ARTICLE IN PRESS

Journal of Econometrics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Gender differences in sorting on wages and risk[☆]

Kurt Lavetti[a,*], Ian M. Schmutte[b]

[a] Ohio State University, Department of Economics and NBER, United States
[b] University of Georgia, Department of Economics, United States

ARTICLE INFO

Article history:
Received 22 May 2020
Received in revised form 28 May 2022
Accepted 10 June 2022
Available online xxxx

JEL classification:
J31
J16
J17

Keywords:
Wage differentials
Gender
Occupational sorting

ABSTRACT

We use Brazilian matched employer–employee data to provide new evidence on gender disparities in labor market sorting on wages and workplace safety. We show that women and men sort in highly disparate, but systematic, ways on the basis of physical risk, despite sorting almost identically on financial risk. To understand what factors might explain these gender differences in sorting, we begin by ruling out the possibility that men and women receive different compensating wage differentials for risk. We find that women do earn smaller establishment wage premia than men on average, but this difference is also unlikely to explain the sorting patterns. Despite having little direct effect on wages, sorting on safety may have large indirect effects on the wage gap through its outsized influence on the segregation of women and men across establishments. This segregation results in men earning substantially higher establishment wage premia, which explains 28% of the entire gender wage gap in Brazil.

1. Introduction

As gender disparities in education and other human capital measures have closed in recent decades, so too has the extent to which skills contribute to earnings differences between women and men (Blau and Kahn, 2017; Goldin, 2014). Yet this narrowing of the skills gap has had little impact on the segregation of men and women in the workforce. Holding participation and labor force attachment constant, women in many countries tend to work in lower-paying occupations (Cortes and Pan, 2018), industries (Goldin et al., 2017), and firms (Bayard et al., 2003; Card et al., 2016). This occupational segregation, along with accompanying differences in the ways in which workers are compensated in different jobs, now explains the majority of the gender gap in earnings (Blau and Kahn, 2017). It remains unclear why gender differences in labor market sorting have not converged along with the skills gap. Potential explanations include differences in preferences for job attributes, attitudes toward risk, and discriminatory constraints on employment opportunities.

We focus on risk as a particularly salient explanation. A robust experimental literature has shown that women tend to be more averse to risk than men (Bertrand, 2011). To date, only a few studies have directly measured the impact of risk on gender differences in job assignment. DeLeire and Levy (2004) use family structure as a proxy for changes in aversion to physical risk (the risk of injury or death), and argue that the risk of a fatal workplace injury can explain about one-fourth of occupational gender segregation in CPS data.[1] Wiswall and Zafar (2018) document differences in mens' and

☆ We are grateful for helpful comments from Martha Bailey, Tom DeLeire, Judy Hellerstein, Chris Moser, Larry Kahn, Claudio Labanca, and seminar participants at the University of Maryland and Monash University.

* Corresponding author.
E-mail addresses: lavetti.1@osu.edu (K. Lavetti), schmutte@uga.edu (I.M. Schmutte).

1 Bonin et al. (2007) also find that among German men, self-reported risk aversion is correlated with occupational sorting.

https://doi.org/10.1016/j.jeconom.2022.06.012
0304-4076/

Please cite this article as: K. Lavetti and I.M. Schmutte, Gender differences in sorting on wages and risk. Journal of Econometrics (2022), https://doi.org/10.1016/j.jeconom.2022.06.012.

womens' preferences for employment risk using a hypothetical choice experiment, and show that these experimentally-elicited measures of preferences can predict actual job choices. These studies rely largely on cross-sectional variation in job assignment, and are not designed to identify how physical or financial risks affect mobility patterns between jobs, nor how sorting on the basis of risk contributes to the segregation of men and women in the labor market and the gender gap in earnings.

We study the role of workplace safety in driving observed differences in labor market sorting of women and men. We provide a new approach to graphically characterizing job transition patterns in matched longitudinal employer–employee data using vector gradient fields. Applying this approach to data from Brazil, we show that sorting on physical risk is highly systematic, and roughly as salient as sorting on establishment wage premia, a topic that has received far greater attention in the literature. Interpreting these gradient fields from the perspective of a multidimensional job-ladder, in which workers increase utility by moving to jobs that offer higher pay or greater safety, the evidence suggests that men and women ascend job ladders in different directions, with men ascending primarily in the wage direction and women ascending more in the safety direction.

We evaluate several potential explanations to understand why these sorting patterns are so different. First, we evaluate whether the patterns are unique, or whether other job amenities generate similar gender disparities in sorting. Female-dominated jobs may tend to be associated with a bundle of amenities that differs from male-dominated bundle, and many of the correlated amenities in these bundles may appear to generate disparate sorting patterns. Although we observe a limited set of job amenities in our administrative data, we can benchmark the patterns we observe relative to sorting on measures of financial risk, such as within-job earnings volatility. Prior literature has documented that women have greater financial risk aversion, suggesting that jobs with high earnings risk may be more likely to be male-dominated. However, when we apply the same approach to characterize sorting on earnings volatility, we find that men and women sort almost identically. This suggests that the known gender differences in financial risk preferences are not strong enough to generate a clear disparity in labor market sorting, and in contrast physical risk, though not necessarily a causal factor, is at least a very salient proxy for a related set of factors (such as the risk of injury or illness) that drive the sorting of men and women across jobs.

A second potential explanation is that women may not be equally compensated for accepting jobs with higher risk. This hypothesis is implied by much of the prior literature on gender disparities in compensating wage differentials for workplace injuries (Hersch, 1998; Leeth and Ruser, 2003). To evaluate this potential explanation, we revisit the theory of hedonic wages and empirical methods for estimating gender differences in compensating wage differentials. According to hedonic wage theory, if women have stronger preferences for safety, they should also face a weakly higher marginal price for safety, and hence, a larger compensating wage differential than men (Rosen, 1974). Further, if men and women are substitutes in the labor market, then the equilibrium is determined by the preferences of a common "marginal worker", and so men and women will face the same compensating differential. Relative to this theory, much of the empirical literature on compensating wage differentials for fatal risk is puzzling, as is suggests that women face *lower* marginal prices of safety.

We replicate the gender gap in compensating differentials documented in prior work, and show that it can be explained by two important modeling and measurement issues. First, previous estimates do not account for the nonrandom assignment of workers to jobs on the basis of unobserved worker, establishment, and occupational characteristics, or allow this assignment process to itself be gender-specific. We build upon extensions of the model from Abowd et al. (1999) to relax these restrictions. Second, we calculate fatality risk separately for men and women in narrow industry by occupation by time cells, providing a level of detail not previously available.[2] After correcting these estimation and measurement issues, the gender gap in compensating wage differentials is reduced to a precisely-estimated zero in our setting. Since there are no gender differences in the implicit wage price of safety, this cannot explain the observed sorting patterns.

Third, we consider whether the sorting patterns can be explained by differences in establishment-level rents paid to men and women. Following Card et al. (2016), we consider the possibility that women may have more limited opportunities for upward wage mobility from sorting across firms on the basis of wage premia, or rents.[3] If so, movements toward safety might be a natural response by women to this difference in wage premia. We find that in Brazil, women do receive lower employer-specific rents, but on average differences in rents explain less than one percent of the raw gap in log wages. However, holding establishment assignment fixed, if women were paid the same rents as their male co-workers, their wages would actually *decrease* by 0.8 percentage points, and a similar pattern holds from the perspective of male workers.[4] This suggests that although differences in rents paid to women and men by the same employer are an important component of equilibrium wages, their impact on the earnings gap is secondary to that of sorting *across* employers.

Motivated by this finding, we show that sorting on workplace safety is connected to gender differences in the establishment component of compensation. In jobs (industry–occupation pairs) in which women are safer than men, there

[2] We follow Tsai et al. (2011) and Lalive (2003), who showed that using overly coarse industry–occupation summaries of fatality risk generate substantial aggregation bias when estimating compensating wage differentials.

[3] We follow Card et al. (2016) in using the term 'rents' to describe the firm wage premium estimated from the AKM decomposition.

[4] Specifically, if men were paid the same rents as their female co-workers, their wages would also decrease by 1.5 percentage points. See Appendix Table A.4.

is substantially greater gender segregation across establishments. A consequence of this segregation is that women end up employed in firms that on average pay lower wages to all workers, both female and male. We find that this differential sorting across establishments explains 28 percent of the total gender wage gap in Brazil, about twice the share found by Card et al. (2016) in Portugal.[5]

Our paper contributes to a growing literature on the role of workplace sorting in the propagation of the gender wage gap. Card et al. (2016) discuss, and discard, the possibility that observed differences in firm-specific pay for women and men in Portugal reflect compensating differentials for hours requirements. Hotz et al. (2017) study the contribution of family-friendly policies to career outcomes for women immediately after childbirth. We are among the first to focus on worker sorting on the basis of a specific, observable, job disamenity – workplace safety – that has been shown in prior work to be highly salient in women's labor market decisions (Hersch, 1998).

While our primary interest is in sorting across firms, our analysis also complements recent work on the role of sorting and segregation within firms. Survey evidence by Wiswall and Zafar (2018) suggests that women and men have very different preferences for workplace amenities, and that women are less inclined to sort into jobs based on earnings growth. Mas and Pallais (2017) use experimental variation to recover preferences for job amenities, and find important gender differences in the willingness-to-pay to avoid irregular work schedules. Goldin et al. (2017) show the gender gap is driven in part by differences in pay growth within firms, after controlling for sorting across firms. Merlino et al. (2018) find gender wage gaps may be driven by differences in the propensity for women to be promoted within the firm, while Bartolucci (2013) shows women have lower bargaining power than men, but are also more mobile. Hellerstein and Neumark (2008) showed that racial segregation within firms cannot be explained by educational differences, suggesting sorting could be driven either by worker or employer preferences. Unlike these studies, we are uniquely able to consider comprehensive patterns of labor market sorting, and to estimate the importance of sorting on job safety relative to wages.

