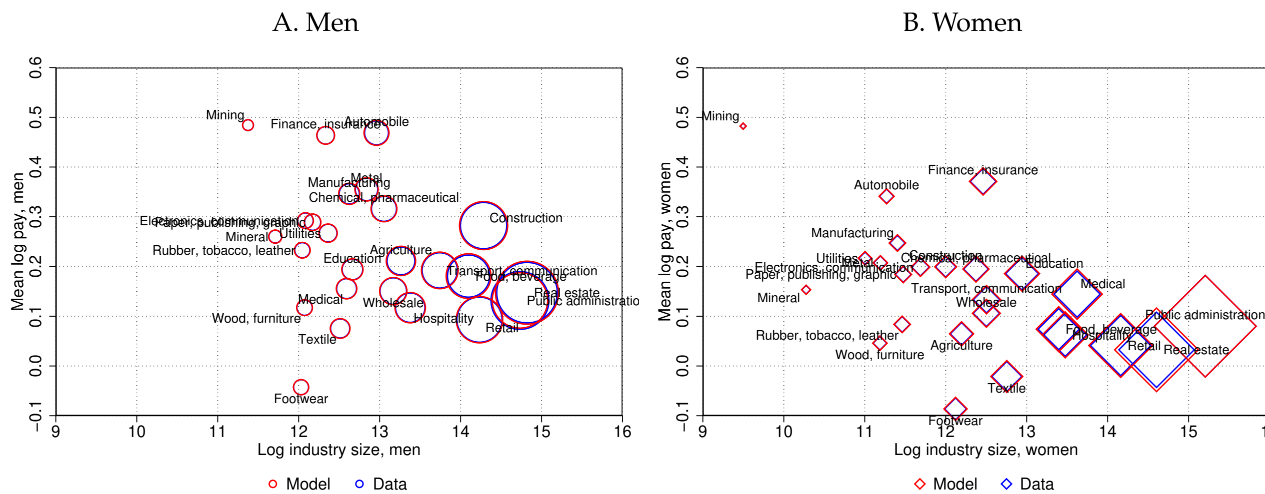

Note: This figure plots the model fit in terms of mean log pay against log industry size across 25 sectors, separately for men in Panel A and for women in Panel B. Marker sizes represent employment weights. *Source:* Model estimates based on RAIS, 2007–2014.

E.8 Relationship to Other Employer Rank Measures

In this Appendix, we show that our model is consistent with other methods to measure employer ranks, such as poaching ranks (Bagger and Lentz, 2019) or PageRanks (Page et al., 1998; Sorkin, 2018). To this end, we provide the following Proposition.

Proposition 7 (Employer Ranks using Worker Flows). *The ranking $r \in (0,1)$ can also be identified from worker flows between employers.*

Proof. We go through each alternative rank measure.

Poaching Ranks. The *poaching rank* (Bagger and Lentz, 2019) of every firm j is defined as

$$\text{Poaching rank}_j = \frac{\text{number of E-to-E hires}_j}{\text{number of all hires}_j} \tag{E.4}$$

$$= \frac{\text{number of E-to-E hires}_j}{\text{number of E-to-E hires}_j + \text{number of U-to-E hires}_j}, \tag{E.5}$$

which in our model can be rewritten as

$$\text{Poaching rank}_j = \frac{\lambda^E G_j + \lambda^G (1-u)}{\lambda^E G_j + \lambda^G + \lambda^U u} \tag{E.6}$$

The poaching rank in equation (E.6) is a monotonic transformation of the utility rank of a firm because it is strictly increasing in the cumulative employment distribution G_j, which is precisely the employment-weighted rank of firm j in the pool of all firms.[72] This proves that, in our model, the Poaching rank is monotonically increasing in firm rank.

[72] Note also that G_j itself is a monotonic transformation of the cumulative offer distribution F_j, which is the vacancy-weighted rank of firm j in the pool of all firms.

PageRanks. Another alternative is to exploit the full pattern of worker flows between employers in order to construct each employer's *PageRank* (Page et al., 1998; Sorkin, 2018), which represents the utility rank of a firm among the pool of all firms.

Let r_j be the rank of firm j. Let $f_j := v_j/V$ be the recruiting intensity of firm j.

Retentions. Let h_j be the mass of workers that firm j retains from one period to the next, defined as

$$h_j := \Big(1 - \delta - \lambda^G - \lambda^E(1 - F_j)\Big)\, l_j. \tag{E.7}$$

Here, we have used the fact that firms are atomistic. This fact implies, for example, that workers hit with an involuntary job offer at rate λ^G leave their current employer with probability one (as opposed to the complement of a firm's share of all vacancies in the economy, $1 - f_j$).

Separations. Let $o_{i,j}$ be the *outflow* of workers from firm i to firm j. There are two cases: Either $r_i > r_j$, in which case

$$o_{i,j} = \lambda^G l_i f_j, \tag{E.8}$$

or else $r_i < r_j$, in which case

$$o_{i,j} = (\lambda^E + \lambda^G) l_i f_j. \tag{E.9}$$

Markov Transition Matrix. The Markov transition matrix of the model can be written as follows. Rows are indexed by i and represent workers' current employer. Columns are indexed by j and represent workers' future employers. Fixing the row corresponding to a given firm i, the entries of this row are made up of the diagonal retention probability

$$\rho_i = \frac{h_i}{l_i}, \quad \forall j \in \{1, \ldots, N\} \tag{E.10}$$

and the off-diagonal separation probabilities

$$p_{i,j} = \frac{o_{i,j}}{l_i} \quad \forall j \neq i, j > 0. \tag{E.11}$$

Finally, we add the nonemployment state to the Markov transition matrix as firm 0, where the probability of transiting from any firm i to nonemployment is $p_{i,0} = p_0 = \delta$; the probability of finding a job in any firm j when nonemployed is $p_{0,j} = f_j(\lambda^G + \lambda^U)$; and the probability of remaining in nonemployment (i.e., nonemployment's retention probability) is $\rho_0 = p_{0,0} = (1 - \lambda^G - \lambda_U)$. To summarize, the Markov transition matrix looks as follows:

$$P = \begin{bmatrix} \rho_0 & p_{0,1} & p_{0,2} & \cdots & p_{1,N} \\ p_0 & \rho_1 & p_{1,2} & \cdots & p_{1,N} \\ p_0 & p_{2,1} & \rho_2 & \cdots & p_{2,N} \\ \vdots & \vdots & \ddots & \vdots & \\ p_0 & p_{N,1} & p_{N,2} & \cdots & \rho_N \end{bmatrix} \tag{E.12}$$

If we write the system $P' \times l = l$, where l is the $N + 1 \times 1$ vector of firm sizes that includes nonemployment as an additional "firm", it is easy to see that the above system is the discrete-version equivalent of the continuous Kolmogorov forward equation for the evolution of firm size in the model:

$$\dot{l}_j = \left[-\delta - \lambda^E \left[1 - F_j\right] - \lambda^G\right] l_j + \left[\frac{u + (1-u)s^E G_r + s^G}{u + (1-u)s^E + s^G}\right] v_j q. \tag{E.13}$$

where we can use the discrete time, discrete firms equivalents:

$$(1-u) = \sum_{j \in \{1/N,\dots,1\}} l_j, \tag{E.14}$$

$$G_j = \frac{1}{1-u} \sum_{i \in \{1/N,\dots,r\}} l_i, \tag{E.15}$$

$$q = m(\theta)/V = \lambda^U \left[\frac{u + (1-u)s^E + s^G}{V}\right] \tag{E.16}$$

$$F_j = \sum_{k \in \{1,\dots,N\}} (f_k \mathbf{1}[r(j) > r(k)]) \tag{E.17}$$

$$f_j = \frac{v_j}{V} \tag{E.18}$$

so we can rewrite that, in steady state,

$$0 = \left[-\delta - \lambda^E (1 - F_r) - \lambda^G\right] l_j + \left[u + (1-u)s^E G_r + s^G\right] \lambda^U \frac{v}{V}. \tag{E.19}$$

$$\left[\delta + \lambda^G + \lambda^E (1 - F_r)\right] l_j = \frac{v}{V} \left[\lambda^U u + \lambda^E (1-u) G_j + \lambda^G\right] \tag{E.20}$$

