Gender Differences in Earnings and Sorting on Occupational Safety*

Kurt Lavetti†
Ian M. Schmutte‡

May 13, 2017

Abstract

Using Brazilian matched employer-employee data, we adapt the AKM wage decomposition to estimate compensating differentials for occupational safety, allowing for gender-specific sorting on unobserved worker and establishment characteristics. After correcting for endogenous sorting and measurement errors, using the most detailed fatality measures to-date, we conclude the female and male VSLs are nearly identical, contrary to previous findings. We decompose earnings and safety differences to show sorting on safety explains 19% of the gender-wage gap, roughly the same share as sorting across establishments. When changing jobs, we show that women move more strongly towards safety, whereas men move towards higher-paying establishments.

JEL Classifications: J31, J16, J17
Keywords: Wage Differentials, Gender, Occupational Sorting, Value of Statistical Life

*We are grateful for helpful discussions with Martha Bailey, Tom DeLeire, and Larry Kahn.
†Ohio State University, Department of Economics. `lavetti.1@osu.edu`
‡University of Georgia, Department of Economics. `schmutte@uga.edu`

1 Introduction

In this paper we provide new evidence on the role of gender differences in workplace safety on estimates of compensating wage differentials. We are interested in the extent to which men and women have different wages because (1) women sort into jobs differently than men on the basis of occupational safety and earnings (2) within the same type of job, women and men face different risks of death, and (3) women confront a different marginal implicit price for occupational safety than do men. These three possibilities are clearly linked in the hedonic wage theory Rosen (1974). Under that model, if women have stronger preferences for safety than men, they will sort into safer jobs and receive a lower wage. Furthermore, if women face discrimination, or if the hedonic pricing function is nonlinear, women may confront a higher marginal price for safety than men.

Despite clear evidence that women have very different preferences for safety than men, and that these preferences drive a substantial portion of occupational sorting (Hersch 1998; DeLeire and Levy 2004), there is surprisingly little research focused on separately estimating compensating wage differentials for female workers, particularly when compared to the very large literature on this topic for men. This dearth of research could lead to gender bias in decisions about public policies that affect safety and the risk of death. In cost-benefit analyses of such policies, benefits are commonly based in part on estimates of workers' implicit willingness to trade off wages to reduce the risk of fatal injury on the job. In policy evaluation, estimated compensating wage differentials for the risk of fatal injury are normalized to measure the value of a statistical life (VSL). The VSL can be thought of as the amount of money a utilitarian social planner would be willing to pay to save one life. However, if women and men have different preferences for safety, policy decisions based on existing VSL estimates may be biased toward the preferences of male workers.[1]

To fill this gap in the literature, we use matched employer-employee data from Brazil to estimate a very precise female-specific VSL from hedonic wage models using the complete population of women employed full-time in the formal sector between 2003–2010. A key challenge to this approach to estimating the VSL is to find appropriate measures of the risk of fatal injury. Our analysis is based on the most detailed, gender-

[1]To the extent that there is publication selection bias in journals against zero or negative results, the combination of lower VSLs and larger standard errors for women could amplify a gender-imbalance in information about the VSL (Viscusi 2015).

specific measures of fatal workplace injury to date. The Brazilian data record whether each job ends due to a fatal workplace-related injury, and thus serve as a census of fatal occupational injuries. We use this information to construct gender-by-occupation-by-industry aggregate fatality rates. We find that women sort into much safer industry-occupation cells. Furthermore, women are much safer than men on average, even when employed in the same cell. We are able to show that 62% of the increased safety is explained by women sorting, within a cell, into different establishments than men. This is the first such evidence of its kind we are aware of.

Our analysis addresses two related, but distinct, sources of measurement error that afflict prior work. First, it is common for studies to use measures of fatality risk that combine information across genders[2]. Second, most research in this area tends to measure fatality risk by using industry- or occupation-specific aggregates. In some of the only studies focused on estimating compensating wage differentials for female workers, Hersch (1998) and Leeth and Ruser (2003) show that women face considerably lower rates of injury and fatality than men when employed in the same industry or occupation. Furthermore, failing to account for gender differences in safety leads to bias in estimates of compensating wage differentials for women. We also find women and men have different risk of fatal injury when employed in the same type of job. In our analysis, we find that for women 91% of the total variation in fatality rates occurs across 2-digit industry codes and within 3-digit occupation codes in our setting, whereas for men this same component explains only 34% of variation. Hence, using male data to measure female fatality risk will generate considerable measurement error. Furthermore, combining fatality data across industries or occupations is also a source of non-ignorable aggregation bias, as observed by Tsai, Liu and Hammitt (2011) and Lalive (2003).

After correcting for measurement error in fatality rates, we estimate that the female VSL is nearly identical to that of men: 3.77 million Brazilian reais (SE 0.02 million) compared to 3.72 million (SE 0.01 million) for men. Because the VSL is calculated as the product of the estimated compensating wage differential and the average wage, the similarity in male and female VSL estimates masks two nearly offsetting effects—women pay a higher price than men to obtain less risky employment, but the 18.3

[2] For example, Costa and Kahn (2004) and Aldy and Smyth (2014) estimate a VSL for female workers as part of an analysis not exclusively focused on gender differences in compensating differentials, and so do not use measures of fatality risk that differ by gender.

percent gender wage gap in Brazil reduces the female VSL by almost exactly the same proportion. Failing to account for gender differences in fatality rates has a very strong effect on the estimated compensating wage differential, and hence on the VSL. When fatality rates are measured by pooling both genders, the estimated female VSL is 150 percent larger at 5.66 million reais.

We use both measures to estimate the cross-sectional and panel data models that are common in the literature (Hersch 1998; Kniesner, Viscusi, Woock and Ziliak 2012) as well as a model that accounts for sorting across firms on the basis of unobserved employer-specific characteristics introduced by Lavetti and Schmutte (2016). Using a male-only sample in Brazil, Lavetti and Schmutte (2016) show that adjusting for unobserved employer-specific characteristics corrects for endogenous mobility bias, resulting in estimates of the VSL that are an order of magnitude larger than those based on conventional panel data estimators (implying VSLs of 3.84 million Brazilian reais versus 0.32 million reais, respectively). We find similar patterns of endogenous mobility bias for women, which cause within-worker estimates of the female VSL to be about 80% lower than our preferred estimate. Most important, from our perspective, is that we find failing to account for gender differences in preferences causes the estimated compensating wage differential to be under-estimated for female workers, as implied by DeLeire and Levy (2004). On the other hand, failing to account for measurement error in fatality rates leads gender differences in the VSL to be overstated. Hence, appropriate modeling and measurement are both critical for avoiding gender bias in policies around health and safety that use the VSL as an input to cost-benefit calculations.

Our paper is also closely related to a growing literature using matched data to study differences in the way men and women sort across jobs. Bayard, Hellerstein, Neumark and Troske (2003) show that gender segregation across occupations and establishments contributes to gender inequality in the U.S. More recently, Card, Cardoso and Kline (2016) also find that the sorting of men into higher-paying establishments can explain around 20 percent of the overall gender wage gap, and suggest this is partly driven by differences in rent-sharing between women and men. Relatedly, Bartolucci (2013) suggests women have lower bargaining power than men, but are also less productive and more mobile. Merlino, Parrotta and Pozzoli (2014) finds these gender differences in sorting may be driven by differences in the propensity for women to be promoted within the firm. In an interesting recent paper, Barth, Goldin, Kerr and Olivetti (2017) find the gender gap is driven in part by differences in pay growth within firms, after controlling for sorting across firms. None of these papers considers sorting on the basis of an observed disamenity like risk of fatal injury. Because we can anchor our analysis in a specific, observed, amenity, we can focus on the relative importance in earnings of sorting across establishments versus sorting across occupations on the basis of the risk of fatal injury.[3]

We find that women earn 3.2% less than men because they are more likely to be employed at low-wage establishments, which is 18% of the overall wage gap. The effect of sorting across jobs with different risk of fatal injury explains about the same share, 19%, of the total gender wage-gap as does the sorting across establishments. These striking gender differences in sorting are evident in both cross-sections and in longitudinal job changes. When women change jobs they are much more likely than men to take new jobs with lower fatality risk. Furthermore, the share of the total change in earnings that can be explained by the change in fatality risk, relative to the share from movements to higher-paying establishments, is twice that of men.

2 Data and Descriptive Statistics

We use matched employer-employee data from Brazil's *Relação Anual de Informações Sociais* (RAIS) from 2003–2010. RAIS is an 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. The information contained in RAIS is reported by establishments, and participation is compulsory because the data are used to calculate a constitutionally-mandated annual worker bonus that is equal to one month's earnings. In addition to providing a census of all formal-sector jobs, the RAIS also contains a census of every occupational fatality linked to each job.

In each year, every establishment reports information for each employment contract that was active at any point during the year, so that the unit of observation in the data is a contract-year. The data include characteristics of the worker, establishment, and 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,

[3]Card et al. (2016) briefly consider whether their results can be explained by differences across firms in hours requirements and their associated compensating differentials, but do not estimate those differentials or characterize the extent of sorting on such hours differences.

contracted weekly hours, occupation, dates of hire and separation, and the cause of separation, which includes three detailed codes identifying work-related fatalities.