2. Data and descriptive statistics

We use matched employer–employee data from Brazil's *Relação Anual de Informações Sociais* (RAIS) from 2003–2010 (Ministério do Trabalho e Emprego, 2017). The RAIS is a compulsory administrative census of every formal-sector job, conducted annually by the Brazilian Ministry of Labor and Employment for the purpose of administering tax and transfer programs. We use these data to measure earnings, workplace sorting, and occupational fatality rates.

2.1. Preparation of the estimation sample

The RAIS contains information reported by establishments for each employment contract that was active at any point during the year. Hence, the unit of observation is a contract-year. For each contract-year, RAIS reports the characteristics of the worker, the establishment, and the job. For each worker, the key variables are date of birth, race, education, and gender, along with a unique identifier. For each establishment we observe the industry, location, and number of employees. For each job we observe average monthly earnings, contracted weekly hours, occupation, dates of hire and separation, and the cause of separation. The data on cause of separation includes three detailed codes identifying work-related fatalities.

Starting with the raw RAIS data, we construct a worker-year panel focused on individuals aged 23–65 who are employed for at least 1000 h per year in at least one full-time job, defined as 30 or more contracted hours per week.[6] When a worker holds multiple jobs in the same year, we restrict attention to the job with the highest estimated annual earnings.[7] We exclude government jobs, and temporary employment contacts. As we describe below, to reduce error in measuring fatality rates, we also exclude jobs in 2-digit industry by 3-digit occupation cells that contain fewer than 10,000 full-time full-year-equivalent workers of either gender.

Following Card et al. (2016), we estimate establishment-specific wage effects separately for each gender. This is only possible in establishments where at least one worker of each gender either leaves to, or arrives from, another establishment. In the language of Abowd et al. (2002), we restrict analysis to establishments in the largest connected component of the realized mobility networks generated by both male job transitions and by female job transitions. For women and men, taken separately, the largest connected components contain 91% of observations in the attached dominant jobs sample. The intersection of these sets, the dual connected set, contains 8.2 million observations for women, or 91% of the largest female connected set, and 14.6 million observations for men, 72% of the largest male connected set.

The RAIS records whether each job ends due to a fatal work-related accident, and thus serves as a census of individual fatal occupational injuries.[8] We use this information to construct gender-specific fatality rates in each year for 11,440

5 Within establishments in Brazil, we show that women and men earn roughly similar wages conditional on their occupation, measured, and unmeasured characteristics.

6 Blau and Kahn (2017) show that part of the gender wage gap can be attributed to a persistent gender division of household labor, and resulting differences in labor force attachment. Our sample restrictions reduce the influence of gender division of labor, but focus on the subset of women with strong attachment.

7 This approach follows Abowd et al. (1999), Woodcock (2008), and Card et al. (2013). The data only report average monthly earnings, so the estimated annual earnings are equal to monthly earnings times the number of months employed.

8 The causes of separation that we use to indicate fatal occupation accidents translate approximately to "deceased, at work accident" and "deceased, work related illness".

Table 1
Summary statistics.

	Full Population		Attached Dominant jobs		Dual Connected set	
	Women	Men	Women	Men	Women	Men
Age	35.83	35.46	38.36	38.64	38.20	38.37
Race white	0.57	0.56	0.69	0.64	0.68	0.63
Less than high school	0.32	0.53	0.38	0.58	0.36	0.52
High school	0.41	0.33	0.40	0.33	0.39	0.36
Some college	0.05	0.03	0.06	0.03	0.06	0.04
College or more	0.21	0.10	0.16	0.07	0.18	0.09
Contracted weekly hours	39.46	42.03	42.05	43.04	41.63	42.79
Log hourly wage	1.33	1.42	1.44	1.61	1.54	1.72
Total experience (Years)	18.35	18.98	22.12	23.43	21.80	22.89
Job tenure (Months)	62.36	51.69	89.56	82.93	90.91	84.66
Pooled fatality rate (per 100,000)	0.02	0.07	0.02	0.08	0.02	0.07
Gender-specific fatality rate	0.01	0.08	0.01	0.08	0.01	0.08
Zero fatality rate	0.26	0.15	0.35	0.12	0.37	0.13
N	134,361,238	194,907,785	11,419,266	22,234,188	8,193,244	14,567,312

Notes: Attached dominant jobs sample includes jobs with the highest annual earnings for each worker between ages 23–65, with 30 or more contracted hours per week, excluding government jobs, temporary jobs, and jobs in 2-digit industry by 3-digit occupation cells that have fewer than 10,000 full-time full-year equivalent workers in the three-year moving average window used to calculate fatality rates. Attached workers are defined as those with a dominant job meeting this sample criteria in each year between 2005–2010. Fatality Rate is measured in deaths per 1,000 full-time full-year equivalent workers. Dual connected set includes the subset of attached dominant jobs at establishments in the intersection of the largest connected sets of establishments for men and women.

two-digit industry by three-digit occupation cells.[9] Since fatalities are uncommon, we follow the literature in combining annual data to construct three-year moving averages of the number of fatalities per 100,000 full-time full-year-equivalent workers. In contrast, prior research has generally been limited to industry- or occupation-specific risk measures that do not differentiate by gender.[10]

2.2. Descriptive statistics

Table 1 reports descriptive statistics by gender for the full population, the subset of dominant jobs among attached workers, and the dual connected set used as the main analysis sample (see Appendix Table A.1 for additional summary statistics.) Among female workers, relative to the population, workers in the attached jobs sample and dual connected set are slightly older, more likely to be white, less educated, have more experience and job tenure, higher hourly wages, and are more likely to be in jobs with zero observed fatalities. These differences are due primarily to our restriction in both samples to jobs in industry–occupation pairs with more than 10,000 full-time full-year-equivalent workers, which we impose to avoid imprecisely estimating fatality rates in industry–occupation cells with very few workers. The average hourly wage is about 18 (17) percent higher for men than for women in the dual connected set (attached dominant jobs sample), about twice as large as the 9 percent gender-wage gap in the full population.

When calculated using fatality rates that pool both genders, the average fatality rate for women is 2.6 deaths per 100,000 full-time full-year workers. The average job for a male worker is roughly three times as dangerous (7.3 deaths per 100,000 full-time full-year workers.) Using gender-pooled fatality rates introduces considerable measurement error, which, as we will show, has important implications for understanding gender differences in labor market sorting. We find that women are considerably safer than men even within narrowly-defined industry–occupation cells. Using a more precise gender-industry–occupation measure, we find that the average fatality rate across jobs held by women is only 1.2 deaths per 100,000 full-time full-year workers. In addition, women are much more likely than men to be employed in industry–occupation cells with zero observed fatalities—37 percent of observations in the female analysis sample are jobs with a measured fatality rate of zero, versus 13 percent for men.

3. Gender differences in workplace safety

Our granular measure of fatality risk reveals large differences in job safety between women and men.[11] Fig. 1 depicts the joint distribution of the male and female fatality rates by industry–occupation cell, where the sizes of circles are

[9] There is a change in industry classification systems during our study period. The RAIS codes industries using the CNAE95 (1.0) classification system for all years of our study (Classificação Brasileiro de Ocupações: Downloads - 5.1.0, 2020). Beginning in 2005, industry is also recorded in a separate variable using the CNAE 2.0 classification system. For 2003 and 2004, we convert the CNAE 1.0 codes to CNAE 2.0 codes using a crosswalk obtained from Brazil's IBGE (Classificação Nacional de Atividades Econômicas, 2018). Occupations are classified using the 2002 vintage CBO codes in every year of our analysis (Classificação Brasileiro de Ocupações: Downloads - 5.1.0, 2020).

[10] One notable exception is Hersch (1998), which uses gender-specific measures of injury risk within coarse occupation categories, but does not disaggregate by industry.

[11] Appendix Table A.2 provides summary statistics on industry and occupation-level average fatality rates for men and women.

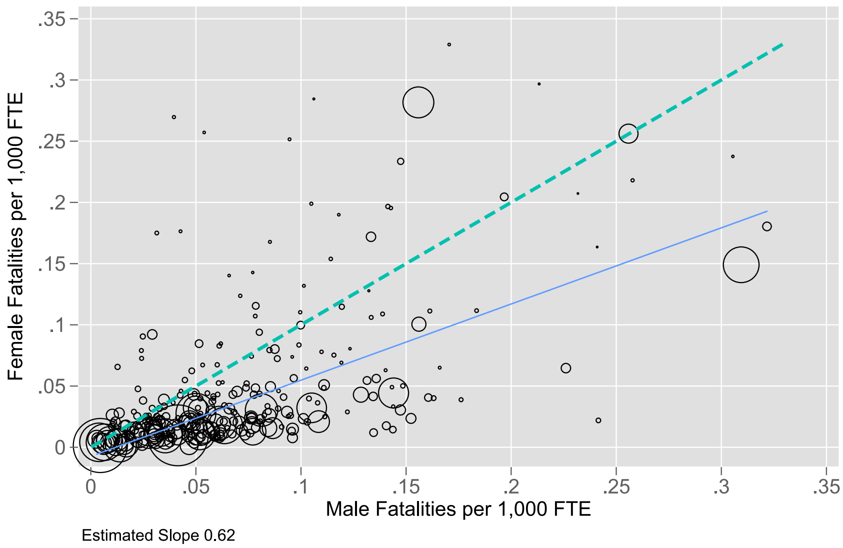

Fig. 1. Female versus male fatality rates by industry and occupation.
Notes: Fatality rates are gender-specific averages by 2-digit industry and 3-digit occupation pooled over all years from 2003–2010, measured in deaths per 1000 FTE worker-years. Figure includes cells with strictly positive estimated fatality rates and is truncated at the 99th percentile of fatality rates (0.33 fatalities per 1000 FTEs). Circle sizes are proportional to the total number of workers in the industry–occupation cell. Dashed line is 45-degree line.

proportional to the total number of hours worked in the cell. If there were no difference on average between the safety of men and women within a job-cell, the circles would lie near the dashed line in the figure. Instead, the figure shows that the female fatality rate is about 62% of the male rate in the same industry–occupation cell. These findings are similar to those of Hersch (1998), which shows in the U.S. women have much lower rates of illness and injury. Without task-level data we cannot directly test the underlying causes of this gender disparity, but the patterns are potentially consistent with women being more cautious than men, or with women and men performing different tasks even within narrowly-defined jobs. We revisit these potential explanations in Section 8.2.