Using the definition of F_j and G_j for discrete firms, we can write:

$$\left[\delta + \lambda^G + \lambda^E \left(1 - \sum_{k \in \{1,\dots,N\}} (f_k \mathbf{1}[r(j) > r(k)])\right)\right] l_j = f_j \left[\lambda^U u + \lambda^G + \lambda^E \sum_{k \in \{1,\dots,N\}} (l \mathbf{1}[r(j) > r(k)])\right]. \tag{E.21}$$

Focusing on a single row, the system $P' \times l = l$ can be written as

$$f_j \left[\lambda^U l_0 + \lambda^G + \lambda^E \sum_{k \in \{1,\dots,N\}} (l_k \mathbf{1}[r(j) > r(k)])\right] \tag{E.22}$$

$$+(1 - \delta - \lambda^G - \lambda^E (1 - \sum_{k \in \{1,\dots,N\}} [f_k \mathbf{1}[r(j) > r(k)]]))l_j = l_j. \tag{E.23}$$

By substituting $l_0 = u$ by definition, it is immediately apparent that equations (E.21) and (E.23) are identical, proving that this matrix implementation of PageRank solves simultaneously for model sizes l_j that are implied by worker flows, and that are consistent with the model ranking. Solving for l_j yields

$$l_j = f_j \left[\frac{\lambda^U u + \lambda^G + \lambda^E \sum_{k \in \{1,\dots,N\}} (l_k \mathbf{1}[r(j) > r(k)])}{\delta + \lambda^G + \lambda^E \left(1 - \sum_{k \in \{1,\dots,N\}} (f_k \mathbf{1}[r(j) > r(k)])\right)}\right] \tag{E.24}$$

which makes explicit that the steady-state size l_j of every firm j depends on the relative rank in the ladder, where firms higher up the job ladder hire from a larger set of firms. Further rearranging and substituting for G_j leads us to the discrete equivalent of the steady state firm size in equation (16):

$$l_j = f_j \left[\frac{1}{\delta + \lambda^G + \lambda^E(1 - F_j)} \right]^2 (u + s^G)\lambda^U(\delta + \lambda^G + \lambda^E) \tag{E.25}$$

In short, if we construct the matrix of firm-to-firm flows in the data and solve the following equation,

$$P'l = l, \tag{E.26}$$

where the vector l contains the steady-state firm sizes of equation (E.25), which include hiring intensity f_j. Therefore, dividing firm size l by hiring intensity will give us the rank-implied firm size $\hat{r}_j$. This approach is similar in spirit to Sorkin (2018), who also uses worker flows between firms to identify firm ranks. However, there is an important distinction between our construction of PageRanks and that of Sorkin (2018). Specifically, our definition of PageRank is closer to the original one due to Page et al. (1998), in the sense that our modified adjacency matrix is a Markov transition matrix (i.e., the elements in each row sum to 1), whereas this is not the case in the discrete-choice setting of Sorkin (2018). Therefore, we have proved that we can use the pattern of worker flows to identify ranks in our model.

□

E.9 Comparison across Different Employer Rank Measures

For robustness, we computed a variant of PageRanks (Page et al., 1998; Sorkin, 2018) and poaching ranks (Bagger and Lentz, 2019) as described in Appendix E.8.[73] The two measures—PageRanks and poaching ranks—are related in so far as poaching ranks can be viewed as the result of a single application of the PageRanks' transition matrix to an initial guess of ones for all employers and zero for nonemployment. Viewed this way, PageRanks are obtained as the infinite (or converged) iteration of the same transition matrix that generated poaching ranks in the first iteration. See also the discussion in Lentz (2024). To be transparent, we will compare both PageRanks and poaching ranks to our baseline rank measure that is employer size.

The results of comparing our baseline rank measure of size to both PageRanks and poaching ranks are presented in Table E.2 below. We find that our model-consistent PageRank estimates and firm sizes align quite closely in the data, with correlations always above 0.87. Poaching Ranks are still positively correlated to both size and PageRank, albeit more weakly.

To see what factors load differently onto each employer rank measure, we compare our alternative employer rank measures across sectors. Specifically, we compute the mean PageRank (poaching index) ranks onto mean employer size by two-digit sector in the RAIS data. We then assess to what extent the two different rank measures align overall, and how the degree of alignment varies across sectors. Figure E.7 shows the results of this exercise. Panel A projects Pagerank ranks onto employer size ranks by sector for men, while Panel B does the same for women. Overall, there is a strong correlation between the two measures, as evidenced by the steep slope and relatively little dispersion

[73] Importantly, the variant of PageRanks that we use is consistent with our labor market model and closely parallels the original PageRanks proposed by Page et al. (1998). It is worth noting that our variant of PageRanks differs from that introduced by Sorkin (2018) in that our measure reflects the fixed-point weights assigned to each employer based on the empirical transition matrix—i.e., the adjacency matrix normalized such that all rows sum to unity. This is in contrast to the approach in Sorkin (2018) that constructs an alternative set of PageRanks defined based on the adjacency matrix normalized such that all columns (i.e., not rows) sum to unity. While the latter is the appropriate normalization in the discrete-choice setting of Sorkin (2018), we show in Appendix E.8 of our paper that the original definition of PageRanks due to Page et al. (1998) is the one consistent with our labor market model. In addition, we compute poaching ranks following the exact definition in (Bagger and Lentz, 2019).

Table E.2. Rank correlations between alternative employer rank measures, by gender

	Men			Women		
	Size	PageRank	Poaching rank	Size	PageRank	Poaching rank
Size	1.000			1.000		
PageRank	0.872	1.000		0.886	1.000	
Poaching rank	0.258	0.263	1.000	0.199	0.217	1.000

Note: This table reports Spearman rank correlations between three alternative employer rank measures, separately by gender. The three rank measures are employer size as used in the paper's baseline analysis, PageRanks (Page et al., 1998; Sorkin, 2018), and poaching ranks (Bagger and Lentz, 2019). See Appendix E.8 for details of how these are constructed. *Source:* RAIS, 2007–2014.

around the linear best-fit line. PageRanks and employer size also agree on the relative ranking of employers in several key sectors—e.g., the public sector for women is ranked near the very top according to both PageRanks and employer size. Panels C and D repeat the same analysis for the relationship between poaching ranks and employer size ranks across sectors for men and women, respectively. The slope is notably flatter and dispersion is notably greater, compared to the comparison between PageRanks and employer size in the panels A–B, consistent with the rank correlations reported in Table E.2 above. Nevertheless, there is a strong correlation between the two measures. Still, there are some important exceptions—e.g., the public sector for women is ranked near the top according to employer size but only above average according to poaching ranks. This may reflect differences in the quantity of worker outflows as well as the quality of worker in- and outflows of various sectors. It is reassuring that PageRanks, which take into account the entire empirical transition matrix of worker flows between employers, line up remarkably closely with our baseline analysis based on employer size.

We re-estimate our entire model based on a cardinal interpretation of PageRanks instead of employer size. To this end, we leverage the interpretation of PageRanks as the steady-state employer size implied by worker flows between firms (Page et al., 1998). This may be a preferable measure of employer rank if the cross-sectional distribution of employer sizes is contaminated by empirical factors such as history dependence—e.g., some employers used to be desirable to work and thus large, but they are no longer desirable to work at now after the realization of some shocks, though they remain large during a transition to their new steady-state size. Table E.3 compares the estimation results based on PageRanks to our baseline estimation results based on employer size as a measure of employer ranks. We find that the main properties of our estimates are remarkably stable. In particular, amenities estimated using PageRanks are very strongly correlated to our baseline amenity estimates. Productivity is also strongly correlated, though it is estimated to be more dispersed when using PageRanks than under our baseline estimation. Similarly, gender wedges are estimated to be more dispersed when using PageRanks to rank firms.