2.1 Estimation Sample

We use these data to construct a matched worker-establishment panel of jobs. Starting with the raw data, we restrict attention to a single dominant job for every worker in every year. This approach follows Abowd, Kramarz and Margolis (1999), Woodcock (2008), and Card, Heining and Kline (2013). Defining expected earnings as the product of the monthly contracted earnings rate times the number of months employed in the year, we define the *dominant job* for each worker in a given year as the job with highest expected earnings. Dropping jobs that are not dominant eliminates approximately 13 percent of observations.

Using the census of fatalities linked to jobs, we construct gender-specific average fatality rates in each year for each of the 11,440 two-digit industry by three-digit occupation cells. As is common in the literature (Kniesner et al. 2012; Lavetti and Schmutte 2016), we combine these annual rates into three-year moving averages to increase the precision of our estimated fatality rates. We measure fatality rates by the gender-specific moving average number of fatalities per 100,000 full-time full-year-equivalent workers.

The population of interest includes dominant jobs held by any worker aged between 23 and 65. We stratify the population by gender, and then draw separate analysis samples for men and for women. In this paper, we focus primarily on women, but report results for men for comparison. Within each gender we restrict the analysis sample to full-time jobs (at least 30 contracted hours per week) reported by establishments that employ at least two workers during the year. We exclude government jobs, as work location is not accurately reported. We also exclude jobs where the employee is on a temporary contact. Following Lavetti and Schmutte (2016), we 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). There is a trade-off in our aggregation between, on one hand, constructing more precise measures of fatality rates, and potentially increasing measurement error if cells are too small to offer sufficient information about fatality rates, particularly since fatal accidents are very rare events. For analysis, we Winsorize the data at the 1st and 99th percentiles of the log wage distribution within

Table 1: Summary Statistics

	Female		Male	
	Population	Analysis Sample	Population	Analysis Sample
Age	35.83	35.24	35.46	36.13
Race *branco* (White)	0.57	0.65	0.56	0.58
Less than High School	0.29	0.33	0.47	0.49
High School	0.42	0.46	0.37	0.40
Some College	0.06	0.07	0.04	0.04
College or More	0.24	0.14	0.12	0.07
Contracted Weekly Hours	39.46	42.59	42.03	43.32
Log Hourly Wage	1.33	1.17	1.42	1.37
Total Experience (Years)	18.35	18.00	18.98	19.71
Job Tenure (Months)	62.36	44.97	51.69	44.99
Pooled Fatality Rate (per 100,000)	2.14	2.62	7.04	7.26
Gender-Specific Fatality Rate	1.04	1.24	7.79	8.05
Zero Fatality Rate (Percent)	0.26	0.32	0.15	0.11
Number of Observations	134,361,238	47,096,981	194,907,785	79,162,274

Notes: Means shown are calculated using the 100% Brazilian RAIS data from 2003-2010. 'Female Population' includes all dominant jobs held by women between ages 23 and 65. Dominant jobs are defined as the job with the highest earnings in a given year. 'Female Sample' includes all dominant jobs held by women between ages 23-65 with at least 30 contracted hours per week, excluding government jobs and temporary jobs, held at establishments with at least two workers, in 2-digit industry by 3-digit occupation cells with a total of at least 10,000 full-time full-year equivalent workers, with hourly earnings between the 1st and 99th percentiles of the Female Sample earnings distribution, and fatality rates below the 99th percentile of the overall fatality rate distribution. Fatality Rate is measured as the number of fatalities per 100,000 full-time full-year-equivalent workers. 'Male Population' and 'Male Sample' are defined similarly.

each gender, and at the 99th percentile of the overall population fatality rate. Imposing these restrictions yields an analysis sample for women with 47 million observations.

2.2 Descriptive Statistics

Table 1 reports descriptive statistics for the full population and for the analysis sample separately by gender. Among female workers, relative to the population, observations in the analysis sample involve workers that are less educated, have lower seniority, lower wages, and are employed in riskier jobs. We focus on several key distinctions between the male and female analysis sample. These differences are due primarily to our restriction to jobs in industry-occupation pairs with more than 10,000 full-time full-year-equivalent workers.

First, the average wage is about 20 percent higher for men than for women. This will mechanically cause the estimated value of statistical life to be considerably lower for women than for men, even if both genders have the same preferences for safety. Second, 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.) Third, as we have noted, women are considerably safer than men even within narrowly-defined industry-occupation cells. The average female-specific fatality rate across jobs held by women is 1.2 deaths per 100,000 full-time full-year workers, or 1.8 deaths conditional on the fatality rate being greater than zero. Hence, using the gender-pooled measure overstates the average female fatality rate by a factor of two. Finally, women are much more likely than men to be employed in industry-occupation cells with zero observed fatalities. Fully 32 percent of job-years in the female analysis sample are in jobs with a measured fatality rate of zero, versus 11 percent for men. In our analysis, we will allow the estimated wage to vary discontinuously at fatality risk equal to zero. Otherwise, the estimated compensating wage differential may be biased, as discussed in Lavetti and Schmutte (2016), who also show how jobs with small or zero fatality risk can be used to detect model misspecification.

3 Gender Differences in Occupational Fatality Rates

Here we show that measures that pool information on fatal injuries across men and women will measure female fatality risk with considerable error. We find that women are employed in jobs that are safer for both men and women. Furthermore, women have much lower rates of fatal injury than men when employed in the same industry and occupation. Here, we document these empirical patterns and try to understand why women are safer than men when hired in the same job. We also characterize the size and distribution of implied measurement error.

3.1 Women are Safer Than Men

Women sort into safer industry-occupation cells than men. As DeLeire and Levy (2004) argue, such sorting is consistent with women having stronger preferences for safety, on average. Figure 1a 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.

We next ask whether sorting is a primarily between- or within-establishment phenomenon. A decomposition of the variance in female fatality rate shows that 71% of the total variation in fatality rates occurs across 3-digit occupations within establishments. Figure 1b shows that within an establishment, doubling the occupation-level female fatality rate is associated with a 5.7% reduction in the share of female employees in that occupation relative to the overall female share of employment in the establishment.

We also find that women are much less likely to die on the job than are men in the same industry-occupation cell. These findings are similar to those of Hersch (1998), which shows for the U.S. that women have much lower rates of illness and injury. Figure 2 depicts the joint distribution of the male and female fatality rates by industry-occupation cell, where the sizes of circles are 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, one would expect the correlation between these variables to be approximately one, consistent with the dashed line in the figure. Instead, the analysis suggests the female fatality rate in the same industry-occupation cell increases is 62% of the male rate.

There are several reasons women might be safer than men within the same industry-occupation cell. One possibility is that women are more cautious at performing similar tasks as men. Another is that even within a narrow industry-occupation cell, men and women are assigned to different types tasks with different safety hazards.

In the absence of task-specific data, it is not possible to distinguish whether women have lower incidence of workplace fatality because of unmeasured sorting, task assignment, or because women behave differently when performing similar tasks.[4] Figure 3 provides evidence that measured differences in fatality rates between men and women

[4]Cortes and Pan (2017), 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. A similar type of sorting could potentially occur within occupations that include a range of separable tasks.

Figure 1: Female Worker Shares versus Log Fatality Rate

(a) Aggregate Distribution

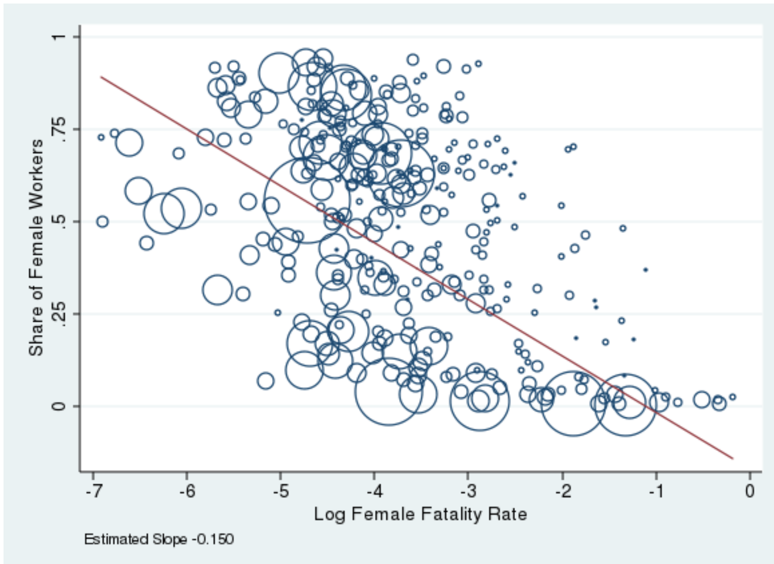

(b) Within-Establishment Distribution

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 1,000 FTFY workers. Figure includes cells with positive female fatality rates and is truncated at the 99th percentile of fatality rates. Circle sizes are proportional to the total number of workers in the industry-occupation cell.