Most prior research on gender differences in safety has been limited to industry- or occupation-level risk measures that do not differentiate by gender. In Brazil, we find that 91% of the total variation in female fatality rates occurs within 2-digit industries, and 89% of the total variation occurs within 3-digit occupations. This implies that in our empirical context, industry–occupation interactions explain the vast majority of the variation in female fatality rates. The nature of this measurement problem is itself gender-specific—among men only 68% of total variation occurs within 2-digit industry codes, and only 34% occurs within 3-digit occupations. Pooling men and women to calculate the fatality rate generates substantial measurement error that is strongly negatively correlated with the gender-pooled rate (Appendix Figure A.1).

In addition to being safer within industry–occupation cells, women also sort more strongly across cells on the basis of safety. This is consistent with well-documented gender differences in risk preferences generally (Blau and Kahn, 2017; Bertrand, 2011). Fig. 2(a) plots the female share of total hours worked in an industry–occupation cell against the log female fatality rate in that cell. Doubling the female fatality rate is associated with a 15.0% reduction in the female employment share, suggesting that women and men sort very differently on occupational safety.

We also find the growth of female employment over time is disproportionately concentrated in jobs that are safer (Appendix Figure A.2). During the period we study the female share of employment increased by about 2%, and this increase was driven by the entry of women into relatively safer jobs, with no net change in the female employment share in dangerous jobs.

4. Job-to-job sorting

To graphically characterize how women and men sort through the labor market, we construct vector gradient fields that depict changes in characteristics between origin and destination jobs. Fig. 3 shows how women and men sort on the basis of physical risk, measured by occupational fatality rates, and wages.

The vertical axis measures quintiles of the wage distribution. The horizontal axis measures the distribution of fatality risk, with a zero column indicating all jobs originating in cells with zero observed fatalities, and the remaining columns indicating the five quintiles of the fatality rate distribution conditional on the rate being nonzero. For each wage and risk quintile pair, we plot a vector that shows the average change in wages and risk observed when workers make job-to-job transitions originating from that quintile pair. The gradient vectors originating from each quintile pair are proportional

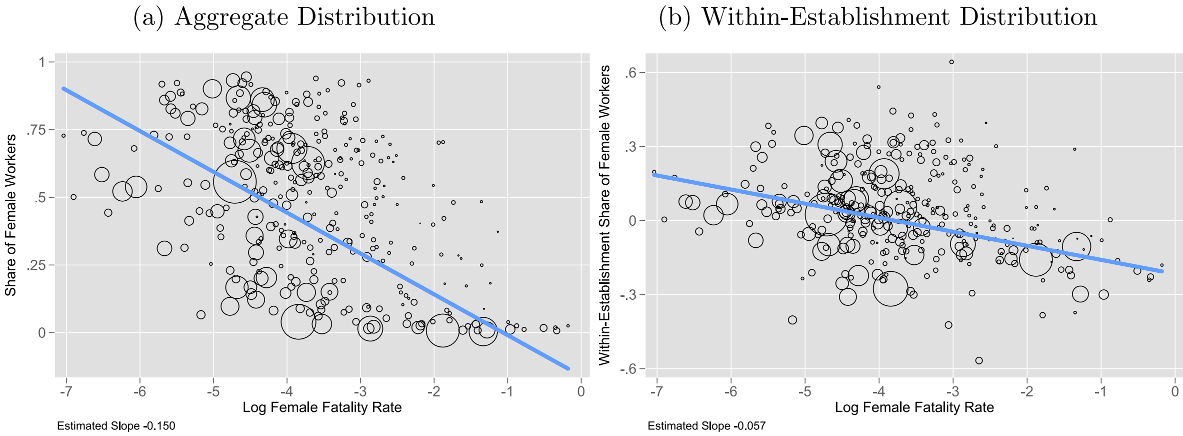

Fig. 2. Female worker shares versus log fatality rate.
Notes: Fatality rates are female-specific averages by 2-digit industry and 3-digit occupation pooled over all years from 2003–2010, measured in log of deaths per 1000 FTFY workers. The sample includes the full population of women with positive fatality rates, with fatality rates truncated at the 99th percentile for ease of presentation. Circle sizes are proportional to the total number of workers in the industry–occupation cell.

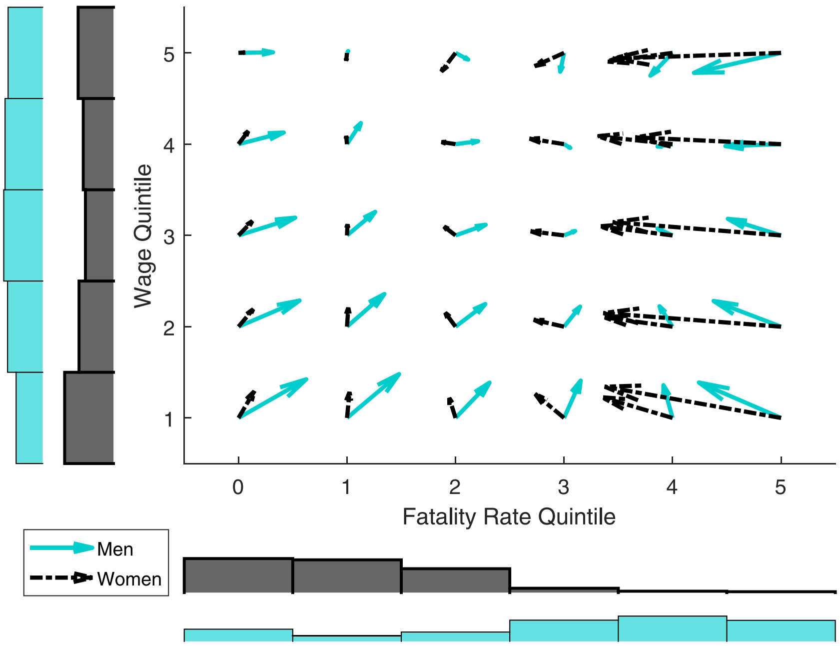

Fig. 3. Job-to-job transition gradient field: Physical risk.
Notes: Sample includes attached full-time women and men between ages 23–59 employed in the largest dual connected set of establishments, and excludes origin jobs with zero fatality rate. Quintiles defined based on combined female and male distributions, using gender-specific fatality rates. Vectors indicate directions and relative magnitudes of changes in wages and fatality rates associated with job-to-job transitions originating from each cell. Marginal density functions depicted at origin jobs.

in length to the magnitude of the change in fatality rates or wages, normalized by the respective standard deviation of changes in fatality rates or wages. For example, the figure shows that when women exit jobs in the first quintile of the wage distribution and third quintile of the fatality rate distribution, they tend to move up and to the left, indicating that the destination job is safer and pays a higher wage. In contrast, men originating from the same quintile pair move up and to the right, and the vertical component of the change is larger than that for women, implying that men exiting these jobs have larger wage gains than women, but also accept an increase in risk on average.

The gradient field shows that for nearly all origin points, when women change jobs they move very systematically toward the left – toward jobs with lower fatality rates – with little change in the wage dimension. Of course, when originating in jobs with zero observed fatalities, the average change in fatality risk is mechanically non-negative for both men and women. Even women originating in the first quintile of the fatality rate distribution move vertically, with very little increase in physical risk on average. In contrast, when men change jobs, risk increases when they originate at safer jobs, and decreases when they originate at riskier jobs—a pattern consistent with mean reversion. Moreover, men in low-wage jobs move noticeably further in the vertical direction, while women experience more modest wage increases. When originating in safe jobs, women's vectors have shorter length, indicating little change in characteristics regardless of the wage. In contrast, men originating in similar jobs move much further from their origin job, on average. The gradient fields reveal a stark gender difference in the way that men and women sort in the labor market on the basis of wages relative to physical risk. Interpreting these fields from the perspective of a multidimensional job-ladder, the patterns suggest that men tend to ascend more strongly in the wage dimension, while women ascending more in the safety direction.

This contrast in mobility patterns contributes to the differences in the marginal wage distributions plotted to the left of the graph. The histograms on the left show that women are most likely to originate from jobs in the lowest wage quintile, while men are relatively over-represented in higher wage quintiles. This difference in the wage distributions corresponds to an 18 percentage point average gender wage gap in Brazil. Similarly, the histograms below the horizontal axis show the marginal distributions of fatality rates for men and women. They indicate that women are concentrated in jobs with zero observed fatalities, or in the first and second deciles of fatality risk, while men are overrepresented in higher deciles.

This figure provides a concise way of graphically describing job mobility patterns. Of course, these vector gradients are only stylized summary statistics of origin and destination jobs, and should not be interpreted to imply that all workers of a given gender sort through the labor market in the same way, nor do they indicate a causal mechanism behind these sorting patterns. The patterns are, however, consistent with the literature suggesting women are more averse to physical risk than men (DeLeire and Levy, 2004).

One possible interpretation of the figure is that physical risk is only a proxy for a bundle of correlated job characteristics over which men and women have different preferences. For example, it was been well-documented in the literature that women have greater financial risk aversion than men (Bertrand, 2011). Based on these preference estimates, one expects to see that women sort away from jobs with greater earnings risk. If earnings risk and physical risk are positively correlated in the labor market, differences in financial risk preferences could potentially explain the sorting patterns in Fig. 3.