Building on these new model estimates, Table E.4 recomputes the Kitaga-Oaxaca-Blinder decomposition of gender gaps in pay, amenities, and total compensation based on PageRanks as an alternative employer rank.[74]

[74] It is worth noting that poaching ranks (Bagger and Lentz, 2019) have no cardinal interpretation in our model. The reason is intuitive: firm size and PageRanks both reflect different notions of the *number* of workers that a firm manages to attract and retain in steady state, while poaching ranks reflect the *share* of workers that a firm poaches from competitors *relative to all hires at that firm*. As such, poaching ranks normalize by total hiring, which removes cardinal information from this employer rank, making it impossible to recover firms' labor demand that can be mapped into composite productivity in our equilibrium framework. Thus, the same model inversion that we use for firm size in the baseline analysis (Table 11) and for PageRanks in a robustness exercise (Table E.4) is technically not possible for poaching ranks.

Figure E.7. Comparing alternative employer rank measures across sectors, by gender

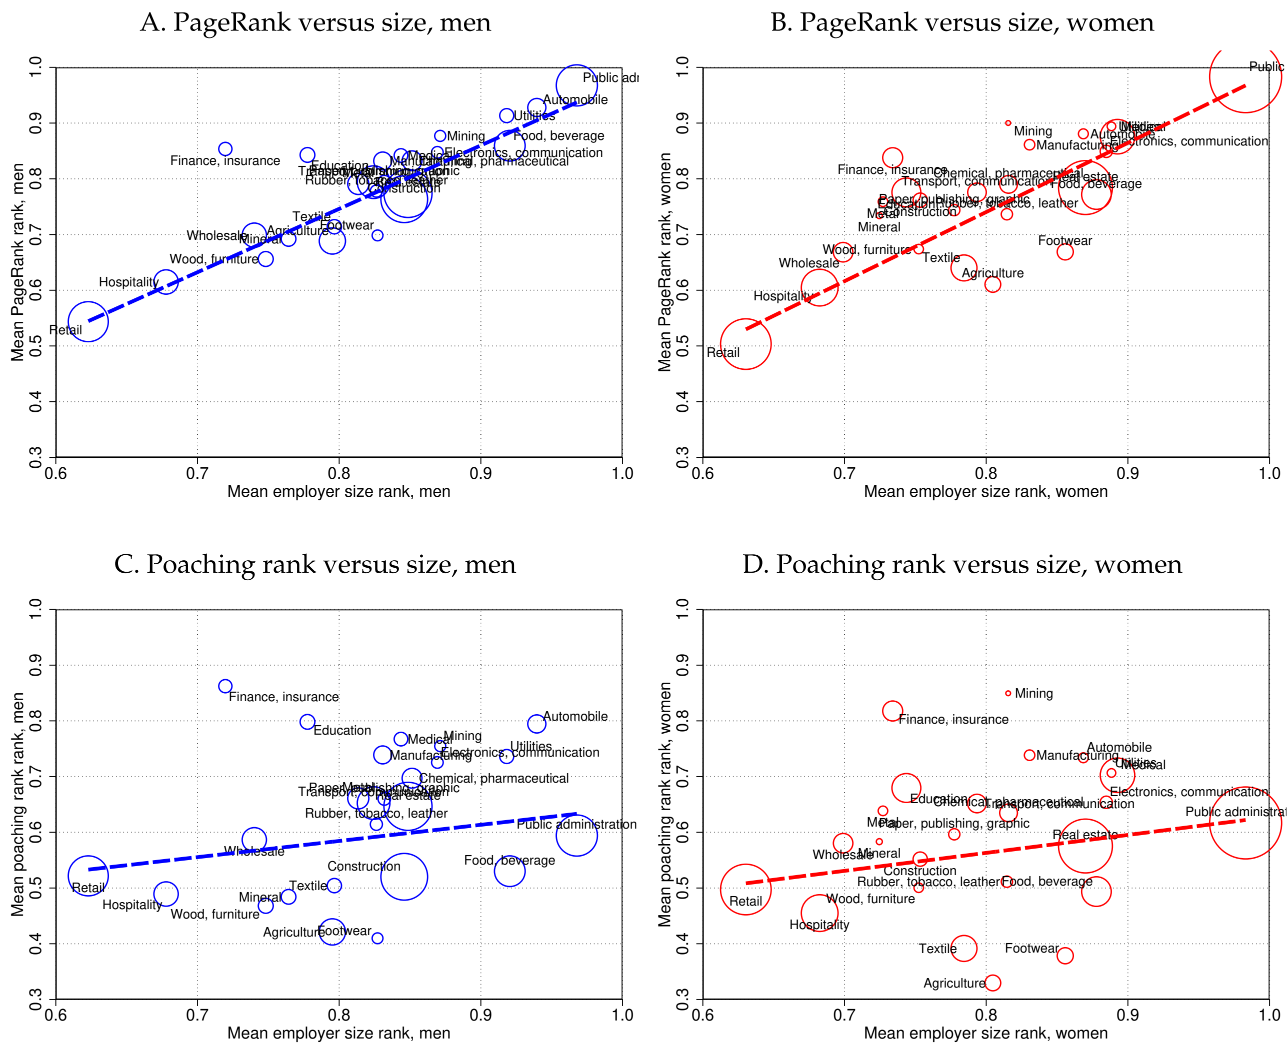

Note: This figure shows the means of alternative employer ranks—either PageRanks (Page et al., 1998; Sorkin, 2018) or poaching ranks (Bagger and Lentz, 2019)—against employer size ranks, which is the baseline rank measure in the paper. See Appendix E.8 for details of how these are constructed. *Source:* RAIS, 2007–2014.

Table E.3. Estimation results under PageRanks as alternative measure of employer ranks

	Baseline	PageRank-based
η_a	5.728	6.985
η_v	2.063	2.500
Mean of p	2.887	3.148
Variance of p	4.666	8.426
Mean of a_m	1.196	1.515
Variance of a_m	0.089	0.128
Mean of a_f	1.243	1.440
Variance of a_f	0.081	0.140
Mean of τ	-0.015	-0.126
Variance of τ	0.096	0.277
Correlation between p_{baseline} and alternative		0.860
Correlation between $a_{m,\text{baseline}}$ and alternative		0.998
Correlation between $a_{f,\text{baseline}}$ and alternative		0.999
Correlation between τ_{baseline} and alternative		0.801

Note: This table compares estimation results from our baseline model to one based on PageRanks as an alternative measure of employer ranks. *Source:* Model estimates based on RAIS, 2007–2014.

Table E.4. Kitagawa-Oaxaca-Blinder decompositions—alternative employer ranks based on PageRanks

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.089	78.7	0.024	21.3
Amenity-valuation	-0.079	-0.089	112.3	0.010	-12.3
Total compensation	0.034	0.000	0.0	0.034	100.0

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) when re-estimating our model based on PageRanks as an alternative employer rank measure. *Source:* Model estimates based on RAIS, 2007–2014.

The results based on PageRanks as a measure of employer ranks are in line with our baseline results shown in Table 11 of the main text. Comparing the two sets of results, we find that all of our results are qualitatively unchanged across the two estimation methods. In particular, we continue to find an amenity-valuation gap equal in favor of women, resulting in a total-compensation gap that is less than half of the gender pay gap and almost entirely within (rather than between) employers.

Altogether, the analysis above suggests that our baseline analysis based on size as a measure of employer rank is robust to alternative, worker-flow-based measures of employer ranks and that we get very similar results under such an alternative estimation approach.

E.10 Estimation Results Under Heterogeneous Separation Rates

To estimate our model with separation rates that are heterogeneous across firms, we use the modified procedure detailed in Appendix D.12. Our results are summarized in Table E.5. We find that estimated productivity distribution is more dispersed when separation rates are allowed to be heterogeneous across firms. Estimated amenities are also more dispersed; however, their estimated mean is lower, suggesting that a share of utility that is explained by amenities in the baseline model can be explained by heterogeneity in separation rates. We also find that heterogeneous separation rates differ more for men than for women between firms, which affects men's results more than women's. Thus, the gap in amenities between men and women widens compared to our baseline estimation. Finally, all estimates are strongly correlated to our baseline estimates, suggesting that though heterogeneity in separation rates qualitatively affects our results, quantitatively, the difference is relatively small.