Figure 2: Female versus Male Fatality Rates by Industry and Occupation

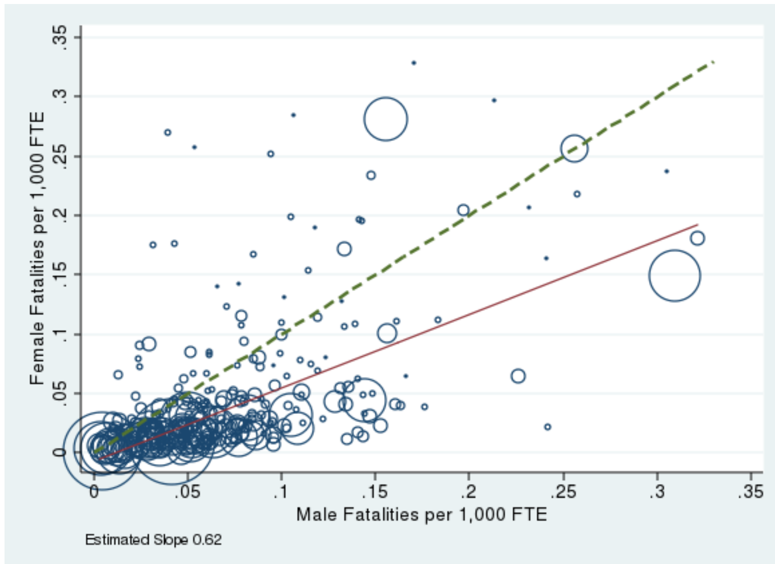

Notes: Fatality rates are gender-specific averages by 2-digit industry and 3-digit occupation pooled over all years from 2003–2010, measured in deaths per 1,000 FTE worker-years. Figure includes cells with strictly positive estimated fatality rates and is truncated at the 99th percentile of fatality rates (0.33 fatalities per 1,000 FTEs). Circle sizes are proportional to the total number of workers in the industry-occupation cell. Dashed line is 45-degree line.

in the same industry-occupation are larger in industry-occupations with more gender segregation. In the figure, the horizontal axis measures the difference between female and male fatality rates (so from left to right female fatality rates increase relative to male rates). The vertical axis measures segregation using the dissimilarity index across establishment-occupations within each industry-occupation cell. The dissimilarity 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 index ranges from zero, in which case there is no excess segregation, to one, in which case there is total segregation.[5]

[5] The dissimilarity index for an industry-occupation cell is computed 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

Figure 3: Differences by Gender in Fatality Risk and Establishment Segregation

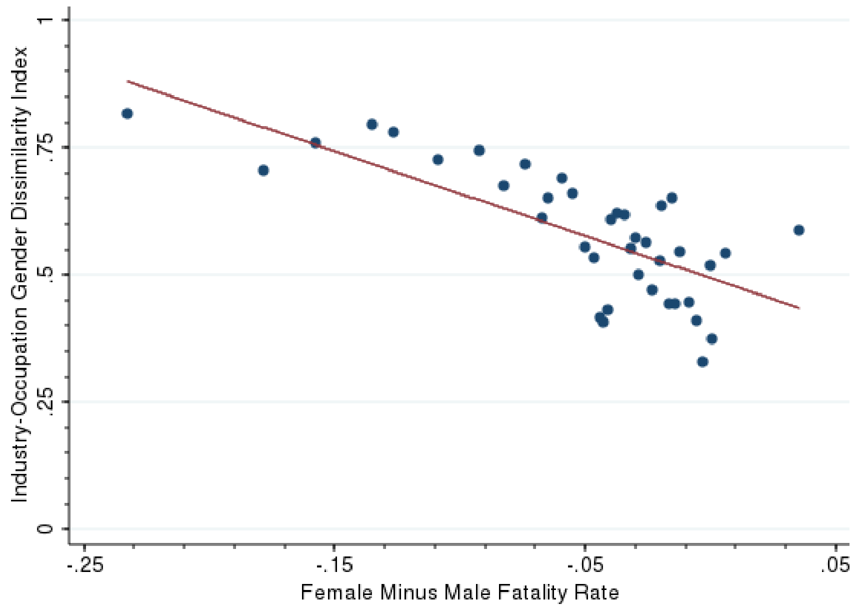

Notes: Vertical axis is the dissimilarity index for male and female workers within a 2-digit industry and 3-digit occupation. The horizontal axis measures the difference between the female and male fatality rate in deaths per 1,000 FTFY workers.

The figure shows that when women are much safer than men in the same narrowly-defined job, there tends to be much stronger gender segregation, as measured by the dissimilarity index. This finding is consistent with women sorting into safer establishments within a given industry-occupation. Note that this figure includes cells for which both men and women have zero fatality risk, yielding the mass at zero. Interestingly, the pattern appears to be roughly symmetric—for the small number of cells in which women have a higher measured fatality rate than men (to the right of zero), there is also some evidence of a positive relationship between the safety gap and gender segregation in jobs. In either case, these facts suggest that whatever makes women safer is associated with gender sorting across employers within cells. To be clear, this does not rule out the possibility that women are also safer because they exhibit greater caution on the job.

We also find the growth of female employment is disproportionately concentrated

equivalent definitions for m_ℓ and M.

Figure 4: Intertemporal Changes in Employment Gender Shares by Fatality Rate

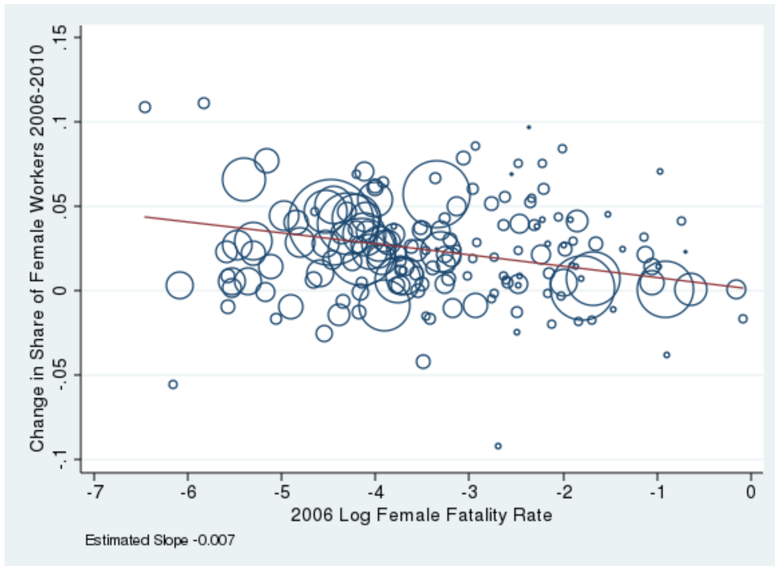

Notes: Vertical axis is the change between 2006-2010 in the employment-weighted share of female workers by 2-digit industry and 3-digit occupation. Horizontal axis is the log of the female-specific fatality rate in the 2-digit industry and 3-digit occupation in 2006, measured in deaths per 1,000 FTE worker-years. Circle sizes are proportional to the total number of full-time workers of both genders in the industry-occupation cell.

in jobs that are safer. Figure 4 plots the change between 2006 and 2010 in the share of full-time workers who are female within each industry-occupation cell versus the fatality rate level at the beginning of the period. Overall during this five-year period the female share of total employment increased by about 2%, and this increase was driven primarily by entry of women into relatively safer jobs, with no net change in the female share on average in the most dangerous jobs. This evidence suggests at entry, women prefer safer jobs, and their ability to find them has increased over the sample period. Note that over this period, the Brazilian economy enjoyed a trade boom that improved labor market conditions for all workers. Gaddis and Pieters (forthcoming) indicate this boom did not translate into a closure of the gender wage gap, but it may have closed the gender utility gap if, as our analysis in the 5.2 suggests, women place high value on their ability to shift into safer employment.

3.2 Measurement Error

Failure to properly aggregate fatality data can substantially contaminate estimates of the VSL (Lalive 2003; Tsai et al. 2011; Kniesner et al. 2012). We address aggregation bias that results from pooling data across industries, across occupations, and across genders. Most prior work has been unable to measure fatality risk by gender, and instead use gender-pooled average industry and/or occupation rates. Measurement error arises since men and women in the same industry-occupation cell have different fatality risk, as indicated by Figure 2. This suggests that using gender-specific fatality rates could substantially improve measurement. However, as we show, the error distribution is non-classical, making it difficult to predict the direction of bias.

Of the studies that have been able to calculate female-specific fatality rates, the risk measures used are either industry or occupation aggregates. Our decomposition of the variance in female fatality rates reveals that 91% of the total variation in fatality rates occurs within 2-digit industries, and 89% of the total variation occurs within 3-digit occupations. This suggests that using either industry average rates or occupation average rates, as opposed to industry-by-occupation pairs, leaves the vast majority of the total variation as measurement error. The nature of this measurement problem itself is also gender-specific, as the variance components of female fatality rates are sharply different than those for men, for whom 68% of total variation occurs within 2-digit industries and only 34% occurs within 3-digit occupations.