Fig. 4 shows the same job to job transition gradient field over wages and financial risk.[12] To measure financial risk, we first calculate monthly earnings in December divided by annual earnings for all full-time workers in the sample. We then calculate the standard deviation of this ratio at the establishment by three-digit occupation by year level. This measures fluctuations in earnings across workers within a narrowly defined group of jobs in the same establishment, and in the same year.

Fig. 4 shows that there is virtually no difference between men and women in Brazil in sorting on the basis of earnings volatility, suggesting that the sorting patterns in Fig. 3 are not picking up gender differences in financial risk aversion that may be correlated with physical risk. Correspondingly, the marginal distributions of earnings volatility shown below the figure are mostly overlapping, indicating very little aggregate gender difference in the distribution of earnings risk for men and women.

Although we cannot observe many specific job amenities in our administrative data, the contrasts between Figs. 3 and 4 are informative. Whereas researchers have documented clear differences in preferences for financial risk, these differences are not substantial enough to generate a disparity in job-to-job sorting patterns in the labor market. Although there are fewer directly elicited measures of gender-specific preferences over the risk of death, men and women clearly sort very differently with respect to fatal injury risk. The contrast between the figures suggests that there are meaningful differences between jobs that are related to these two measures of risk, as opposed to each measure being spuriously correlated with the same large bundle of job attributes. Still, each figure should be interpreted as depicting sorting patterns over the corresponding risk measure and correlated job characteristics.

5. Wage decomposition model

To understand whether these gender differences in sorting on physical risk are driven by differences in how men and women are compensated, we decompose wages into components associated with the price of risk and firm-specific wage premia. We use the empirical estimation approach proposed in Lavetti and Schmutte (2018), an extension of the earnings decomposition model from Abowd et al. (1999) that we refer to as the *orthogonal match effects* (OME) estimator. This two-step specification allows time-varying characteristics x_{it} to be arbitrarily correlated with unobserved worker, establishment, and job-match effects. We include a hedonic component in the model, allowing fatality rates to affect wages and to be correlated with latent worker and establishment effects. Controlling for establishment heterogeneity in wages matters because higher-paying employers tend to offer safer jobs, and workers frequently move toward jobs with better wages and better amenities. Importantly for our application, this specification also controls for all unobserved establishment-level job amenities (including, for example, average differences in injury rates, as well as the set of family-friendly workplace policies studied by Hotz et al. (2017)). The inclusion of establishment effects therefore mitigates a common source of omitted variable bias in estimates of compensating differentials.

[12] There is no zero column in this figure because the share of jobs with exactly zero measured earnings volatility is negligible. These jobs are excluded from the figure.

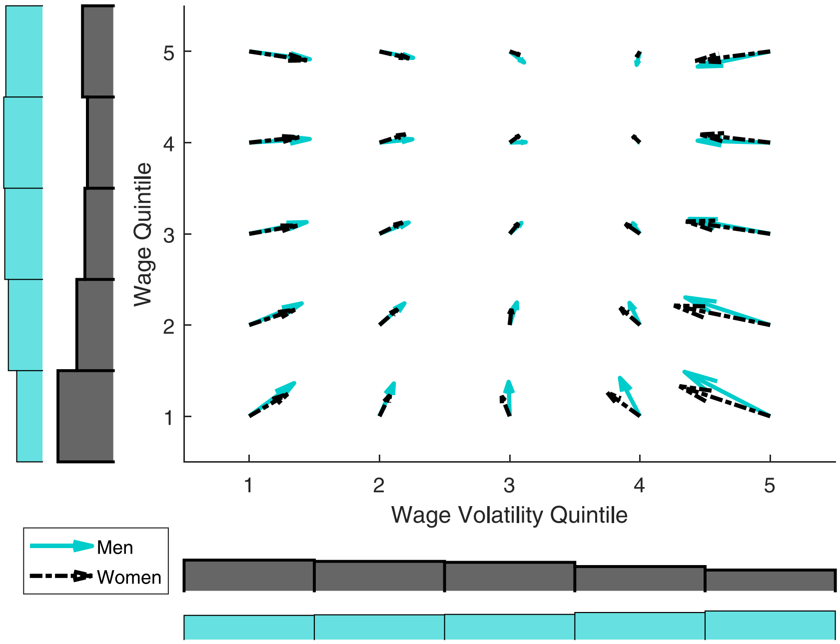

Fig. 4. Job-to-job transition gradient field: Financial risk.
Notes: Sample includes attached full-time women and men between ages 23–59 employed in the largest dual connected set of establishments, and excludes origin jobs with zero fatality rate. Quintiles defined based on combined female and male distributions, using gender-specific fatality rates. Vectors indicate directions and relative magnitudes of changes in wages and fatality rates associated with job-to-job transitions originating from each cell. Marginal density functions depicted at origin jobs.

5.1. Specification

In a first step of the model, we remove the effect of time-varying observables using within-match variation by estimating:

$$w_{it} = x_{it}\beta^g + \tilde{\gamma}^g a_{c(i,t),t} + \Phi_{i,Jk(i,t)} + \epsilon_{it} \tag{1}$$

where the dependent variable, w_{it}, is the log hourly wage of worker i in year t. The vector of control variables, x_{it}, includes a set of indicators for each year of labor market experience (censored at 30 years) fully interacted with gender, and year effects. The coefficient vector β^g denotes that each coefficient in β is free to vary by gender. $a_{c(i,t),t}$ is our gender-specific measure of the risk of fatal injury when the worker is employed in year t in the industry–occupation cell given by the function $c(i,t)$ which maps panel indices to cells. $\Phi_{i,Jk(i,t)}$ denotes the match effect between worker i and the Jk establishment-occupation pair at which worker i is employed in period t.

In the second stage, we estimate a decomposition of the components of variation in wages across jobs, while conditioning on the unobserved characteristics of establishments that may affect job mobility choices:

$$w_{it} - x_{it}\widehat{\beta^g} = \pi^g z_{it} + \gamma a_{c(i,t),t} + \gamma^f a_{c(i,t),t} * Female_i + \theta_i + \psi^g_{J(i,t)} + \varepsilon_{it} \tag{2}$$

This second stage is the two-way fixed effects decomposition popularized by Abowd et al. (1999). Variation in this model that contributes to the identification of γ comes from job changes in which a worker moves across industry–occupation cells, which could occur either within or across establishments. z_{it} includes gender-by-year effects and one-digit occupation code effects interacted with gender, to account for broad differences in job characteristics other than safety. In this model, the unobserved error depends on a component that is common to all observations for the same worker, θ_i, and a common establishment wage effect $\psi^g_{J(i,t)}$ that can vary arbitrarily by gender. The index function $J(i,t)$ maps panel observations of the job held by worker i in year t to establishment identifiers.

5.2. Establishment wage effects

Since we are interested in decomposing within-establishment gender wage gaps from the contribution to the mean gender wage gap of differential sorting across establishments, we have to ensure that level differences in $\widehat{\psi^g_{J(i,t)}}$ for men and women have a valid interpretation. In general, gender-specific estimates of $\psi^g_{J(i,t)}$ from the model described above cannot be interpreted as level differences in wages because the set of female-establishment cells is disconnected from the set of male-establishment cells. That is, we never observe a single worker transitioning from receiving the female treatment effect to receiving the male treatment effect. Without such a transition, under the standard normalization (Abowd et al., 2002), the gender-specific establishment effects have a mean of zero within each gender.

We follow Card et al. (2016) in assuming that $\mathbb{E}\left[\psi^g_{J(i,t)} = 0\right]$ in industries with the lowest average firm effects.[13] That is, we restrict the average firm effects to be zero in the industries with the lowest average firm effects, and this restriction is imposed separately for each gender. The theoretical intuition behind this assumption is that if establishment wage premia reflect rents paid to workers, then in perfectly competitive industries the wage premium levels should be zero for both men and women. By assuming that at least one industry is perfectly competitive, the minimum industry average value of $\psi^g_{J(i,t)} = 0$ for both men and women. This allows a direct comparison of the levels of $\psi^g_{J(i,t)}$ for men and women at all employers under the maintained model assumption that $\psi^g_{J(i,t)}$ is gender-specific.[14]

Under this assumption there is no restriction that average establishment effects must equal zero for each gender. To implement this restriction, we first estimate Eq. (2), calculate the average $\psi^g_{J(i,t)}$ from this model by two-digit industry, and identify the five industries with the lowest average establishment effects. Using only observations from these five industries, we calculate the gender-specific average values of $\psi^g_{J(i,t)}$, which we denote μ^M and μ^F. We then construct adjusted firm effects for all firms, which equal $\psi^g_{J(i,t)} - \mu^M$ for male workers and $\psi^g_{J(i,t)} - \mu^F$ for female workers. These adjusted firm effects used in analyses in which we compare estimates of $\psi^g_{J(i,t)}$ levels across genders, specifically Figs. 6 and 9. Note that since this adjustment is imposed after estimating Eq. (2), it has no effect on our estimates of compensating wage differentials, nor does it impact the way that we characterize sorting patterns within genders.

These five industries include, from lowest to highest average values of $\psi^g_{J(i,t)}$, petroleum coke production (50% women); sports, recreation, and leisure activities (59% women); printing and reproduction (79% women); chemical manufacturing (23% women); and food, lodging, and hospitality (49% women). We obtain similar estimates when we implement this assumption using only the single industry with the lowest average establishment effect. Appendix B presents additional sensitivity analyses related to this assumption.

5.3. Exogeneity conditions

Lavetti and Schmutte (2018) provide a detailed derivation of the OME model specification, including its relationship to a structural log wage equation that characterizes the equilibrium of a labor market with differentiated employers and frictional search, and show a wide range of empirical evidence demonstrating support for the exogeneity conditions among men in Brazil. Here we briefly supplement the evidence from that paper to show that the OME model is also a reasonable representation of the wage data for our sample of women.