Table E.5. Estimation results under heterogeneous separation rates

	Baseline	Heterogeneous separations
η_a	5.728	5.309
η_v	2.063	1.925
Mean of p	2.887	2.854
Variance of p	4.666	4.423
Mean of a_m	1.196	1.108
Variance of $\log a_m$	0.089	0.092
Mean of a_f	1.243	1.166
Variance of a_f	0.081	0.066
Mean of τ	−0.015	0.003
Variance of τ	0.096	0.083
Correlation between p_{baseline} and alternative		1.000
Correlation between $a_{m,\text{baseline}}$ and alternative		0.995
Correlation between $a_{f,\text{baseline}}$ and alternative		0.998
Correlation between τ_{baseline} and alternative		0.999

Note: This table compares estimation results from our baseline analysis with homogeneous separation rate (δ) across firms to the case in which separation rates ($\delta(r)$) are allowed to be nonincreasing in firm ranks. *Source:* Model estimates based on RAIS, 2007–2014.

To what extent do these differences in estimation results affect the substantive insights from our analysis? To answer this question, Table E.6 below shows the results of Kitagawa-Oaxaca-Blinder decompositions based on the model with heterogeneous firm-level separation rates ($\delta(r)$).

Comparing the two sets of results—i.e., those in the model with heterogeneous firm-level separation rates in Table E.6 compared to those in our baseline model in Table 11—we find two main results. First, the amenity-valuation gap is slightly more pronounced, implying that amenities play a slightly

Table E.6. Kitagawa-Oaxaca-Blinder decompositions—heterogeneous firm-level separation rates ($\delta(r)$)

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.089	78.7	0.024	21.3
Amenity-valuation	−0.079	−0.089	112.3	0.010	−12.3
Total compensation	0.034	0.000	0.0	0.034	100.0

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) when re-estimating our model with heterogeneous firm-level separation rates ($\delta(r)$) that are nonincreasing across firm ranks r. *Source:* Model estimates based on RAIS, 2007–2014.

more important role in explaining gender differences in total compensation. Second, the between- versus within-split of all three compensation concepts is similar across both models. Altogether, this analysis suggests that a model with heterogeneous firm-level separation rates yields results that are largely in line and, if anything, slightly more pronounced than what we find using our baseline model.

E.11 Sensitivity of Estimation Results to Alternative Amenity Cost Shares

Here, we further investigate the robustness of our results with respect to our estimation target of the cost share of amenities in value added, taken from Bieri et al. (2023). We acknowledge that this is a notoriously difficult moment to find. While an amenity cost share of value added of 8 percent is maybe a reasonable target—absent a better target from the literature—it is useful to check the sensitivity of our results to alternative targets by re-estimating our model to different amenity cost shares and evaluating the resulting model predictions.

We repeated our estimation for three different targets for the cost share of amenities: 4 percent, 12 percent and 16 percent. Our results are summarized in Table E.7, where we compare all these results to the baseline. Obviously, the assumed value of the cost share of amenities makes a difference for the estimated value of η_a, which is decreasing in the assumed cost share. However, we find that different values of the cost share of amenities—and therefore different values of the elasticity of amenity costs η_a—make relatively little difference for our estimates of productivity. We also find that the different values of the cost share of amenities make no difference at all for our estimates of amenities, which we discuss below. A higher cost share of amenities leads to a higher estimated average productivity, with the biggest difference in mean productivity from our baseline being less than 14 percent when the cost share of amenity doubles, a case we do not believe to be empirically relevant. The variance of gender wedges is slightly more affected in relative terms, but the impact remains small. All sets of estimates remain strongly correlated across estimation sets. We take this as evidence that our results are not strongly influenced by the specific choice of target for the economy-wide cost share of amenities.

It is important to note that the distribution of amenities remains exactly unchanged across estimation sets for both genders, as the elasticity of the amenity creation cost is not used in the estimation of profits per worker, which determine estimated utility and in turn estimated amenities. Intuitively, the elasticity of the amenity cost function does not affect our estimation of composite productivity $\tilde{p}$ and thus does not affect our estimation of total compensation x and amenity valuations βa.[75] To confirm this intuition, after re-estimating our model under alternative targets for the amenity share, we assess the robustness of our Kitagawa-Oaxaca-Blinder decomposition of pay, amenities, and total

[75]The elasticity of the amenity cost function matters only for the split of composite productivity $\tilde{p}$ into productivity p and the net surplus due to amenity valuations net of the amenity cost $\beta a - c^a(a)$.

Table E.7. Robustness: Estimation results under different assumed cost shares of amenities

Assumed cost share value	0.08 (Baseline)	0.04	0.12	0.16
η_a	5.728	11.457	3.819	2.864
η_v	2.063	2.063	2.063	2.063
Mean of p	1.529	1.423	1.634	1.740
Variance of p	0.245	0.257	0.235	0.226
Mean of a_m	1.211	1.211	1.211	1.211
Variance of a_m	0.071	0.071	0.071	0.071
Mean of a_f	1.172	1.172	1.172	1.172
Variance of a_f	0.062	0.062	0.062	0.062
Mean of τ	0.068	0.069	0.067	0.066
Variance of τ	0.014	0.017	0.012	0.011
Correlation between p_{baseline} and alternative		0.999	0.999	0.996
Correlation between $a_{m,\text{baseline}}$ and alternative		1.000	1.000	1.000
Correlation between $a_{f,\text{baseline}}$ and alternative		1.000	1.000	1.000
Correlation between τ_{baseline} and alternative		0.998	0.998	0.992

Notes: This table shows summary statistics for the distributions of estimated productivity p, amenities a_m and a_f, gender wedge τ and economy-wide parameters η_a and η_v, depending on the assumed value of the cost share of amenities. *Source:* Model estimates based on RAIS, 2007–2014.

compensation. Table E.8 shows the results from this decomposition based on half of the baseline's amenity share target (i.e., 4 percent). The results are identical to those in our baseline.

Table E.8. Kitagawa-Oaxaca-Blinder decompositions—half baseline's amenity share target

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.089	78.7	0.024	21.3
Amenity-valuation	−0.073	−0.087	118.4	0.014	−18.4
Total compensation	0.040	0.002	5.3	0.038	94.7

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) when re-estimating our model with half the baseline's amenity share target (i.e., 0.04). *Source:* Model estimates based on RAIS, 2007–2014.

Analogously, Table E.9 shows the results from this decomposition based on double the baseline's amenity share target (i.e., 16 percent). Again, the results are identical to those in our baseline.

Altogether, this gives us confidence that—although there is significant uncertainty around the estimate of the amenity cost elasticity η_a based on the assumed amenity cost share—our main conclusions are robust to its exact choice.

Table E.9. Kitagawa-Oaxaca-Blinder decompositions—double baseline's amenity share target

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.089	78.7	0.024	21.3
Amenity-valuation	−0.073	−0.087	118.4	0.014	−18.4
Total compensation	0.040	0.002	5.3	0.038	94.7

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) when re-estimating our model with double the baseline's amenity share target (i.e., 0.16). *Source:* Model estimates based on RAIS, 2007–2014.

F Gender-Specific Compensation Structures Across Employers Appendix

F.1 Alternative Kitagawa-Oaxaca-Blinder Decompositions

Recall that in Table B.1 of Appendix B.2, we presented alternative Kitagawa-Oaxaca-Blinder decompositions of the gender log pay gap. Here, we present analogous decompositions for the gender gaps in amenities (Table F.1) and utility (Table F.2).

Table F.1. Alternative Kitagawa-Oaxaca-Blinder decompositions of the gender gap in amenities

		Between-employer gap		Within-employer gap	
	Gender log amenities gap	Level	Share (%)	Level	Share (%)
Decomposition 1	−0.026	−0.092	348.6	0.065	−248.6
Decomposition 2	−0.026	−0.118	447.2	0.091	−347.2

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the overall gender log amenities gap into a between-employer log amenities gap and a within-employer log amenities gap. Decomposition 1 corresponds to equation (B.2) and uses women's estimates of log amenities for computing the between-employer component. Decomposition 2 corresponds to equation (B.3) and uses men's estimates of log amenities for computing the between-employer component. *Source:* Model estimates based on RAIS, 2007–2014.