To assess the importance of this measurement issue in understanding gender differences in compensating wage differentials, we estimate the distributions of measurement error that would arise if we were to calculate fatality rates using either the gender-pooled industry-by-occupation measure or the female-specific industry average. Figure 5a shows a kernel density estimate of the distribution of measurement error in the gender-pooled measure relative to 2-digit industry by 3-digit occupation rate, in our analysis sample of female workers, defined in Section 2. The distribution has a mean of –0.016, or about 129% of the female-specific mean fatality rate, again indicating that combined fatality rates significantly overstate the propensity for fatal injury among female workers. The distribution is also negatively skewed (-2.57) and leptokurtic (21.79).

Figure 5c plots the joint distribution of this measurement error and the average pooled fatality rate, for the sample of female workers. The lower bound of the slope of these points is –1, which occurs if only men die within an industry-occupation cell,

Figure 5: Distributions of Female Fatality Rate Measurement Error

Kernel Density of Measurement Error

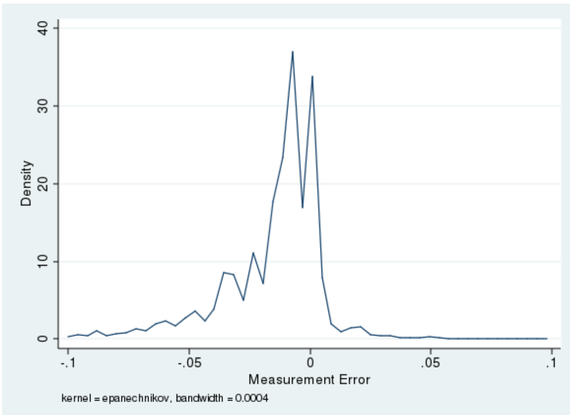

(a) Gender-Pooled Fatality Rate

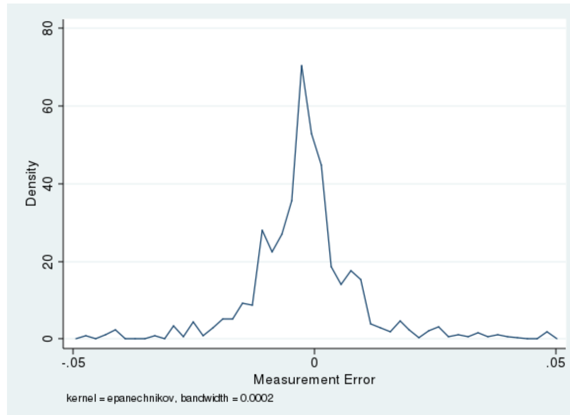

(b) Female Industry Average Fatality Rate

Joint Distribution of Measurement Error and Mis-Measured Fatality Rate

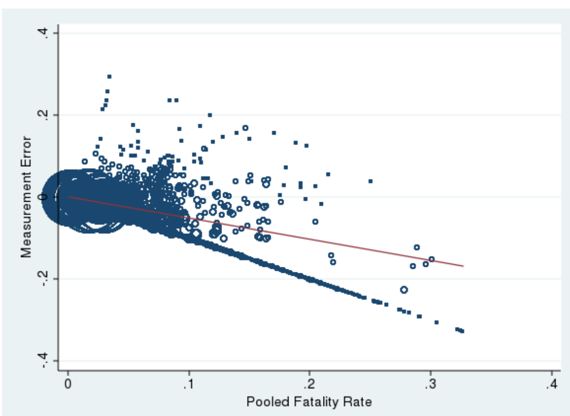

(c) Gender-Pooled Fatality Rate

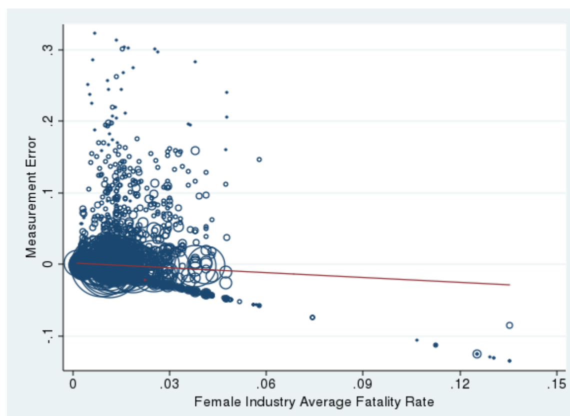

(d) Female Industry Average Fatality Rate

Notes: Figures 5a and 5b plot the kernel density of the gender by 3-digit occupation by 2-digit industry fatality rate minus the gender-pooled 3-digit occupation by 2-digit industry fatality rate (5a) and minus the female 2-digit average industry fatality rate (5b), all measured as three-year moving averages. Figures 5c and 5d plot the joint distributions of measurement error corresponding to Figures 5a and 5b, respectively, versus the gender-pooled 3-digit occupation by 2-digit industry fatality rate (5c) and minus the female 2-digit average industry fatality rate (5d), all measured as three-year moving averages. Sample includes women in cells with positive fatality rates and is truncated at the 99th percentile of fatality rates. Circle sizes in Figures 5c and 5d are proportional to the number of female workers in the industry-occupation cell.

so the female-specific fatality rate is zero and the measurement error in the female sample is the negative of the pooled fatality rate. The best-fitting OLS prediction has a slope of –0.52, suggesting that the error component introduced by gender aggregation is strongly negatively correlated with the key independent variable in a wage regression using pooled fatality rates. If within-cell gender differences in safety are priced in labor markets, the endogeneity bias caused by this measurement error component is likely to inflate the estimated compensating wage differential per unit of pooled fatal risk.

These measurement concerns are not eliminated by using gender-specific industry average fatality rates. Figure 5b plots the kernel density estimate of the measurement error distribution using female-specific 2-digit industry averages relative to our industry-occupation measure. There is again a large amount of mass in the measurement error distribution with absolute values in excess of the mean fatality rate of 0.012, and the distribution is even more skewed (5.05) and leptokurtic (53.84) than the gender-pooled measure. This, along with the evidence from the variance decomposition, suggests that using female-specific industry average rates, ignoring the substantial interaction components between industries and occupations, does not resolve this important measurement concern.

4 Empirical Model

Following the literature on the estimation of hedonic wage models (Rosen 1974), particularly for risk of death on the job as in Thaler and Rosen (1976), our analysis is based on models of the form:

$$y_{it} = x_{it}\beta + \gamma a_{c(i,t),t} + \eta_{it}. \tag{1}$$

Our data are structured as a panel over workers, but with dependence across workers employed in the same establishment. The dependent variable, y_{it}, is the log hourly wage. The key regressor, $a_{c(i,t)}$, is a measure of the risk of fatal injury when the worker is employed in the industry-occupation cell given by the function $c(i,t)$ which maps panel indices to cells. A key feature of our analysis is contrasting estimates of this model when $a_{c(i,t)}$ is based on our gender-specific industry-occupation measure of fatality risk relative to alternative measures from the literature. The vector of controls, x_{it}, includes measures of labor market experience, tenure, education, and job assignment described in more detail below.

We estimate three models with different, and nested, assumptions on the unexplained component of variation in wages, η_{it}. Following Lavetti and Schmutte (2016), our preferred specification is a generalization of the decomposition popularized by Abowd et al. (1999) in which

$$\eta_{it} = \theta_i + \psi_{J(i,t)} + \varepsilon_{it}.$$

In this model, the unobserved error depends on a component that is common to all observations for the same worker, θ_i, and a component common to each observation in the same establishment, $\psi_{J(i,t)}$. The index function $J(i,t)$ maps panel observations of the job held by worker i in year t to establishment identifiers.

Our variant, which we refer to as the *orthogonal match effects* (OME) model, allows for worker-establishment match effects that are arbitrarily correlated with observables in x_{it}, but which are orthogonal by construction to worker and firm effects (θ and ψ). This model is estimated in two steps, with the first step:

$$w_{it} = x_{it}\beta + \tilde{\gamma} a_{c(i,t),t} + \Phi_{i,M(i,t)} + \epsilon_{it}. \tag{2}$$

where $\Phi_{i,M(i,t)}$ denotes the match effect, where $M(i,t)$ indexes a match between worker i, establishment J, and occupation c in which worker i is employed in period t. Although $\tilde{\gamma}$ is identified in this model, the only source of identifying variation is intertemporal variation in fatality risk within jobs. There is reason to be suspect that within jobs wages may not respond as in a spot market to intertemporal variation in average occupational safety rates, which are often very small and potentially non-salient. Instead, our estimation strategy is to use variation in wage-risk pairs across jobs, while conditioning on the unobserved characteristics of jobs that could impact the decision of workers to switch jobs, in order to avoid estimation bias arising from potentially endogenous job mobility choices.

To do this, we purge the within-job intertemporal variation in wages from the dependent variable, and estimate the second stage model:

$$w_{it} - x_{it}\widehat{\beta} = \gamma a_{c(i,t),t} + Occ1_{it} + \theta_i + \psi_{J(i,t)} + \varepsilon_{it} \tag{3}$$

analogous to the two-way fixed effects model of Abowd et al. (1999). Identifying variation in this model comes from job changes in which a worker moves across industry-occupation cells, which could occur either within or across establishments. We also include controls for one-digit occupation codes, $Occ1_{it}$, to account for broad differences in job characteristics other than risk, such as whether the occupation tends to require manual labor. The advantage of this two-step specification is that it allows x_{it} to be arbitrarily correlated with worker, establishment, and job-match effects. The second stage model then allows fatality rates to be correlated with latent worker and establishment effects, addressing the form of endogenous job mobility that is likely to arise if workers climb a job ladder towards jobs at higher paying establishments. An additional advantage of this specification is that establishment effects also control all of the other establishment-level job amenities besides safety, mitigating a common and pernicious source of omitted variable bias.