Conceptually, there are several classes of potential explanations why the change in residuals may differ for joiners and leavers. First, if establishment wage premia are transitory, joiners and leavers may differ systematically. Second, if job mobility is based in part on transitory log wage shocks, then ε_{it} may be correlated with the assignment of workers to establishments J or industry–occupation cells c. Finally, to the extent job mobility decisions are determined in part on idiosyncratic match effects, since the pure match effect conditional on worker and establishment effects ($\Phi_{i,M(i,t)}|\theta_i, \psi^g_{J(i,t)}$) is contained ε_{it}, such mobility could violate the exogeneity condition.

Fig. 5 replicates the event study diagnostic proposed by Card et al. (2013), which they use to evaluate comparable exogeneity conditions. The figure shows the average wage in each year surrounding a job change, where observations are grouped into cells based on quartiles of the origin and destination wage. The sample is restricted to job change events for which the worker was employed on the origin job and destination job for at least two years. The figure indicates no evidence for the most likely violations of exogenous mobility: (1) systematic declines in residual log wages prior to workers departing, and (2) asymmetric average changes in log wages associated with movements up versus down the distribution of establishment wage effects $\psi^g_{J(i,t)}$.

In addition to the discussion of exogeneity conditions in Lavetti and Schmutte (2018), comparing estimates of β^g for men and women requires an additional assumption that omitted job characteristics do not vary differentially for men and women. Since Eq. (2) controls for gender-specific establishment effects, the exogeneity requirement is that unobserved amenities do not vary differentially for men and women across jobs within establishments in a way that is correlated with fatal injury rates.

6. Can sorting be explained by gender gaps in implicit prices?

One possible explanation for the differences in worker sorting in Fig. 3 is that they could arise due to different incentives, rather than preferences. If women receive smaller wage gains than men when moving between the same pair of firms, or receive lower compensating differentials per unit of fatal risk, then differences in sorting could be straightforward responses to facing different implicit wage prices. In contrast to prior research on this question, we find that the data do not support this explanation.

13 Card et al. (2016) refer to this assumption as a normalization.

14 Abowd et al. (1999) do not make this assumption because they restrict $\psi^g_{J(i,t)}$ to be the same for men and women at all employers.

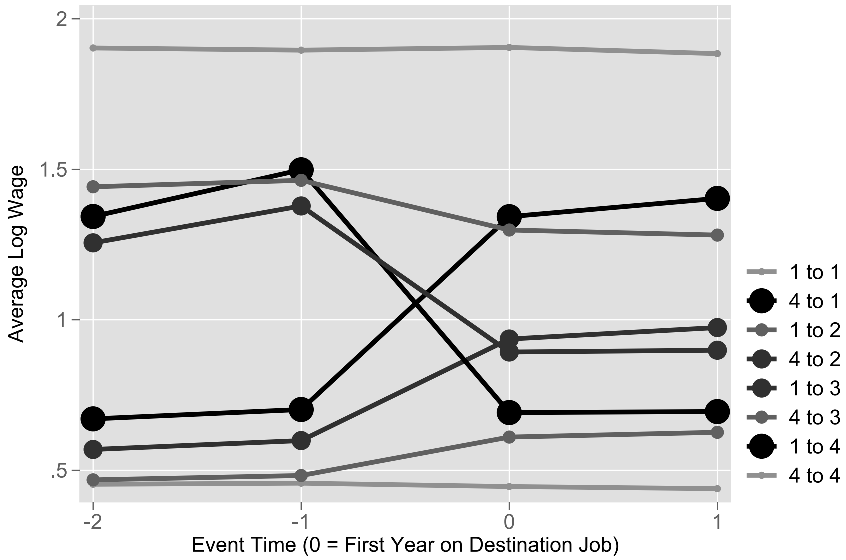

Fig. 5. Event study for female workers who change jobs.
Notes: Mean wages of female workers in the dual connected set who change jobs and held their origin job for at least two years. The jobs are classified by their quartile in the female establishment effect distribution.

Table 2
Gender-specific compensating wage differentials, OME model.

	Fatality rate Industry*Occupation		Fatality Rate Gender*Industry*Occupation		
	(1) Men	(2) Women	(3) Men	(4) Women	(5) Both
Fatality rate	0.233*	0.161*	0.174*	0.174*	0.174*
	(0.002)	(0.005)	(0.002)	(0.005)	(0.002)
Fatality rate*Female					0.001
					(0.005)
VSL (million reais)	3.41	2.06	2.55	2.23	2.43
	[3.34, 3.47]	[1.94, 2.18]	[2.49, 2.60]	[2.11, 2.35]	[2.33, 2.54]
N	13,985,793	8,131,646	13,985,793	8,131,646	22,117,439
R-Sq	0.959	0.970	0.959	0.970	0.971

Notes: Analysis sample includes dominant jobs of attached workers in the dual connected set between ages 23–65, with 30 or more contracted hours per week, at least 1000 total hours annually in all jobs, excluding government jobs, temporary jobs, and jobs in 2-digit industry by 3-digit occupation cells that have fewer than 10,000 FTFY workers in the three-year moving average window used to calculate fatality rates. Column 5 includes both men an women, with gender-specific establishment effects and gender-specific coefficients on all other control variables (year, experience, and occupation). 'Fatality rate' is measured in deaths per 1,000 FTFY workers and is Winsorized at the 99th percentile. Log wages are Winsorized at the 1st and 99th percentiles. VSLs are measured in millions of 2003 reais and calculated at sample mean hourly wage, with 95% confidence intervals reported in brackets. * Indicates significance at the 0.01 level.

Table 2 reports estimates of the compensating wage differential, γ, from our preferred OME specification. In columns (1)–(2) we present estimates separately for men and women using fatality risk measures that pool data from both genders. This gender-pooled measure is used in most previous studies because of the scarcity of the gender-specific fatality data at detailed industry by occupation levels. The coefficient estimate for men, 0.233, suggests that an increase in the average fatality rate of one death per 1000 full-time equivalent worker-years is associated with an approximately 23 percent increase in the hourly wage.[15] The compensating differential for women, 0.161, is 31% lower, consistent with qualitative patterns from prior literature.

In columns (3)–(4) we re-estimate the same specifications using our preferred gender-industry–occupation fatality rates. Column (5) presents estimates from the full combined specification, which includes both women and men, and allows γ (in addition to establishment effects and all other control variables) to be gender-specific. Across all three specifications, the estimated compensating wage differentials for women and men are equal, at $\gamma = 0.174$. In Column (5),

[15] This is estimate is comparable, though slightly larger, than the estimate of 0.17 from Lavetti and Schmutte (2018) due to restricting the sample to the dual connected set of establishments.

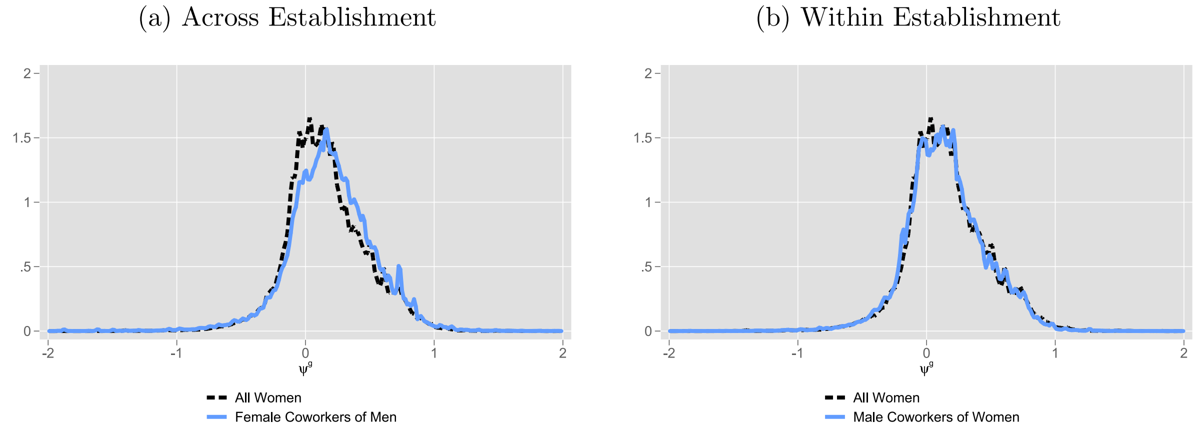

Fig. 6. Decomposition of establishment wage effect.
Notes: Fig. 6(a) plots the kernel density of the distribution of $\psi^g_{J(i,t)}$ for all women (dashed line) and counterfactual kernel density that men would have received if they had been women employed at the same establishment (solid line). This counterfactual distribution is estimated based on the female coworkers of men, weighted by the distribution of jobs held by men. Fig. 6(b) plots the same kernel density for all women (dashed line) compared to the counterfactual kernel density that women would have received if they had been men employed at the same establishment. This counterfactual distribution is estimated based on the male coworkers of women, weighted by the distribution of jobs held by women.

Table 3
Components of Gender Wage Gap.

	Women	Men	Difference
Log wage	1.544	1.727	−0.183
First stage controls (Exp. and Year)	0.670	0.652	0.018
Occupation effects	0.030	0.004	0.027
Person-specific component	−0.109	0.063	−0.172
Establishment assignment	0.160	0.211	−0.051
Within-Estab. gender gap	0.001	−0.000	0.001

Notes: Summaries of components from Eqs. (1) and (2), estimated on the dual connected set, with normalized gender-establishment effects. Occupation effects include combined effect of differences in the distribution of men and women across 1-digit occupation cells, and the conditional average gender gap within occupations.

the estimated deviation from this base parameter for women is 0.001 (SE=0.005), confirming that there is no economically meaningful difference in the compensating wage differential across genders.[16]

This result clarifies somewhat puzzling evidence from previous studies on this topic. Hedonic theory suggests that the equilibrium compensating wage differential should be determined by the preferences of the marginal worker. Since non-segmented labor markets share a common marginal worker, the substitutability of female and male labor generally suggests that large gender differences in compensating wage differentials for safety should be unlikely to occur, even if women and men have different preferences for safety on average. Our findings suggest that more precisely measuring occupational safety, and correcting for the effects of endogenous assignment of workers to establishments, can help explain this puzzle. After implementing these corrections, the compensating wage differentials for men and women are indistinguishable, and therefore they cannot explain the observed gender differences in sorting.