Table F.2. Alternative Kitagawa-Oaxaca-Blinder decompositions of the gender gap in utility

		Between-employer gap		Within-employer gap	
	Gender log utility gap	Level	Share (%)	Level	Share (%)
Decomposition 1	0.046	0.002	4.6	0.044	95.4
Decomposition 2	0.046	−0.013	−28.5	0.059	128.5

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the overall gender log utility gap into a between-employer log utility gap and a within-employer log utility gap. Decomposition 1 corresponds to equation (B.2) and uses women's estimates of log utility for computing the between-employer component. Decomposition 2 corresponds to equation (B.3) and uses men's estimates of log utility for computing the between-employer component. *Source:* Model estimates based on RAIS, 2007–2014.

F.2 Different Margins of Gender Discrimination

In classical theories of taste-based discrimination (e.g., Becker, 1971), employers receive disutility from employing certain population groups, which affects the employer's recruiting decisions vis-à-vis these groups. In our framework, the gender wedge τ plays precisely this role. In addition, our framework features frictional pay dispersion across employers and allows employers to choose amenities separately for each gender.[76] This highlights amenities as a novel margin of employer "discrimination" (i.e., unequal treatment) across genders. Under the hypothesis that employers use amenities to differentiate between workers, one would expect gender wedges and women's amenity values a_F to be negatively related. Indeed, we confirm such a negative relationship, shown in Figure F.1.

F.3 The Role of the Public Sector

To investigate the role of the public sector in our analysis, we have classified each employer in our sample into public vs. private. We then complement Figure 6 from the main text by computing the

[76]We have in mind the costly provision of amenities such as health insurance (Dey and Flinn, 2005), job security (Jarosch, 2023), workplace safety (Lavetti and Schmutte, 2018), and protection against sexual harassment (Folke and Rickne, 2022).

Figure F.1. Negative relation between women's amenities and gender wedges

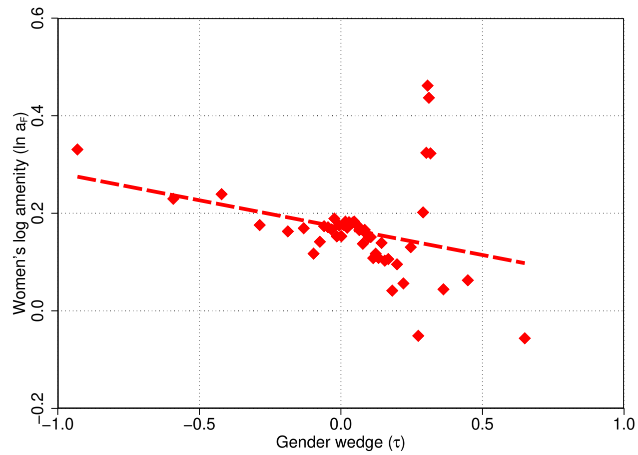

Note: This figure shows a binned scatter plot of women's log amenities ($\ln a_F$) against gender wedges (τ), with the linear best fit line in dashed red being weighted by female employment. *Source:* Model estimates based on RAIS, 2007–2014.

shares of employment in the public sector across firm ranks separately by gender. The results are shown in Figure F.2 below. From this, see that the public sector plays an important role in accounting

Figure F.2. Share of employment in public sector across firm ranks, by gender

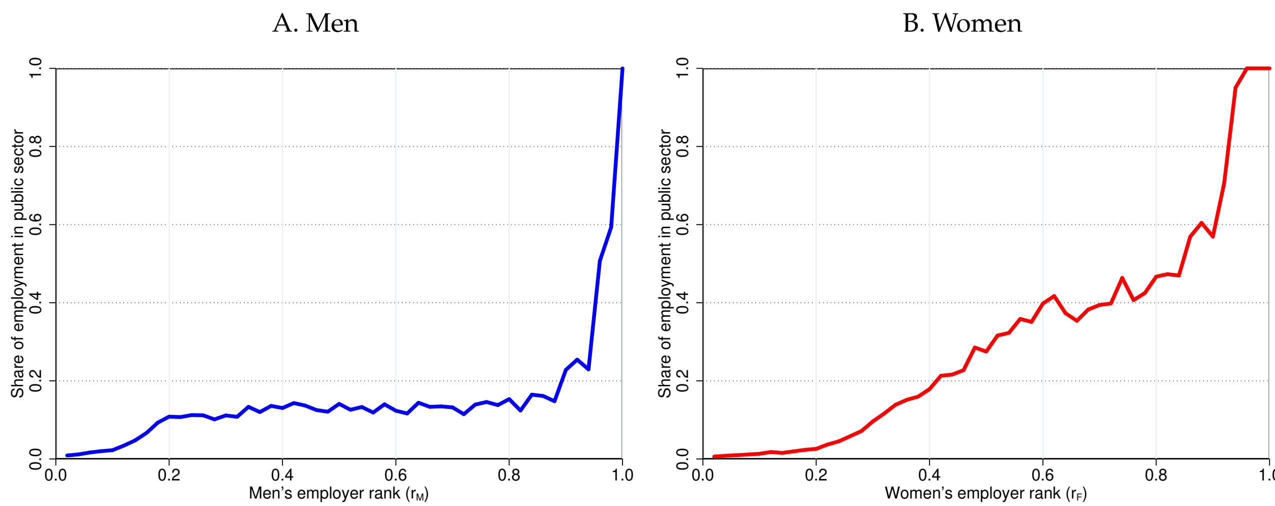

Note: This figure plots the share of employment in the public sector across firm ranks, separately for men in Panel A and for women in Panel B.*Source:* Model estimates based on RAIS, 2007–2014.

for high-ranked employers, and more so for women. Note that the top 4% of employment for women and the top 2% of employment for men are all concentrated in the public sector. Below those thresholds, public- and private-sector firms are mixed. In other words, 96% of the firm ladder for women and 98% of the firm ladder for men feature a combination of public and private-sector employers. This result echoes the findings from Section 7.1 and, in particular, Figure 3 of the main text.

These findings suggest that the public sector clearly plays an important role, complementing existing academic studies on the role of public-sector employment in relation to gender in the labor market. As such, we investigate to what extent our results are shaped by the public vs. private sector distinction by recomputing the same Kitagawa-Oaxca-Blinder decompositions separately for the whole economy, the private sector only, and the public sector only. Table 11 in the main text showed that amenities played an important role in understanding gender gaps in firm pay in the overall economy. Table F.3 below estimates the same Kitagawa-Oaxaca-Blinder decompositions on the subset of workers and firms in Brazil's private sector. We find a baseline gender pay gap of 11.3 log points (i.e., the same as in the whole economy), an amenity-valuation gap of −6.0 log points (i.e., the same sign and only slightly smaller magnitude than in the whole economy), and a total-compensation gap of 5.3 log points (i.e., similar to the whole economy). The split into between vs. within components is qualitatively unchanged.

Table F.3. Kitagawa-Oaxaca-Blinder decompositions—private sector only

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.113	0.074	65.7	0.039	34.3
Amenity-valuation	−0.060	−0.062	104.4	0.003	−4.4
Total compensation	0.053	0.012	22.3	0.041	77.7

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) for the private sector only. *Source:* Model estimates based on RAIS, 2007–2014.

Table F.4 repeats the same decomposition for the subset of workers and firms in Brazil's public sector. We find a baseline gender pay gap of 9.1 log points (i.e., a bit smaller than in the whole economy), an amenity-valuation gap of −3.6 log points (i.e., a bit smaller in magnitude than in the whole economy), and a total-compensation gap of 5.5 log points (i.e., a bit larger than in the whole economy). The split into between vs. within components is qualitatively unchanged with the exception of the amenity-valuation gap, which is more concentrated between rather than within employers in Brazil's public sector. These results are intuitive since we naturally expect pay gaps within the same public-sector organizations to be more muted by design. The public sector is also commonly viewed as a high-amenity sector, especially for women, rationalizing the larger within-employer gap in amenity valuations.

Table F.4. Kitagawa-Oaxaca-Blinder decompositions—public sector only

		Between-employer gap		Within-employer gap	
	Gender gap	Level	Share (%)	Level	Share (%)
Pay	0.091	0.099	109.1	−0.008	−9.1
Amenity-valuation	−0.036	−0.094	263.3	0.059	−163.3
Total compensation	0.055	0.005	8.7	0.050	91.3

Note: This table shows results from the Kitagawa-Oaxaca-Blinder decomposition of the gender gaps in log pay ($\ln w$), log amenity valuations ($\ln \tilde{a} = \ln(x/w)$), and log total compensation ($\ln x$) for the public sector only. *Source:* Model estimates based on RAIS, 2007–2014.