The assumptions required by this model are substantially weaker than those in all previous studies of female compensating wage differentials and the VSL. Since worker and establishment wage effects are nuisance parameters, the key identifying assumption is $E\left[\Phi_{i,M(i,t)} a_{c(i,t),t} \,\middle|\, \theta_i, \Psi_{J(i,t)}\right] = 0$, which says that the fatality risk a worker takes on is not systematically associated with match-specific unobservables that affect the wage, after conditioning on worker and employer characteristics. This assumption is satisfied as long as, within a firm, wages only vary across jobs because of compensating wage differentials. Note that in estimation, we control directly for occupation heterogeneity in wages.

the workers' idiosyncratic match wage component must not be correlated with the industry-occupation cell average fatality rate.

Our OME specification nests two other, more common, specifications. The benchmark estimates in the literature of compensating wage differentials for female workers in Hersch (1998) and Leeth and Ruser (2003) use cross-sectional data and assume η_{it} is independent of fatality risk conditional on x_{it}. We will also estimate this specification to make direct comparison with their analysis. A key problem with this model is the presence of unobserved worker characteristics that affect earnings potential, and therefore the accepted level of risk. To address unobserved ability, Kniesner et al. (2012) estimate within-worker panel data models (for men only) for which $\eta_{it} = \theta_i + \varepsilon_{it}$. The identifying assumption in the within-worker model is that ε is uncorrelated with fatality risk after conditioning on observables included in x_{it} as well as unobserved, but stationary worker characteristics absorbed by θ_i. DeLeire, Khan and Timmins (2013)

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	0.551*	0.878*	0.476*	0.585*	0.447*
	(0.002)	(0.003)	(0.001)	(0.004)	(0.001)
VSL (million reais)	4.31*	5.66*	3.72*	3.77*	3.27*
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)
N	79,060,705	47,197,529	79,060,705	47,197,529	126,258,234
R-Sq.	0.966	0.972	0.966	0.972	0.967

Notes: Analysis sample includes dominant jobs of workers 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 FTFY workers in the three-year moving average window used to calculate fatality rates. '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. * Indicates significance at the 0.01 level.

also estimate VSLs for women and men in the US, using a structural approach to address for endogenous sorting into occupations based on idiosyncratic returns. Like us, they find much larger VSLs for both men and women after correcting for sorting.

Lavetti and Schmutte (2016) provide a very detailed derivation of the OME model specification, along with evidence that the assumptions underlying the pooled and within-worker model are violated, while those of the OME model are supported, for male workers in Brazil. Our primary focus in the present paper is on the implications of mis-measurement of fatality rates and differences in sorting on risk for estimates of the compensating wage differential for women. It turns out these implications are qualitatively robust to model specification, as shown in Section 5.

5 Results

Table 2 reports estimates of compensating wage differentials from our preferred OME specification. Our attention in this table is on the contrast between Columns (1) and (2), for which the fatality rate is measured using combined data from men and women, and Columns (3)–(5), for which the fatality rate is gender-specific. We first estimate the model separately by gender using the overall average industry by occupation fatality rate. We then discuss the implications of estimating the model using gender-specific fatality rates.

5.1 Estimates Based on Fatality Rate Measures Combining Genders

Column (1) presents the estimates for the male sample using the combined fatality rate measure. The coefficient on the fatality rate, 0.551 suggests that doubling the fatality rate relative to the mean (from 0.0726 to 0.1452 deaths per 1,000 FTFY workers) is associated with a 4% increase in wages. The corresponding point estimate for women is substantially higher, at 0.878.

A key concern with these estimates, given the evidence in Section 2, is that the higher coefficient for women may be driven by non-classical measurement error in the fatality rate, as depicted in Figure 5a. If the market compensating wage differential reflects this gender difference in job amenities, the change in fatality risk accepted by women when moving between jobs may be larger or smaller than what is calculated using the combined fatality rate measure. Since the within industry-occupation ratio of female to male fatality rates is lower for safer jobs, and rises with the fatality rate level, we might expect the "true" change in fatality rate to be larger than changes measured using combined-gender fatality rates. On the other hand, since Figure 5a shows that the measurement error in fatality rates in this model could be positive or negative, the sign of the bias in Column (2) is difficult to predict, and could potentially understate the true compensating differential for female workers.

The second row of the table converts the estimated coefficients on fatality risk into VSL estimates at the gender-specific mean log wage. 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 2,000 hours per FTFY worker with 1,000 workers, since γ is the wage effect of an increase in number of fatalities per 1,000 workers. The implied VSL for men, estimated using the male average wage, is 4.31 million 2003 reais, while for women the estimate is still higher, 5.66 million.

Aside from measurement error, there are other reasons the estimated compensating differential may be higher for women. The coefficient could be higher if women tend to hold different jobs than men, female labor is not perfectly substitutable for male labor, and the marginal worker in female jobs has stronger preferences for occupational safety than the marginal worker in male jobs. Alternatively, since the female average wage is lower, if the compensating wage differential per unit of risk is the same for men and women, this level mechanically represents a higher percentage of the average female wage, increasing the female coefficient in this model. Each explanation has a very different implication for how one might use these estimates to measure the marginal benefits of safety for conducting cost-benefit analysis of policies to improve safety. We investigate these alternative explanations more thoroughly in the sections that follow.

5.2 Estimates Based on Gender-Specific Fatality Rates

Columns (3)–(5) of Table 2 present estimates from the same model specifications, but using gender-specific fatality rates. Although the coefficient on the fatality rate remains higher for women, the gender gap in estimated VSLs is nearly eliminated when the fatality rate measures are improved. Since the samples are the same, so is the gender-gap in log wages. Hence the smaller gender gap in VSLs relative to columns 1–2 can only be caused by a change in the estimated coefficients on fatality risk. The estimated female coefficient is still higher than the male coefficient, at 0.59 versus 0.48, but the gender difference in average wages approximately offsets the difference in estimated coefficients, causing the VSL estimates to be nearly indistinguishable: 3.72 million for men (SE 0.01 million), and 3.77 million for women (SE 0.02 million).

If the preferences for occupational safety of the marginal worker, which determines the compensating wage differential in a competitive labor market, are similar across genders, this pattern of estimates would be consistent with women facing lower fatality rates in the same industry-occupation cells as men, and the compensating wage differentials reflecting this gender differences in safety. This implies that using gender-specific fatality rates alleviates substantial bias from measurement error, shrinking the estimated female VSL from 5.66 million to 3.77 million in our empirical setting. As a point of reference for interpreting these VSL figures, we find the literature suggests a range for the VSL in Brazil between 2.14 and 3.10 million 2003 reais. Our estimates are above the high end of this range, but not implausibly so. To obtain this range, we follow the benefit transfer approach recommended by Hammitt and Robinson (2011) as applied by Miller and Façanha (2016). This is necessary because there are no comprehensive studies reporting the VSL for Brazil based on estimates from hedonic wage equations. The benefit transfer method adjusts the VSL estimates for differences in per-capita income since these have large effects on the scale of willingness to pay for reducing the risk of death.[6]

5.3 Alternative Models

For comparison, Table 3 reports estimates based on the pooled (Column 1) and within-worker (Column 2) models. As discussed in Section 4, both models are nested within our OME specification, and are more consistent with the hedonic wage models for risk of fatal injury estimated based on cross-sectional and panel data in prior work. In the table's upper panel, we report estimates using fatality rates calculated by combining data from male and female workers, but estimating the model using only the sample of female jobs. The middle panel uses female-specific industry average fatality rates, omitting occupation-level variation to replicate the measure that has been used in several prior studies of the female VSL. The bottom panel uses female-specific industry-occupation fatality rates.

As discussed in Lavetti and Schmutte (2016) and Lavetti (2015), the worker effects specification amplifies the problems associated with endogenous job mobility choices, causing severe negative bias, similar to the patterns observed for men. Relative to the gender-pooled fatality rate estimates, our industry-occupation estimates are slightly smaller in every specification. This is consistent with evidence from previous studies comparing industry average rates to industry-occupation averages in cross-sectional models, such as Viscusi (2004).

[6] We implement the benefit transfer method as follows. First, we use the preferred range of VSL estimates reported by Viscusi (2015) in his meta-analysis of studies using the hedonic wage method. He gives lower bound of 7.6 million 2013, and an upper bound of 11.0 million U.S. dollars. We convert these to 2003 Brazilian reais using the OECD price index (OECD 2010). Hammitt and Robinson (2011) recommends converting the VSL using the formula:

$$VSL_{Brazil} = VSL_{US} \times \frac{\text{GNI per capita in Brazil}}{\text{GNI per capita in U.S.}}. \tag{4}$$

We obtain PPP-adjusted estimates of per capita gross national income for Brazil and the U.S. from The World Bank (2017). Strictly speaking, Hammitt and Robinson (2011) recommend adjusting the numerator in Equation (4) for the elasticity of willingness-to-pay with respect to income. We follow the World Bank recommendation to set this elasticity to 1. Combining these estimates yields our reported range of estimates for the VSL in Brazil between between 2.14 and 3.10 million 2003 reais.