7. Can sorting be explained by gender differences in firm wage premia?

Tables 3 and 4 present summaries of the estimated components of Eqs. (1) and (2). Table 3 reports Oaxaca–Blinder decomposition of the OME model components, while Table 4 reports the contribution of each model component to the gender-specific wage variance.

The Oaxaca–Blinder decomposition shows that the assignment of women and men to different establishments explains 28 percent of the total gender wage gap in Brazil (5.1 log points of the 18.3 point gap), while differences in establishment wage premia conditional on assignment explain less than 1 percent.[17] If women were paid the male establishment wage

16 We also convert the compensating wage differential estimates to implied values of statistical life (VSL), which are used frequently in public policy. The VSL is computed by converting the estimated semi-elasticity into a marginal effect $\frac{\partial w}{\partial a}$, and rescaling. Specifically, $\widehat{VSL} = \overline{w}\widehat{\gamma} * 2,000,000$, where the scaling factor, 2,000,000 is the product of 2000 h per FTFY worker with 1000 workers, since γ is the wage effect of an increase in number of fatalities per 1000 workers.

17 Appendix Figure A.5 shows the relationship between male and female establishment wage effects.

Table 4
Variance decomposition of log wages.

	Women		Men	
	Component	Share	Component	Share
SD of Log wages	0.74	100%	0.69	100%
SD Worker effects	0.59	63%	0.54	60%
SD Estab-Gender effects	0.31	17%	0.28	16%
SD of $X\beta$	0.11	2%	0.13	4%
SD Residual	0.13	3%	0.14	4%
Cov (θ, ψ^g)	0.03	6%	0.04	8%
Cov $(\theta, X\beta)$	0.01	1%	−0.00	0%
Cov $(\psi^g, X\beta)$	0.00	0%	0.00	0%

Notes: Summaries of components from Eqs. (1) and (2), estimated on the dual connected set, with normalized gender-establishment effects. $X\beta$ includes the combined contribution of both first and second stage observed controls and fatality rates.

effect, holding fixed establishment assignment, their wages would *decrease* by 0.8 percentage point (see Appendix Table A.4). Interestingly, the same pattern is true for men—if men were paid the female establishment wage effect, their wages would also decrease by 1.5 percentage points. That is, on average workers in Brazil are concentrated in jobs in which workers of their gender earn modestly higher establishment wage premia than workers of the opposite gender.

Despite a small net gap within establishments (0.1 log points), sorting across establishments contributes substantially to the gender wage gap. The average female establishment wage premium is 16 log points while the average male establishment premium is 21.1 log points. The male distribution of establishment effects is shifted roughly uniformly to the right relative to the female distribution (Appendix Fig. 6(a)), with a mean difference of 5.1 log points, or 28 percent of the wage gap. Partially offsetting this, women in our sample have higher levels of education than men (the effect of which is absorbed by θ), and hold different occupations, which narrows the wage gap by 2.7 log points.

Table 4 presents a decomposition of the variance of log wages. The wage component shares are quite similar for men and women in Brazil, and the share of the variance explained by establishment assignment is 17% for women and 16% for men.

Fig. 6 shows kernel density distributions of the components of establishment wage effects. Fig. 6(a) depicts the distribution of $\psi^g_{J(i,t)}$ for all women (dashed line) and the counterfactual distribution of $\psi^g_{J(i,t)}$ that men would have received if they had been women, holding fixed the distribution of male jobs (solid line). The difference between the distributions isolates the impact of gender differences in establishment assignment on $\psi^g_{J(i,t)}$. Fig. 6(b) isolates the within-establishment component. The figure depicts the same distribution $\psi^g_{J(i,t)}$ for all women (dashed line) and the counterfactual distribution of $\psi^g_{J(i,t)}$ that women would have received if they had been men, holding fixed the distribution of female jobs (solid line). The figures show that, although within-establishment gender wage effects are not zero, they are small in magnitude, and are therefore unlikely to be the cause of the heterogeneity in sorting in Fig. 3. This conclusion is also supported by the fact that both women and men earn larger wage effects on average than opposite gendered-workers employed at the same establishment.

8. Sorting on safety and establishment-level gender segregation

Our aim in this section is to evaluate whether the gender gap in sorting on safety is related to the assignment of men and women to high or low paying establishments. We begin by formalizing a framework for two-dimensional sorting on safety and establishment-specific pay (ψ). We then present graphical evidence on sorting patterns based on vector gradient fields that can be interpreted from the perspective of the sorting framework. Finally, we connect the gender gap in safety to establishment-level gender segregation, and show suggestive evidence that sorting on safety may contribute indirectly to the large gender gap in ψ.

8.1. Simultaneous sorting on wages and fatality risk

We provide a simple framework to characterize job mobility patterns of workers moving across a two-dimensional attribute space. This framework combines hedonic search with the multidimensional sorting model posed by Lindenlaub and Postel-Vinay (2017), and facilitates the interpretation of our empirical estimates that follow.

Workers may have different relative preferences for wages and for safety. In particular female and male workers may differ on average in their preferences. On the employer's side, jobs are differentiated by ψ, the average log wage premium associated with the establishment, and a, which measures the bundle of job attributes correlated with fatality rates. Jobs are characterized by the bundle $\mathbf{y} = (\psi, a)$. Workers are characterized by $\mathbf{x} = (\theta, g)$, which measures their skills, θ, and gender, g. A match between a worker with attributes $\mathbf{x}$ and job with attributes $\mathbf{y}$ generates surplus $\sigma(\mathbf{x}, \mathbf{y})$. The job surplus, σ, can vary across genders due to differences in relative preferences for safety, skill differences between women and men, or due to employer discrimination. For example, this representation allows, as suggested by Garen (1988), that

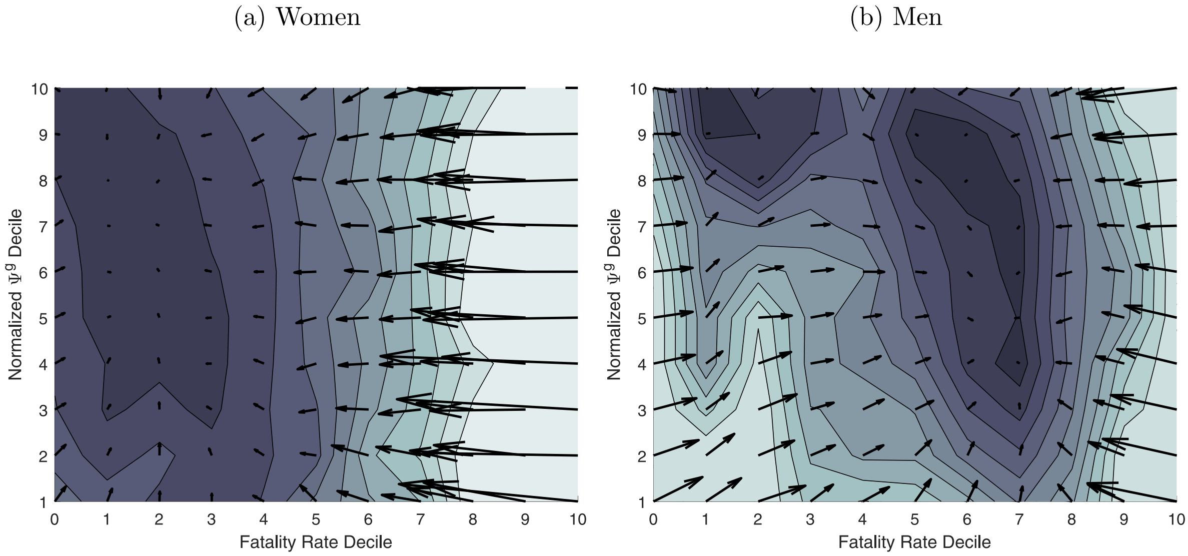

Fig. 7. Job-to-job transition gradient fields. Establishment pay versus risk.
Notes: Samples include attached full-time women/men between ages 23–59 employed in the largest dual connected set of establishments, and excludes origin jobs with zero fatality rate. Deciles defined based on gender-specific wage and fatality rate distributions. Vectors indicate directions and relative, not absolute, magnitudes of changes in wages and fatality rates associated with job-to-job transitions originating from each cell. Contours indicate level sets based on the relative lengths of vectors.

some workers may be relatively more productive in dangerous employment than others, or, conversely, some workers are better at managing risk. In addition, σ could also capture correlations between worker characteristics and average preferences for job attributes. For example, if workers with higher values of θ have stronger average preferences for jobs with y_a, such a match may generate greater surplus.

In an environment characterized by frictional search, employed workers switch between jobs when the surplus at the destination job exceeds that at the origin job: $\sigma(\mathbf{x}, \mathbf{y^d}) > \sigma(\mathbf{x}, \mathbf{y^o})$. This process generates conditional transition probabilities $Pr\,[\mathbf{y}|\mathbf{y^o}, \mathbf{x}]$ that describe the probability a worker moves to a job with bundle $\mathbf{y}$ given he or she has attribute $\mathbf{x}$ and is currently employed in job $\mathbf{y^o}$.

Without additional structure, we can characterize sorting by the vectors $\{\mathbb{E}[\Delta\psi|\mathbf{y^o}, \mathbf{x}], \mathbb{E}[\Delta a|\mathbf{y^o}, \mathbf{x}]\}$, which describe the average change between the destination and origin jobs in wages and fatality rates, conditional on worker characteristics. We empirically estimate these vectors to characterize the full distribution of revealed job change behavior in the ψ and a dimensions as a function of gender, θ, and the characteristics of origin jobs.