F.4 Implications for Productivity

Panel A of Figure F.3 shows the mean productivity p at different rungs of men's and women's firm ladders. Productivity is more steeply increasing in employer ranks for men than for women. The differences are meaningful. For example, men's median-ranked employer is 35 log points more productive than bottom-ranked employers, while the same statistic is only 21 log points for women. Thus, for men more so than for women, improvements in labor market efficiency yield productivity gains. Panel B plots women's productivity net of the gender wedge ($(1-\tau)p$) against men's productivity (p).

Among the least productive firms, gender wedges are close to zero but they steadily increase toward higher productivity levels. Again, the magnitudes are noteworthy. Near the top of the productivity distribution, the average gender wedge accounts for up to 50 log points of productivity.

Figure F.3. Productivity, ranks, and gender wedges

A. Productivity across ranks, by gender

B. Incidence of gender wedges

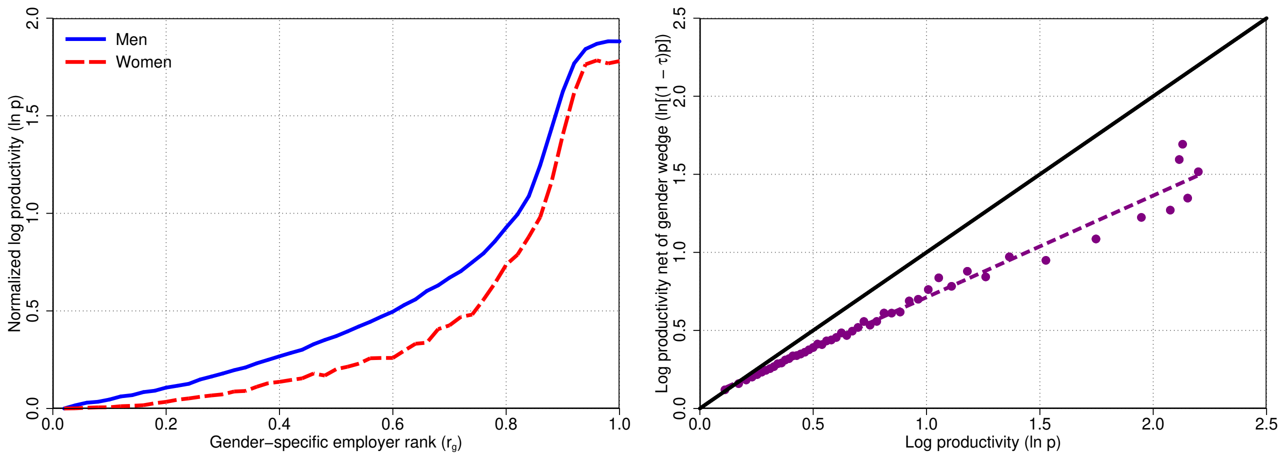

Note: Panel A of this figure shows gender-specific employment-weighted log productivity (ln p) across employer ranks (r_g) separately by gender $g \in \{M, F\}$. Panel B shows a binned scatter plot of women's productivity net of gender wedges ($\ln[(1-\tau)p]$) against men's log productivity (ln p), with the linear best fit line in dashed purple being weighted by total (i.e., male plus female) employment. *Source:* Model estimates based on RAIS, 2007–2014.

While not targeted, our model generates a lower elasticity of pay with respect to productivity for women (0.094) than for men (0.174). In our model, rent sharing is determined endogenously due to the combination of three gender-specific factors that we separately identify: compensating differentials (Rosen, 1986), taste-based discrimination (Becker, 1971), and labor market frictions (Manning, 2003).

F.5 Switching Employment and Compensation Policies Across Genders

Next, we relocate all workers of one gender into the other gender's employers, while keeping compensation policies constant, and vice versa. In an accounting sense, we then ask whether either men or women would prefer the other gender's employment distribution or compensation policies.

We first shift men's employment distribution to that of women: $l_M \mapsto l_F$, the results of which are presented in Panels A–C of Figure F.4. In the baseline, men's CDF of employment is represented by the 45-degree line in each panel. After moving men into women's firms, the solid blue line indicates the simulated CDF of employment. Such a move reduces men's welfare by 0.2 log points overall, consisting of a 8.9 log points loss in pay and a 9.2 log points gain in amenities. Interestingly, workers in the top 15% of the utility distribution actually prefer women's employment distribution. If, instead, we keep men's employment constant but change firm pay and amenities to those for women, men's total compensation decreases by 5.9 log points, consisting of a decrease in amenities by 9.1 log points and a decrease in pay by 4.0 log points. This experiment demonstrates that there are large differences in the treatment of men and women within employers.

Next, we shift women's employment distribution to that of men: $l_F \mapsto l_M$, the results of which are presented in Panels D–F of Figure F.4. As a result of this shift, women's welfare decreases by 1.3 log points, consisting of a 7.3 log points increase in pay but a 11.8 log points decrease in amenities. If, instead, we keep women's employment constant but change firm pay and amenities to those of men, women's amenities increase by 6.5 log points and their pay increases by 2.4 log points. As a result, women's total compensation increases by 4.4 log points, accounting for almost all of the gender utility gap, which again highlights the importance of unequal treatment within employers.

Figure F.4. Outcomes associated with moving men into women's employers and vice versa

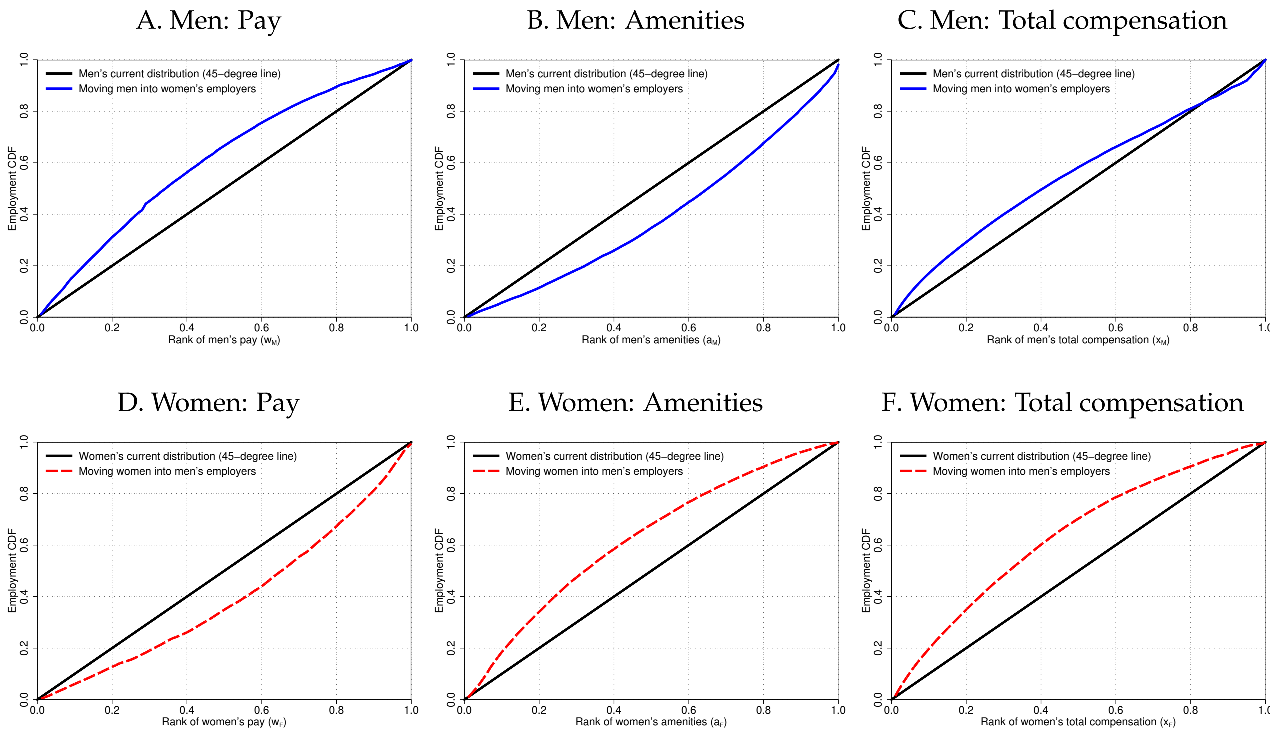

Note: This figure shows men's CDF over pay (w_M) in Panel A, amenities (a_M) in Panel B, and total compensation (x_M) in Panel C in the baseline as the diagonal solid black line and after moving men into women's employers as the solid blue line. Analogously, it shows women's CDF over pay (w_F) in Panel D, amenities (a_F) in Panel E, and total compensation (x_F) in Panel F in the baseline as the diagonal solid black line and after moving women into men's employers as the dashed red line. *Source:* Model estimates based on RAIS, 2007–2014.