Table 3: Comparison of Female Compensating Differential Estimates by Specification and Fatality Rate Measure

	Dependent Variable: ln($Wage$)		
	(1)	(2)	(3)
	Pooled	Worker Effects	OME
	Gender-Pooled Fatality Rate		
Fatality Rate (3-Yr MA)	0.584*	0.124*	0.878*
	(0.003)	(0.003)	(0.003)
VSL (million reais)	3.77*	0.80*	5.66*
	(0.02)	(0.02)	(0.02)
R-Sq	0.578	0.933	0.972
	Female-Specific Industry Fatality Rate		
Fatality Rate (3-Yr MA)	0.139*	–0.028*	4.441*
	(0.011)	(0.009)	(0.012)
VSL (million reais)	0.90*	–0.18*	28.64*
	(0.07)	(0.06)	(0.08)
R-Sq	0.577	0.934	0.972
	Female-Specific Ind.-Occ. Fatality Rate		
Fatality Rate (3-Yr MA)	0.421*	0.122*	0.585*
	(0.004)	(0.003)	(0.004)
VSL (million reais)	2.71*	0.79*	3.77*
	(0.02)	(0.02)	(0.02)
R-Sq	0.579	0.933	0.972
N	47,187,583	47,197,529	47,197,529
Worker Effects	N	Y	Y
Establishment Effects	N	N	Y

Notes: The estimation sample is restricted to women employed on their dominant jobs between ages 23-65, with 30 or more contracted hours per week. We exclude 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. Gender-Pooled Fatality Rate is the 2-digit industry by 3-digit occupation average rate for men and women combined. Female-Specific Industry Fatality Rate is the 2-digit industry average rate for women. Female-Specific Ind.-Occ. Fatality Rate is the 2-digit industry by 3-digit occupation average rate for women. All fatality rates are measured in deaths per 1,000 FTFY workers. We Winsorize the sample at the 99th percentile of overall fatality rates and at the 1st and 99th percentiles of the overall wage distribution in the full analysis sample. * indicates significance at the .01 level.

The importance of using occupation-level data on fatality rates is also clearly evident in the results. In the middle panel, which excludes this occupation-level variation, forcing it to be an error component, the estimates are dramatically different than in each of the other panels. Consistent with the variance decomposition showing that occupation-level variation in safety is much larger for women than for men, this model demonstrates the importance of accounting for occupations when estimating the VSL for women.

6 Sources of Gender Differences in Wages and Accepted Risk

In this section we assess the role of gender differences in safety and sorting on estimated gender gaps in wages and the value of statistical life. We have found that men and women have different wages because (1) they sort into jobs with different fatality risk (2) within the same job cell, women and men face different risks of death, and (3) women face a different marginal price for avoiding risk of fatal injury than do men. To characterize the role of these different channels, we use a standard Oaxaca-Blinder decomposition, where the sorting and behavioral differences are measured as differences in "endowments" and the higher price paid by women to avoid risk is the "coefficient" effect.

Because our analysis incorporates establishment-specific heterogeneity, as in Abowd et al. (1999), we also compare our analysis to the decomposition of gender wage gaps in Portuguese matched data in Card et al. (2016). They find that women sort into lower-paying firms than men, and are less well-paid than men in those firms. They attribute the latter "coefficient" effect to differences in bargaining power.

We conclude the section by asking what drives differences in the estimated VSL between men and women. Our analysis identifies measurement error, estimated compensating wage differentials (or prices) and differences in the overall gender wage gap as the three primary sources of variation.

6.1 The Role of Sorting and Prices in the Gender Wage Gap

Table 4 presents a decomposition of the gender wage gap in Columns (1) and (2). In Columns (3) and (4), we report an equivalent decomposition for just the gap in

Table 4: Oaxaca Decomposition of Gender Wage and Compensating Differential Gaps

	Gender Wage Gap Decomposition (1)	Share of Total (2)	Gender Compensating Differential Gap Decomposition (3)	Share of Total (4)
Total	18.27%		2.23%	
Fatality Rate Levels	3.39%	19%	2.55%	118%
Fatality Rate Prices	–0.16%	–1%	–1.44%	–67%
Fatality Rate Interactions	–0.73%	–4%	1.12%	52%
Establishment Effects	3.24%	18%		
Occupation Effects	–1.30%	–7%		
Worker Effects	–6.55%	–36%		
Within Match Time-Varying Endowments	21.00%	115%		
Other Within Match Differences	0.11%	1%		

Notes: All estimates are based on the OME model specification using the combined male and female analysis samples, with sample size 126,258,234.

compensating differentials. We estimate the OME separately for men and for women, and then compute a Oaxaca-Blinder decomposition using the model estimates.

6.1.1 The Role of Differences in Risk of Fatal Injury

Differences in occupational safety contribute to the wage gap in several ways. First, women sort into safer jobs. Second, conditional on being in the same job, women are less likely to die on the job than are men. Third, our estimates suggest that women pay a different price than men for avoiding risk. Table 4 shows that men earn 3.39 pp more than women due to their sorting into riskier jobs. This sorting across jobs on the basis of risk explains about 19% of the overall net gender wage gap. However, partially offsetting these effects, the gender wage gap shrinks by 0.89 pp because women pay a higher price to avoid a unit of fatality risk, and because of the interaction terms between prices and endowments, consistent with the patterns shown in Table 2. This difference shrinks the entire net gender wage gap by 5%, suggesting that compensation for occupation safety (and the other unobserved job amenities that may load onto the fatality rate) have meaningful effects on overall wage patterns.

Taken altogether, the sorting of women into safer jobs explains as much of the overall gender wage gap as does the sorting of women into jobs with lower firm-specific wage premia. Both processes may have the same root cause. If variation in pay across firms is explained by different amenities, as suggested by Sorkin (2016), it may be that women have stronger preferences for those amenities. However, Card et al. (2016) suggest these differences may also reflect gender differences in rent-sharing. In either case, the effect on wages of women sorting into safer jobs is attenuated by the fact that women pay more in foregone wages to reduce fatality risk. This means that conditional on being in jobs with equal risk, women receive a higher risk-premium than men.

In Columns (3) and (4), we decompose the components of the gender gap in the part of compensation driven by variation in fatal risk. Men receive about 2.23% higher compensating wage differentials than women. Of this, 2.55 pp come from men working in more dangerous occupations and industries, as shown in Figure 1a, and from men having higher fatality rates within cells. However the gap is reduced by 1.44 pp because women receive more compensation per unit of risk they face. Interaction effects between differences in job characteristics and differences in estimated coefficients contribute another 1.12 pp.

Finally, about 21 percentage points of the gender wage gap can be attributed to gender differences in experience, tenure, and experience interacted with race. Partially offsetting this, women have higher estimated worker fixed effects, and this combined with the occupation wage effects and gender differences in assignment to occupations narrows the wage gap by 7.85 pp.[7]

6.1.2 The Role of Inter-Establishment Differences

The overall gender wage gap in our data is 18.27 percent. Differences between men and women in the employer-specific contribution to pay increase the gender wage gap by 3.24 pp. These differences have two sources. Women sort into lower-paying establishments than do men. Second, conditional on being employed in the same firm, women receive lower employer-specific pay than do men, at least on average. The combined sorting and price effects explain 18 percent of the total gender wage gap. This is consistent with the Card, Cardoso and Kline (2016), who find these difference explain 20 percent of the gender wage gap in Portuguese data.

Figure 6 presents a plot of estimated male and female establishment effects within percentiles of the overall distribution of employment matches. The estimated effects are normalized so the overall average establishment effect has mean zero within each population. With this normalization, we observe female establishment effects are highly compressed relative to those for men. This compression is associated with women earning more than men at low-paying establishments, but less than men at high-paying establishments. Although the figure appears remarkably symmetric, as Table 4 shows the net effect of this compression is an increase in the gender wage gap of about 18%. The simple correlation between male and female establishment effects in our setting is 0.58, in line with estimates from Card et al. (2016) using Portuguese data.

6.2 The Relative Importance of Fatality Risk and Establishment Effects

Card et al. (2016) use data on establishment productivity to decompose gender differences in earnings attributable to sorting as opposed to bargaining, and find that the majority of the difference is explained by sorting. This finding is consistent with a

[7]Note that, as pointed out by Card et al. (2016), the person effects may not be well-identified relative to experience effects. It is perhaps more accurate to say that 72 percent of the total gap is due to the combined effects of experience interacted with race, tenure, occupation, and unobserved worker heterogeneity.