Figs. 7(a) and 7(b) display gradient fields describing the average change in establishment effects $\psi^g_{J(i,t)}$, and in fatality rate, a, at starting points defined by the deciles of $\psi^g_{J(i,t)}$ and a. We plot these separately for women and men. To help visualize differences in the magnitudes of the gradient vectors, we plot the field of gradients over the level sets describing gradient magnitudes (gradient vector norms). The darker shaded areas identify level sets in which the gradient vectors have the smallest magnitudes, indicating that origin and destination jobs tend to have very similar characteristics, on average. The far left columns of each graph include origin jobs in which the fatality rate was zero, and the deciles of fatality rates are calculated conditional on the rate being nonzero.

The most striking contrast between women and men is in movement toward safety. In each origin cell, the average change in job characteristics among women is almost never toward more dangerous jobs.[18] Even women employed in the first decile of fatality rates, and at the lowest-paying firms, do not accept meaningful increases in risk. By contrast, men tend to move toward more dangerous jobs, even when originating at high-wage firms.

Overall, women's job movements are more strongly characterized by horizontal movements, consistent with a job ladder model in which women seek safer jobs and experience fewer changes toward high paying firms. Men move toward riskier jobs when originating below the 7th decile of the a distribution, but the direction of sorting switches at the top of the distribution. The magnitude of female gradients (vector norms) are on average much smaller than those for men, implying that origin and destination jobs are more similar. The vector projection of the female gradient in the vertical dimension is also smaller on average, suggesting that women are less likely to move toward establishments with higher average pay.

Figs. 8(a) to 8(d) present transition gradient fields broken out by groups of women and men with estimated person effects (θ) above and below the mean. These figures provide additional information about the contrast between the

[18] Of course, in the far left column where the origin fatality rate is zero, the change in risk is mechanically always slightly positive on average.

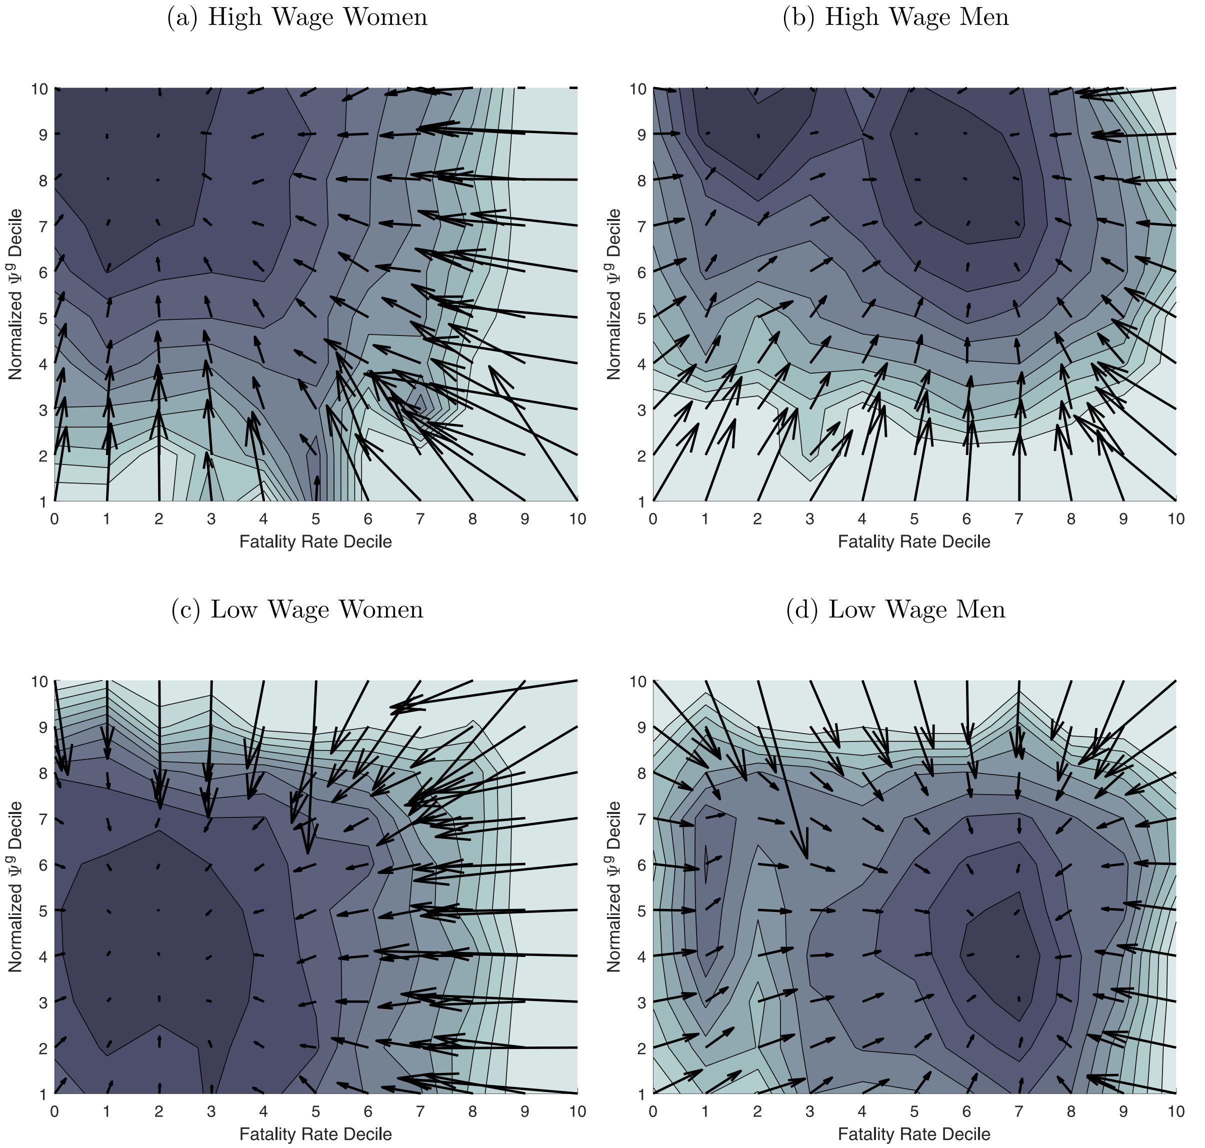

Fig. 8. Job-to-job transition gradient fields: Establishment pay versus physical risk. High wage versus low wage workers.
Notes: Samples include attached full-time women/men between ages 23–59 employed in the largest dual connected set of establishments, and excludes origin jobs with zero fatality rate. Deciles defined based on gender-specific wage and fatality rate distributions. Vectors indicate directions and relative, not absolute, magnitudes of changes in wages and fatality rates associated with job-to-job transitions originating from each cell. Contours indicate level sets based on the relative lengths of vectors.

mobility of high-wage versus low-wage workers that is common across genders. Job movements of high-wage women and men are strongly in the direction of higher-paying firms, consistent with job transitions contributing to positive assortative matching. By contrast, low-wage women and men generally move more strongly in safety (horizontal) dimension. Low wage women and men are also both observed to move from higher-wage to lower-wage firms, which could either reflect assortative matching, or involuntary job transitions. However, we find very similar sorting patterns when restricting the sample to transitions initiated by worker resignations, which are more likely to be voluntary separations (Appendix Figure A.7). Consistent with the overall contrast in sorting between high-wage and low-wage workers, our estimates imply a positive correlation between θ and $\psi^g_{J(i,t)}$. Appendix Figure A.3, however, shows that no such correlation is evident when the transition gradient fields are defined by earnings volatility rather than physical safety. Again, in these figures men and women appear to sort very similarly.

One caveat to interpreting these figures is that they assume the presence of some form of labor market friction, otherwise workers would always choose jobs with the highest combination of amenities. However, if frictions differ by gender then mobility patterns should not necessarily be interpreted as productive complementarities or differences in

Fig. 9. Differences by gender in fatality risk and establishment segregation.
Notes: Vertical axis in both figures is the average establishment dissimilarity index for female and male workers within a 2-digit industry and 3-digit occupation. Horizontal axis in Fig. 9(a) is the female minus male fatality rate, measured in deaths per 1000 FTFY workers. Horizontal axis in Fig. 9(b) is the female minus male average normalized ψ in the corresponding industry–occupation cell. Figures are based on the dual connected sample, and include cells for which both women and men have zero fatalities.

average preferences. For example, gender-based hiring discrimination could alter gradients in the ψ dimension. These figures alone do not differentiate between the underlying causes of sorting.

To evaluate whether these patterns are potentially driven by unmeasured establishment-level heterogeneity in worker safety, we test whether establishment-occupation cells with a higher share of female employees than the industry–occupation average have lower fatality rates (for either men or women). We find no relationship between establishment-level fatality rates and gender shares. This suggests that establishment-level variation in safety within-industries does not contribute to the observed sorting patterns.

8.2. Sorting on safety, establishment segregation, and sorting on wages

To what extent does sorting on safety contribute to the wage gap through its influence on establishment assignment? The evidence we have presented suggests that answering this question may help shed new light on why gender differences in labor market sorting have not converged along with the closing of the skills gap. A challenge for future research on this topic is that further progress at understanding the causal influences behind sorting may require a linkage between data on job mobility networks and a large array of observable job amenities. With limited information about amenities we can provide only associative evidence.

Fig. 9 depicts three bivariate relationships between occupational safety, the degree of segregation of men and women across employers, and the firm wage component of earnings (ψ). To construct the figure, we first measure segregation using the dissimilarity index across establishment-occupations within each industry–occupation cell. The dissimilarity index is defined as:

$$D = \frac{1}{2}\sum_{\ell=1}^{K}\left|\frac{f_\ell}{F} - \frac{m_\ell}{M}\right|$$

where K is the number of establishments within the cell, f_ℓ is the number of females employed in the given occupation in establishment ℓ, and F is the number of females employed in the cell, with equivalent definitions for m_ℓ and M. The index can be interpreted as the share of workers of either gender who must be reallocated to make the share of female workers in each establishment-occupation identical to the overall female share in the industry–occupation cell. The values range from zero, in which case there is no excess segregation, to one, in which case there is total segregation.