G Equilibrium Counterfactuals Appendix

G.1 Numerical Solution Algorithm to Solve Baseline Equilibrium

Firstly, we feed 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 $\{p, a_M, a_F, \tau, c_M^{v,0}, c_F^{v,0}\}$. Then, we rank firms according to composite productivity $\tilde{p}_g$ for each gender. This is useful because, as stated in Lemma 3, firms with higher composite productivity $\tilde{p}_g$ offer higher utility x_g.

We must first find the equilibrium level of aggregate vacancies V_g. We invert the equation for the offer arrival rate from unemployment in (14) to obtain:

$$V_g = U_g \left(\frac{\lambda_g^U}{\chi_g} \right)^{1/\alpha}. \tag{G.1}$$

We recursively apply the discrete version of the FOCs of the firm, as in equations (D.58) and (D.61):

$$\Delta F_g(x(\tilde{p}_{gr})) = \left[\frac{1}{c_g^{v,0}} \frac{T_g(\tilde{p}_{gr-1} - x(\tilde{p}_{gr-1}))}{\left(\delta_g + \lambda_g^G + \lambda_g^E(1 - F_g(x_g(\tilde{p}_{gr-1})))\right)^2} \right]^{\frac{1}{\eta^v - 1}} \frac{1}{V_g N} \tag{G.2}$$

$$\Delta x_g(\tilde{p}_{gr}) = \frac{2\lambda_g^E(\tilde{p}_{gr-1} - x_g(\tilde{p}_{gr-1}))}{\delta_g + \lambda_g^G + \lambda_g^E(1 - F_g(x_g(\tilde{p}_{gr-1})))} \left[\frac{1}{c_g^{v,0}} \frac{T_g(\tilde{p}_{gr-1} - x_g(\tilde{p}_{gr-1}))}{\left(\delta_g + \lambda_g^G + \lambda_g^E(1 - F_g(x_g(\tilde{p}_{gr-1})))\right)^2} \right]^{\frac{1}{\eta^v - 1}} \frac{1}{V_g N}, \tag{G.3}$$

where N is the total number of firms in our data. Then, we calculate total vacancies obtained in equilibrium as $V_g^* = \sum_r v_g(\tilde{p}_{gr})/N$. We solve the algorithm setting the initial conditions $x_g(\tilde{p}_{g0}) = \phi_g$ and $F_g(\tilde{p}_{g0}) = 0$, and we loop over the vacancy cost shifter $c_g^{v,0}$ until we obtain that $V_g^* = V_g$.[77] In the baseline simulation in which we plug in the parameter estimates from the data, our solution algorithm produces gender-specific firm-level flow utility $x_g(\tilde{p}_{gr})$, easily converted to wages $w_g(\tilde{p}_{gr}) = x_g(\tilde{p}_{gr}) - a_g(\tilde{p}_{gr})$, gender-specific firm-level recruiting intensities $v_g(\tilde{p}_{gr})$ and gender-specific firm-level ranks in the offer distribution $F_g(x_g(\tilde{p}_{gr}))$ that are identical to those observed in the data, up to machine precision.

In all counterfactual simulations, we keep constant all firm-level parameters and the economy-wide parameters that are not explicitly mentioned as modified in the exercise, including the gender-specific cost shifters $c_g^{v,0}$. Also, we recalculate the workers' outside option in the counterfactuals by finding a new solution to equation (4) by gender. This allows the equilibrium job-finding probability λ_g^U, and therefore equilibrium employment, to respond to changes in the economy in counterfactual simulations.

G.2 Numerical Solution Algorithm to Solve Policy Counterfactuals

When we simulate the equal pay policy and the equal hiring policy, we can no longer rely on the model prediction that firms with higher composite productivity $(1 - \tau_g)p + a_g - c_g^a(a_g)$ will post higher flow utility and vacancies separately by submarket. 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,

[77] Alternatively, we could plug in the cost shifter we have estimated in Step 5 of our identification strategy, but this approach is equivalent in the baseline solution and allows us to generalize to cases in which we set λ_g^U to a value that is different from the one we estimated in the data.

we solve the following firm profit-maximization problem for the equal pay policy:

$$\max_{w,a_M,a_F,v_M,v_F} \quad \Bigg\{ T_M v_M (p - w - c^a_M(a_M)) \left(\frac{1}{\delta_M + \lambda^G_M + \lambda^E_M (1 - F_M(w + a_M))} \right)^2 \tag{G.4}$$

$$+ T_F v_F ((1-\tau)p - w - c^a_F(a_F)) \left(\frac{1}{\delta_F + \lambda^G_F + \lambda^E_g (1 - F_F(w + a_F))} \right)^2 \tag{G.5}$$

$$- c^v_M(v_M) - c^v_F(v_F) \Bigg\}, \tag{G.6}$$

and the following for the equal hiring policy:

$$\max_{x_M,x_F,a_M,a_F,v} \quad \Bigg\{ T_M v (\tilde{p}_M - x_M) \left(\frac{1}{\delta_M + \lambda^G_M + \lambda^E_M (1 - F_M(x_M))} \right)^2 \tag{G.7}$$

$$+ T_F v (\tilde{p}_F - x_F) \left(\frac{1}{\delta_F + \lambda^G_F + \lambda^E_g (1 - F_F(x_F))} \right)^2 \tag{G.8}$$

$$- c^v_M(v) - c^v_F(v) \Bigg\}, \tag{G.9}$$

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 counterfactual policy equilibrium. In the equal-pay policy, firms can still hire both genders, only one gender, or neither. In the equal hiring policy, dual-gender firms are forced to post the same number of vacancies across genders.

Denote by $\mathcal{F}_g$ the mapping from $\{F_M, F_F\}$ to the offer distribution for gender $g \in \{M, F\}$ implied by firms' optimizing behavior. We solve the following system of functional equations:

$$\mathcal{F}_M(F^*_M, F^*_F) = F^*_M \tag{G.10}$$

$$\mathcal{F}_F(F^*_M, F^*_F) = F^*_F, \tag{G.11}$$

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. In the equal pay policy case, this is 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. In the equal hiring policy case, this is because when firms decide how many vacancies to post, they balance the profits per contacted worker of both genders.

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 (12), 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.

G.3 Other Counterfactual Equilibrium Simulations

We expand the analysis in Section 8 by performing three more structural counterfactuals that illustrate the relative contributions of taste-based discrimination (Becker, 1971), labor market frictions (Manning, 2003), and their interplay. First, we remove employer heterogeneity in gender wedges. Second, we remove gender differences in labor market efficiency. Third, we remove all gender differences. Table G.1 summarizes our results, alongside the counterfactual we performed in Section 8 where we removed amenity differences across employers.