Figure 6: Binned Scatterplot of Within-Establishment Male versus Female Establishment Wage Effects

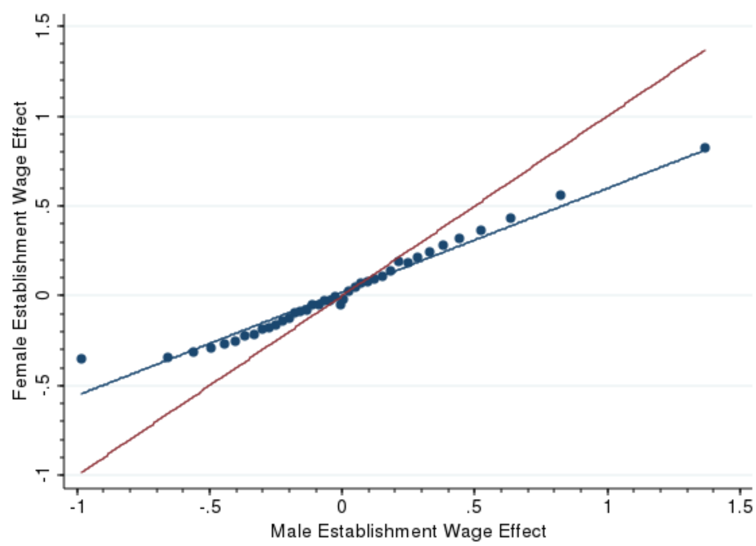

Notes: A binned scatterplot of the estimated male and female establishment effects. The red line is a 45-degree line, and the blue line is the predicted linear fit, which has slope 0.58.

much larger literature on gender differences in sorting on specific dimensions of jobs, such as hours, flexibility, competition, and the risk of variation in earnings, which was recently reviewed by Blau and Kahn (2016). One difficulty is assessing the relative importance of multiple dimensions on which workers may sort, is that sorting on the basis of unobserved establishment pay heterogeneity requires administrative matched data, which is generally lacking in detailed information about other job characteristics. Although our data too suffers from this limitation, we are able to quantify one form of job amenity in occupational safety, which we can use to compare gender differences in the importance of sorting on safety versus establishment pay.[8]

One way of describing the relative importance of sorting on safety and establishment pay is to quantify the value of total compensation gains in each of the two directions when workers switch jobs. If all job mobility decisions were primarily characterized by workers moving up the job-ladder towards establishments offering higher pay, one would expect to observe that both women and men tend to exit low paying establishments,

[8]Another observable amenity is hours worked, which Card et al. (2016) study and find that gender differences in this dimension do not explain the earnings gap associated with sorting.

Figure 7: Sorting on Safety Versus Establishment Pay

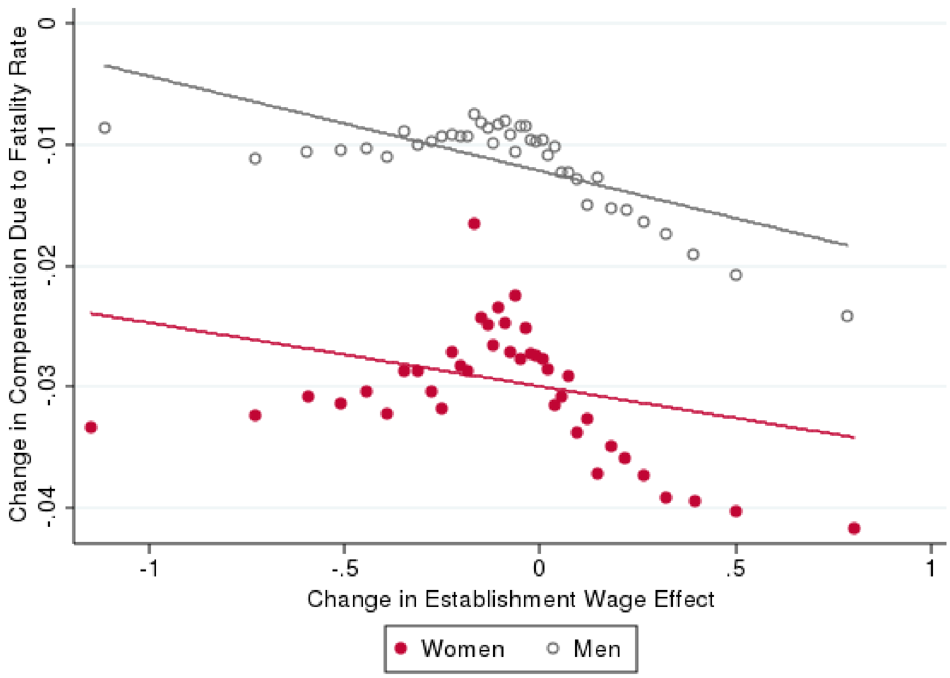

Notes: The figure presents a binned scatterplot of the change in establishment wage effects between origin and destination jobs relative to the change in compensation for fatal risk. The sample is restricted to origin jobs with a risk of death of at least 0.04 per 1,000 FTFY workers.

but have little systematic pattern of exit based on risk. In contrast, if workers only care about moving up a job-ladder in the safety dimension, one would expect workers to be more likely to exit risky jobs, regardless of establishment pay. More realistically, workers may care about both dimensions of the job ladder in potentially different proportions according to gender. In this case, men and women starting at similar origin jobs may tend to climb the ladder towards very different destination jobs.

To empirically quantify the relative importance of these two dimensions of sorting, we begin with the sample of male and female workers employed at jobs with fatality rates of at least 0.04 deaths per 1,000 FTFY workers.[9] Using a minimum risk level ensures that it is possible for workers to move in either direction, which avoids the impact of the zero lower bound constraint. We then plot the change in establishment wage effects relative to the estimated change in compensation for fatal risk, defined as the observed change in risk levels between the origin and destination job multiplied by the gender-specific estimate of the compensating wage differential per unit of risk from Table 2.

Figure 7 depicts this relationship. There are two striking patterns in this plot. The first is the plot for women is shifted downward relative to men, indicating that women make job changes with much larger improvements in safety than do men, and experience larger declines in earnings as a result. On average, when a women exits a job with a fatality rate of at least 0.04 her earnings decrease by about 2.9% because the destination job is safer than the origin job, reducing the total compensating wage differential. When men change jobs their earnings decrease on average by only 1.1% due to movements towards safer jobs. Moreover, regardless of whether women move up or down the establishment pay distribution, they have a much stronger tendency than men to move towards safer jobs. The second pattern is that a job change that results in a large increase in establishment pay tends to result in a larger improvement in safety, on average. This is consistent with the positive correlation between unobserved establishment pay and occupational safety documented by Lavetti and Schmutte (2016). We note that for women there is some non-monotonicity: larger decreases in establishment effects are also associated with larger safety improvements. Among those that move up the establishment pay distribution, female earnings decline at more than twice the rate of male earnings due to changes in compensation for risk, and women gain about 5.3 percentage points less in earnings on the destination job than men because their

[9]Very similar patterns are observed for any choice of fatality rate cutoff.

Figure 8: The Relative Share of Compensation Gains Due to Sorting on Safety

Notes: The figure presents a binned scatterplot of the absolute value of the change in compensation associated with a job change due to changes in safety divided by the sum of the absolute value of the compensation change due to safety and the absolute value of the change in establishment wage effects. The horizontal axis is the establishment wage effect at the origin job. The sample is restricted to origin jobs with a risk of death of at least 0.04 per 1,000 FTFY workers.

movements up the job ladder are directed more strongly in the safety dimension than in the establishment pay dimension.

To more directly depict the gender difference in the relative importance of sorting on the two dimensions, Figure 8 uses the same sample of job changes to plot the fraction of the total change in compensation (in absolute value) due to increases in safety as a function of the origin job establishment wage effect. For women, about 23% of the total change in compensation from an average job change is associated with movements towards safer jobs, which is roughly twice as large as the corresponding share for men. Women employed at lower-paying establishments also tend to have larger improvements in safety when changing jobs relative to compensation gains due to establishment pay.

The difference in the share of compensation gains attributed to sorting on safety is

Table 5: Decomposition of Sources of Gender Variation in VSL

	VSLs based on Female Mean Log Wage		
	Female Fatality Rate	Combined Fatality Rate	Male Fatality Rate
Female Sample	3.77	5.66	2.71
Combined Sample	2.84	3.51	2.87
Male Sample	2.90	3.55	3.07
	VSLs based on Male Mean Log Wage		
	Female Fatality Rate	Combined Fatality Rate	Male Fatality Rate
Female Sample	4.58	6.88	3.27
Combined Sample	3.44	4.26	3.49
Male Sample	3.52	4.31	3.72

Notes: VSL measured in millions of 2003 reais. All estimates are based on the OME model specification estimated on a 20% random sample of workers.

not primarily due to differences in the levels of total compensation gains. To see this, we regress the change in total compensation on establishment by occupation effects, and construct the residuals. We then repeat this step for the change in compensation due to safety, and due to establishment wage effects. Using the same sample of job changes depicted in Figure 8, we find that on average women gain about 0.9% in wages from sorting into higher paying establishments, while men gain about 1.5% on average, conditional on men and women starting in the same origin establishment and occupation pair. However, women on average give up 1.6% of their wages by moving into safer jobs on average, which is about 2.6 times higher than the corresponding value for men (0.6%). The sum of the absolute values of the two components of changes in compensation, which is the denominator in Figure 8, is actually larger for women than it is for men.