Fig. 9(a) shows the relationship between the dissimilarity index and the difference between the female and male fatality rates in the same industry–occupation cell (the values are equal to the female fatality rate minus the male fatality rate.) The figure shows that when there is a larger gap between the safety of men and women in the same industry–occupation cell, men and women are far more segregated across firms in that cell. On the left edge of the graph, which includes jobs in which women experience lower fatality rates than men in the same industry–occupation cell, the lower female fatality rate is not only associated with an increase in the share of women employed in the cell (as Fig. 2(a) suggested), but these women are less likely to have male colleagues, and are instead more segregated at firms with other female workers.

Fig. 9(c) shows that in industry–occupation cells in which men and women are more segregated across firms, the establishment wage effect earned by women is substantially lower. Fig. 9(b) connects the axes of these two bivariate distributions, and shows that on average when the gender gap in safety is larger, women earn meaningfully lower establishment wage effects.

Although these figures only depict associations, and should not be interpreted causally, they offer strong motivation for further work in this area. The slope of the line shown in Fig. 9(b) implies that when the relative safety gap is one standard deviation larger (0.07 units), the relative wage gap is on average 1.8 percentage points larger, which equals 10% of the overall unconditional gender wage gap in Brazil.

Fig. 9(a) also helps explain why women appear to be safer than men within the same industry and occupation. If this pattern could be explained by gender-based task-assignment within industry–occupations, one may have hypothesized that in more gender-segregated establishment-occupation cells there would be less potential for such task-shifting, leading to smaller gender differences in safety. Fig. 9(a) shows the opposite pattern. In the absence of task-specific data, we cannot definitively rule out this possibility in favor of alternative explanations like unmeasured sorting, or women behaving differently than men when performing similar tasks.[19] However, the figure suggests that whatever makes women safer is associated with gender-based sorting across establishments within narrowly defined types of jobs.

9. Discussion

A large literature in labor economics has documented the importance of sorting into high-wage and low-wage firms in explaining career wage dynamics and changes in earnings inequality over time. Our analyses reveal that for many workers, sorting on safety plays a strikingly large role in understanding job dynamics, and is roughly as salient in explaining mobility patterns as sorting on firm-level compensation. Although safety has a relatively small direct impact on wages, it is strongly related to gender disparities in establishment sorting. This has the potential to indirectly increase gender wage inequality, because establishment sorting explains about 28% of the entire gender wage gap in Brazil.[20]

Of course, there are potentially many unobserved amenities that are correlated with safety and vary *within* establishments and occupation groups, and our estimates identify the combined effect of these amenities. To offer a contrast, we show that there is no gender difference in job mobility patterns with respect to earnings risk, despite clearly documented gender differences in aversion to financial risks. Relatedly, Sorkin (2018) uses a revealed preference approach based on job mobility patterns, and infers that the net value of a bundle of unobserved establishment-level amenities explains over half the firm component of earnings variation. Taber and Vejlin (2020) estimate that one-third of job choices would have been different if workers cared only about wages. In all of these studies, unobserved differences in amenities that vary across jobs within establishments have the potential to confound interpretation. Our finding corroborate and extend this evidence to show that a specific observed amenity, safety, is, at a minimum, a strong proxy for the set of amenities that explain workers' choices between jobs.[21] Moreover, the strong degree of safety-related mobility appears to be of first-order importance for understanding how women and men sort differently through the labor market, and why their wages diverge when transitioning jobs.

Our findings are also suggestive that policies aimed at improving workplace safety may alter the patterns of sorting across establishments, potentially reducing gender segregation in labor markets and the wage gap. A caveat to this possible implication is that, despite a 30% decline in the overall occupational fatality rate for women in Brazil between 2005–2010,

19 Cortes and Pan (2018), for example, use O*NET task data and show that women sort into occupations differently than men in part on the basis of the set of tasks associated with an occupation.

20 Across-establishment sorting cannot be explained by establishment-specific differences in safety, conditional on occupation.

21 In many related studies, including this paper, error in the estimation of person and firm wage effects are potential sources of bias.

the gender wage gap remained virtually unchanged. These safety improvements were one-sided, however, as the safety gap between female and male workers actually increased over this period. This leaves open the possibility that broad-based improvements in occupational safety may still have an equalizing force on labor market earnings potential for workers with different safety preferences.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeconom.2022.06.012.

References

Abowd, J.M., Creecy, R.H., Kramarz, F., 2002. Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. Technical Report TP-2002-06, LEHD, U.S. Census Bureau.

Abowd, J.M., Kramarz, F., Margolis, D.N., 1999. High wage workers and high wage firms. Econometrica 67 (2), 251–333.

Bartolucci, C., 2013. Gender wage gaps reconsidered a structural approach using matched employer-employee data. J. Hum. Resour. 48 (4), 998–1034.

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

Bertrand, M., 2011. New perspectives on gender. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, Vol. 4. Elsevier, pp. 1543–1590, chapter 17.

Blau, F.D., Kahn, L.M., 2017. The gender wage gap: Extent, trends, and explanations. J. Econ. Lit. 55 (3), 789–865. http://dx.doi.org/10.1257/jel.20160995.

Bonin, H., Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2007. Cross sectional earnings risk and occupational sorting: The role of risk attitudes. Labour Econ. 14 (6), 926–937.

Card, D., Cardoso, A.R., Kline, P., 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. Q. J. Econ. 131 (2), 633. http://dx.doi.org/10.1093/qje/qjv038.

Card, D., Heining, J., Kline, P., 2013. Workplace heterogeneity and the rise of West German wage inequality. Q. J. Econ. 128 (3), 967–1015.

Classificação Brasileiro de Ocupações: Downloads - 5.1.0, 2020. Classifica Cão Brasileiro de Ocupa Cões, Ministerio do Trabalho, URL http://www.mtecbo.gov.br/cbosite/pages/downloads.jsf. (Accessed on 20 March 2020).

Classificação Nacional de Atividades Econômicas, 2018. URL https://cnae.ibge.gov.br/estrutura/atividades-economicas-estrutura/cnae. (Accessed on 16 October 2020).

Cortes, P., Pan, J., 2018. Occupation and gender. In: Averett, S., Argys, L., Hoffman, S. (Eds.), The Oxford Handbook of Women and the Economy. Oxford University Press.

DeLeire, T., Levy, H., 2004. Worker sorting and the risk of death on the job. J. Labor Econ. 22 (4), 925–953.

Garen, J., 1988. Compensating wage differentials and the endogeneity of job riskiness. Rev. Econ. Stat. 70 (1), 9–16.

Goldin, C., 2014. A grand gender convergence: Its last chapter. Amer. Econ. Rev. 104 (4), 1091–1119.

Goldin, C., Kerr, S., Olivetti, C., Barth, E., 2017. The expanding gender earnings gap: Evidence from the LEHD-2000 census. Am. Econ. Rev.: Pap. Proc. 107 (5), 110–114.

Hellerstein, J.K., Neumark, D., 2008. Workplace segregation in the United States: Race, ethnicity, and skill. Rev. Econ. Stat. 90 (3), 459–477. http://dx.doi.org/10.1162/rest.90.3.459.

Hersch, J., 1998. Compensating differentials for gender-specific job injury risks. Am. Econ. Rev. 88 (3), 598–607.

Hotz, V.J., Johansson, P., Karimi, A., 2017. Parenthood, Family Friendly Firms, and the Gender Gaps in Early Work Careers. Working Paper Series 24173, National Bureau of Economic Research, http://dx.doi.org/10.3386/w24173.

Lalive, R., 2003. Did we overestimate the value of health? J. Risk Uncertain. 27 (2), 171–193. http://dx.doi.org/10.1023/A:1025685024913.

Lavetti, K., Schmutte, I.M., 2018. Estimating Compensating Wage Differentials with Endogenous Job Mobility. Working Paper 29, Labor Dynamics Institute, Cornell University.

Leeth, J.D., Ruser, J., 2003. Compensating wage differentials for fatal and nonfatal injury risk by gender and race. J. Risk Uncertain. 27 (3), 257–277.

Lindenlaub, I., Postel-Vinay, F., 2017. Multidimensional Sorting under Random Search. Technical Report 501, Society for Economic Dynamics, 2017 Meeting Papers.

Mas, A., Pallais, A., 2017. Valuing alternative work arrangements. Amer. Econ. Rev. 107 (12), 3722–3759.

Merlino, L.P., Parrotta, P., Pozzoli, D., 2018. Gender differences in sorting. Ind. Relat.: J. Econ. Soc. 57 (4), 671–709. http://dx.doi.org/10.1111/irel.12216.

Ministério do Trabalho e Emprego, 2017. Base da Relação Anual de Informações Sociais (RAIS): 2003—2010. (Accessed 28 November 2017). [database].

Rosen, S., 1974. Hedonic prices and implicit markets: Product differentiation in pure competition. J. Polit. Econ. 82 (1), 34–55.

Sorkin, I., 2018. Ranking firms using revealed preference. Q. J. Econ. 133 (3), 1331–1393.

Taber, C., Vejlin, R., 2020. Estimation of a Roy/search/compensating differential model of the labor market. Econometrica 88 (3), 1031–1069. http://dx.doi.org/10.3982/ECTA14441.

Tsai, W.-J., Liu, J.-T., Hammitt, J.K., 2011. Aggregation biases in estimates of the value per statistical life: Evidence from longitudinal matched worker–firm data in Taiwan. Environ. Resour. Econ. 49 (3), 425–443.

Wiswall, M., Zafar, B., 2018. Preference for the workplace, human capital, and gender. Q. J. Econ. 133 (1), 457–507.

Woodcock, S., 2008. Wage differentials in the presence of unobserved worker, firm, and match heterogeneity. Labour Econ. 15, 771–793.