Table G.1. Structural decomposition of the gender pay gap

	Baseline	Counterfactuals			
Differences in	(0)	(1)	(2)	(3)	(4)
Amenities	✓		✓	✓	
Gender wedges	✓	✓		✓	
Labor market parameters	✓	✓	✓		
Gender log pay gap	0.109	0.057	−0.040	0.106	0.000
between employers	0.082	0.020	−0.029	0.082	0.000
within employers	0.027	0.037	−0.011	0.025	0.000
Gender log amenities gap	−0.066	−0.011	0.049	−0.065	0.000
between employers	−0.075	−0.010	0.032	−0.075	0.000
within employers	0.009	−0.002	0.017	0.010	−0.000
Gender log utility gap	0.042	0.046	0.009	0.041	−0.000
between employers	0.007	0.010	0.003	0.007	−0.000
within employers	0.035	0.035	0.006	0.034	0.000
Output	1.000	1.016	1.078	0.993	1.049
Worker welfare	1.000	1.014	1.013	1.000	1.015
for men	1.000	1.014	1.000	1.000	1.000
for women	1.000	1.013	1.032	1.000	1.039
Total employment	0.771	0.783	0.786	0.764	0.764
for men	0.764	0.772	0.764	0.764	0.764
for women	0.781	0.799	0.819	0.764	0.764

Note: This table reports results from equilibrium counterfactuals. Log amenities are defined as the log amenity valuations ($\ln \tilde{a} = \ln(x/w)$). The baseline economy (column 0) is compared to counterfactuals without differences in amenities across employers (column 1), without differences in gender wedges across employers (column 2), without gender differences in labor market efficiency (column 3), and without any gender differences in the economy (column 4). *Source:* Model estimates based on RAIS, 2007–2014.

The Role of Employer Preferences over Gender. What would gender inequality be if there were no heterogeneity in gender wedges? We set gender wedges at all firms to the average gender wedge, $\tau \mapsto \overline{\tau}$. As a result, the gender pay gap disappears completely, largely because women relocate toward higher-productivity employers who initially have higher gender wedges and, once removed, compete more for women by increasing their recruiting intensity and pay. As a result, the between-employer pay gap closes and output increases by 7.8 percent. However, more productive employers also offer lower amenities to women leading to an increase in the gender amenities gap by 11.5 log points. Overall, the gender utility gap is reduced by 3.4 out of the original 4.2 log points. While worker welfare of men is unchanged, worker welfare of women increases by 3.2 percent. Taken together, this counterfactual makes clear that amenities interact with gender wedges: low-paying low-productivity firms require both higher amenities and lower gender wedges in order to attract a substantial female workforce.

The Role of Gender-Specific Labor Market Frictions Power. What is the role of gender-specific labor market frictions in shaping gender inequality? We set the parameters guiding women's labor market efficiency so that $\lambda_F^U \mapsto \lambda_M^U$, $s_F^E \mapsto s_M^E$, $s_F^G \mapsto s_M^G$, and $\delta_F \mapsto \delta_M$. As a result, the gender pay gap declines by only 0.2 log points and the gender utility gap by 0.1 log points. The reason behind the small effect is that women receive more on-the-job offers but also separate more frequently from their employers. Thus, their overall speed of climbing the firm ladder barely changes. All in all, while there is large monopsony power in Brazil, differences in labor market fluidity are not a key driver of gender gaps.[78]

The Effects of Moving To a Gender-Neutral Labor Market. What would a labor market with no gender differences in amenities (i.e., $a_F \mapsto a_M$), no gender wedges (i.e., $\tau = 0$), and no gender differences in labor market efficiency (i.e., $\lambda_F^U \mapsto \lambda_M^U$, $s_F^E \mapsto s_M^E$, $s_F^G \mapsto s_M^G$, and $\delta_F \mapsto \delta_M$) look like? Obviously, this closes the gender gaps in pay and amenities. Strikingly, this also increases aggregate output by 4.9 percent and overall worker welfare by 1.5 percent, as women experience greater flow utility by 4.2 log points associated with them moving to more productive firms with higher total compensation. Relative to the counterfactuals that shut down one ingredient at a time, we find strong nonlinearities that reflect the interaction between different dimensions of firm heterogeneity affecting pay, amenities, and employment in equilibrium. For example, changing the wage-amenity composition of total compensation at a firm initially employing few women has a greater impact when also increasing the same firm's composite productivity in a way that leads to more women being hired. In this sense, our equilibrium model features strong complementarities between the different dimensions of firm heterogeneity. In summary, moving to a gender-neutral labor market yields significant output and welfare gains.

[78] This contrasts with Bowlus (1997) and Flinn et al. (2023) who attribute a significant fraction of the U.S. gender wage gap to different labor market behaviors across the sexes using an equilibrium search model *without* compensating differentials.

G.4 Counterfactual Simulations for the Equal-Amenities Policy

We simulate a mandate for all firms to provide equal amenity valuations across workers of both genders: $\beta_M(j)a_M(j) = \beta_F(j)a_F(j)$ for all j. The results from this counterfactual policy simulation are summarized in Table G.2 below.

Table G.2. Effects of simulated equal amenities policy.

	Baseline	Equal-amenities policy
	(0)	(1)
Gender log pay gap	0.109	0.095
between employers	0.082	0.030
within employers	0.027	0.065
Gender total amenities gap	−0.031	−0.040
between employers	−0.084	−0.040
within employers	0.052	0.000
Gender log utility gap	0.042	0.029
between employers	0.007	−0.001
within employers	0.035	0.030
Output	1.000	0.979
Worker welfare	1.000	0.989
for men	1.000	0.990
for women	1.000	0.987
Total employment	0.771	0.763
for men	0.764	0.760
for women	0.781	0.768

Note: Table reports results from a counterfactual policy experiments. Total amenities are defined as amenity valuations ($\beta_g a_g$). Baseline results (column 0) are compared against the economy with an equal-amenities policy (column 1). *Source:* Model estimates based on RAIS, 2007–2014.

We find that, when firms are mandated to offer the same amenities to men and women, both the gender pay gap and the gender utility gap decrease. As firms need now to compromise on the amenities they offer to men and women, those firms that were previously attractive to women due to high amenities become less so. Thus, women move to higher-paying firms, resulting in a lower between-firms gap but a larger within-firms gap, as higher-paying firms tend to be more discriminatory. We find this policy to be even more detrimental to welfare, output and employment than the equal-hiring policy. Output declines by 2.1 percent, while welfare declines by 1.1 percent, and more so for women.

G.5 Counterfactual Simulations for a Restricted Equal Hiring Policy

We simulate a scenario in which an equal hiring policy is imposed only on a subset $\mathbb{J}$ of firms: $v_M(j) = v_F(j)$ for $j \in \mathbb{J}$. The idea is to capture a policy that might be imposed, for instance, only on the public sector of the economy. To this end, we construct a dummy variable that takes value 1 if a firm is classified as operating in the public sector and zero otherwise, and we simulate the equilibrium of our model when only this subset of firms is constrained by the equal hiring policy. Our results are summarized in Table G.3.

We find that, when limited to the public sector, the equal hiring policy leads to a decline in the gender pay gap by 0.9 percent, exclusively driven by a decline in the between-firms gender pay gap. However, this is accompanied by a reduction of the amenities gap, which is in favour of women, by 0.9 percent. As a result, the overall utility gap remains unchanged. The reason is that the public sector already disproportionally hires women: imposing an equal hiring policy on the public sector increases the representation of men in these firms, while women relocate to other firms that offer higher wages but lower amenities.

Table G.3. Effects of simulated equal hiring policy, restricted to the public sector.

	Baseline	Equal-hiring	Equal-hiring, public only
	(0)	(1)	(2)
Gender log pay gap	0.109	0.034	0.100
between employers	0.082	0.006	0.073
within employers	0.027	0.028	0.027
Gender log amenities gap	−0.066	0.011	−0.057
between employers	−0.075	−0.006	−0.068
within employers	0.009	0.017	0.010
Gender log utility gap	0.042	0.045	0.042
between employers	0.007	0.000	0.005
within employers	0.035	0.045	0.037
Output	1.000	0.997	0.999
Worker welfare	1.000	0.992	0.999
for men	1.000	0.991	0.999
for women	1.000	0.993	0.999
Total employment	0.771	0.764	0.771
for men	0.764	0.722	0.764
for women	0.781	0.825	0.782

Note: Table reports results from a counterfactual policy experiments. Log amenities are defined as the log amenity valuations ($\ln \tilde{a} = \ln(x/w)$). Baseline results (column 0) are compared against the economy with an equal-hiring policy (column 1) and an economy where the equal-hiring policy is enforced on the public sector only (column 1). *Source:* Model estimates based on RAIS, 2007–2014.