7 Decomposing Gender Variation in the VSL

Finally we present in Table 5 counterfactual estimates of the VSL to highlight the roles played by measurement error, differences in estimated compensating wage differentials, and differences in the overall gender wage gap. We present these estimates of the sensitivity of the VSL to counterfactual wage and fatality rate levels by gender to guide policymakers in adjusting cost-benefit analyses to account for the gender imbalance in the publication of VSL studies. As a benchmark for interpreting the variation in our counterfactual estimates, recall that in Section 5.2 we find VSL estimates in Brazil based on prior literature range between 2.14 and 3.10 million 2003 reais. The upper panel of the table shows the VSL estimates using the female average log wage, and either the female-specific, male-specific, or combined average fatality rate, applied to the gender-specific samples. The lower panel presents the same estimates using the male average log wage. For example, using the female average wage and female-specific fatality rate applied to the female sample gives a VSL of 3.77 million, as in Table 2. However, if we were to eliminate the gender wage gap and apply the male average wage, this increase in earnings would raise the female VSL to 4.58 million. If instead we were to apply the male fatality rate to the female sample, this would decrease the average compensation per unit of risk, and reduce the VSL to 2.71 million. The comparison suggests that each of these sources of gender differences, the safety gap and wage gap, have effects of similar magnitude on the female VSL. Failure to account for gender differences in the compensating wage differential causes the female VSL to be to be under-estimated. At the same time, failure to account for measurement error in fatality rates leads the female VSL to be overstated. Together, these findings suggest that appropriate modeling and measurement are equally critical to avoiding gender bias in policies using VSL in cost-benefit calculations.

8 Conclusion

Women and men may have different preferences for job amenities, including workplace safety, and this can help explain part of the observed gap in labor market earnings between men and women. Furthermore, the different preferences of men and women for safety should be factored into the evaluation of public policies that reduce the risk of injury and death. Research on these related topics has been hindered by poor measurement of workplace amenities and difficulties identifying compensating wage differentials. We develop the most detailed measures to date of gender-specific fatality rates. Using these data, we find that differences in workplace safety between men and women are an important contributor to the gender wage gap.

Our data indicate women working in the same occupation and same industry as men experience about 40% lower fatality rates. It is not clear from our data whether this differences arises from women being more careful when assigned the same tasks as men, or because there are unobserved differences in the assignment of tasks for jobs in the same industry and occupation. Our data very clearly show that women sort into industry-occupation cells with lower rates of fatal injury, consistent with DeLeire and Levy (2004) and suggesting women have stronger preferences for workplace safety.

Measurement error associated with pooling genders in the calculation of average fatality rates tends to inflate estimates of the value of statistical life. Removing this source of error, we estimate the female VSL in Brazil to be very similar to the male VSL, 3.77 versus 3.72 million reais. This similarity in estimates, however, arises from important offsetting gender differences that affect the calculation. The gender wage gap tends to increase the male VSL relative to the female VSL, since the VSL scales linearly with average wages in the standard log-linear hedonic wage model. However, women tend to earn a higher compensating wage differential per unit of fatal risk accepted, even after correcting for measurement error. These two components roughly offset each other, leading to a negligibly small gender-VSL gap. Estimating an Oaxaca decomposition of the gender compensating differential gap for fatal risk, we find similar offsetting effects: men earn significantly higher compensating wage differentials for fatal risk because they accept higher risk, but women earn more per unit of risk, narrowing the aggregate gender differential to 2.2 percent of overall average wages.

We also show new patterns on the nature of job sorting between men and women. The unique feature of our setting is that we use matched employer-employee data to estimate unobserved establishment-level contributions to pay, and are able to compare the importance of sorting on establishment pay to that of sorting on an important job amenity, occupational safety. Our analyses reveal that when women change jobs they sort much more strongly in the direction of improvements in safety, and experience smaller gains than men in establishment pay. This pattern is consistent with a model of the labor market characterized by a job-ladder, in which women climb in a different direction than men, favoring improvements in occupational safety over increases in pay. The gender difference in sorting along each of these dimensions tends to increase the gender wage gap.

In addition to contributing a more nuanced perspective on gender differences in occupational safety, we hope these findings will spur future research on this topic that may begin to ameliorate an overall gender-imbalance in labor market estimates of the VSL, which have substantial effects on public safety spending decisions. Our estimates are from a nationally-representative database of Brazil's formal-sector workforce. Brazil has a smaller gender wage gap than some developed countries, including the U.S., which may be due in part to differences in returns to education and experience. Likewise, workplace fatality risk in Brazil is higher as is the social context in which people choose whether to enter dangerous jobs. These observations urge caution when extrapolating directly from our analysis to other contexts.

References

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

Aldy, J. E. and Smyth, S. (2014). Heterogeneity in the value of statistical life, *Working Paper 20206*, National Bureau of Economic Research.

Barth, E., Goldin, C., Kerr, S. and Olivetti, C. (2017). The expanding gender earnings gap: Evidence from the LEHD-2000 Census, *American Economic Review: Papers and Proceedings* .

Bartolucci, C. (2013). Gender wage gaps reconsidered a structural approach using matched employer-employee data, *Journal of Human Resources* **48**(4): 998–1034.

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

Blau, F. and Kahn, L. (2016). The gender wage gap: Extent, trends, and explanations, *Forthcoming at The Journal of Economic Literature* .

Card, D., Cardoso, A. R. and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women, *The Quarterly Journal of Economics* **131**(2): 633.

Card, D., Heining, J. and Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality, *The Quarterly Journal of Economics* **128**(3): 967–1015.

Cortes, P. and Pan, J. (2017). Occupation and gender, *IZA Working Paper 10672* .

Costa, D. L. and Kahn, M. E. (2004). Changes in the value of life: 1940-1980, *Journal of Risk and Uncertainty* **29**(2): 159–80.

DeLeire, T., Khan, S. and Timmins, C. (2013). Roy model sorting and nonrandom selection in the valuation of a statistical life, *International Economic Review* **54**(1): 279–306.

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

Gaddis, I. and Pieters, J. (forthcoming). The gendered labor market impacts of trade liberalization: Evidence from Brazil, *Journal of Human Resources* .

Hammitt, J. K. and Robinson, L. A. (2011). The income elasticity of the value per statistical life: Transferring estimates between high and low income populations, *Journal of Benefit-Cost Analysis* **2**(1): 1–29.

Hersch, J. (1998). Compensating differentials for gender-specific job injury risks, *The American Economic Review* **88**(3): 598–607.

Kniesner, T. J., Viscusi, W. K., Woock, C. and Ziliak, J. P. (2012). The value of a statistical life: Evidence from panel data, *The Review of Economics and Statistics* **94**(1): 74–87.

Lalive, R. (2003). Did we overestimate the value of health?, *Journal of Risk and Uncertainty* **27**(2): 171–193.

Lavetti, K. (2015). The estimation of compensating wage differentials: Lessons from the *Deadliest Catch*.

Lavetti, K. and Schmutte, I. M. (2016). Estimating compensating wage differentials with endogenous job mobility, *Working Paper 29*, Labor Dynamics Institute, Cornell University.

Leeth, J. D. and Ruser, J. (2003). Compensating wage differentials for fatal and nonfatal injury risk by gender and race, *Journal of Risk and Uncertainty* **27**(3): 257–77.

Merlino, L., Parrotta, P. and Pozzoli, D. (2014). Gender differences in sorting, *Research Memorandum 022*, Maastricht University, Graduate School of Business and Economics (GSBE).

Miller, J. and Façanha, C. (2016). Cost-benefit analysis of brazil's heavy-duty emission standards (p-8), *white paper*, The International Council on Clean Transportation.
URL: *http://www.theicct.org/cost-benefit-analysis-brazil-HDV-emission-standards-p-8*

OECD (2010). Main economic indicators - complete database. Accessed: 2017 March 29.
URL: *http://dx.doi.org/10.1787/data-00052-en*

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition, *Journal of Political Economy* **82**(1): 34–55.

Sorkin, I. (2016). Ranking firms using revealed preference, *2016 Meeting Papers 66*, Society for Economic Dynamics.

Thaler, R. H. and Rosen, S. (1976). The value of saving a life: Evidence from the labor market, *Household Production and Consumption*, National Bureau of Economic Research, Inc, pp. 265–302.

The World Bank (2017). GNI per capita, PPP (constant 2011 international $). Accessed: 2017 March 29.
URL: *http://data.worldbank.org/indicator/NY.GNP.PCAP.PP.KD*

Tsai, W.-J., Liu, J.-T. and Hammitt, J. K. (2011). Aggregation biases in estimates of the value per statistical life: Evidence from longitudinal matched worker–firm data in taiwan, *Environmental & Resource Economics* **49**(3): 425–443.

Viscusi, W. K. (2004). The value of life: Estimates with risks by occupation and industry, *Economic Inquiry* **42**(1): 29–48.

Viscusi, W. K. (2015). The role of publication selection bias in estimates of the value of a statistical life, *American Journal of Health Economics* **1**(1): 27–52.

Woodcock, S. (2008). Wage differentials in the presence of unobserved worker, firm, and match heterogeneity, *Labour Economics* **15**: 771–793.