Estimating Compensating Wage Differentials with Endogenous Job Mobility*

Kurt Lavetti†
Ian M. Schmutte‡

October 23, 2017

Abstract

We demonstrate a strategy for using matched employer-employee data to correct endogenous job mobility bias when estimating compensating wage differentials. Applied to fatality rates in the census of jobs in Brazil between 2003-2010, we show that common approaches to eliminating ability bias amplify bias from nonrandom assignment to jobs. We establish empirical conditions under which our estimates are equivalent to preferences in a model of frictional search over differentiated firms, and show the identifying conditions are supported in Brazil. Our results suggest that standard panel models are misspecified, and that our proposed model eliminates latent ability and endogenous mobility biases.

*We are grateful for discussion and helpful comments from John Abowd, David Blau, Molly Candon, Chris Cornwell, Andrew Friedson, Kaj Gittings, Joni Hersch, Henry Hyatt, Lauren Jones, Josh Kinsler, Pat Kline, Francis Kramarz, Mark Kutzbach, Rick Mansfield, Jesse Rothstein, Meghan Skira, Chris Taber, and seminar participants at the Berkeley labor lunch, the U.S. Census Bureau, EALE, LERA, Ohio State, NYU, University of Illinois Urbana-Champaign, University of Illinois-Chicago, University of Kentucky, University of Maryland, University of South Carolina, University of Georgia, University of Wisconsin-Madison, and WEAI. Jason Rivera and Rodrigo Saurin provided helpful research assistance.

†Ohio State University, Department of Economics. `lavetti.1@osu.edu`

‡University of Georgia, Department of Economics. `schmutte@uga.edu`

1 Introduction

Despite a vast empirical and theoretical literature, there is little consensus regarding whether and how job amenities affect wages. Improved measurement of compensating wage differentials can shed light on many fundamental questions in labor economics. For example, tradeoffs between wages and job amenities affect our understanding of the welfare and policy consequences of earnings inequality. Estimates of compensating wage differentials are also frequently used to evaluate public policies. Wage differentials for the risk of fatal injury, for example, are used to calculate the value of statistical life, which affects tens of billions of dollars in federal spending annually on public safety policies in the US alone. Finally, a leading explanation for the pervasive difficulty in measuring compensating differentials is that there are systematic imperfections in the labor market. Thus, the study of compensating wage differentials can also contribute to our understanding of the competitive structure of the labor market. Given its prominent role, it is not surprising that the concept of wage differentials was one of the first ideas in economics, and that it remains one of the most widely studied.

The main challenge for estimation is that workers sort into jobs on the basis of factors that are not observed in most datasets. Rosen's seminal work provided conditions under which the equilibrium compensating differentials and underlying preferences of workers for job amenities could be recovered from observational data on wages and amenities when markets are competitive (Rosen 1974). Subsequently, Hwang, Reed and Hubbard (1992) showed that sorting on ability should, in theory, induce a downward bias in cross-sectional estimates of compensating wage differentials for disamenities. However, the standard panel data correction for unobserved worker ability suggests the opposite—that cross-sectional estimates are upward biased (Brown 1980; Kniesner, Viscusi, Woock and Ziliak 2012). The theory of hedonic search (Hwang et al. 1998; Lang and Majumdar 2004) argues this apparent contradiction arises from unobserved differences in pay across firms. A recent empirical literature imposes the structure of these models on the joint distribution of wages and job durations, but lacks sufficient information to account for unobserved firm characteristics or job amenities (Bonhomme and Jolivet 2009; Dey and Flinn 2005; Sullivan and To 2014).

Our main contribution is to show that when firms offer differentiated jobs, the direct effect on wages from moving between jobs with different amenities can be identified using matched longitudinal employer-employee data to control for establishment heterogeneity. We apply our approach to estimate compensating wage differentials for occupational fatality risk using comprehensive administrative data from Brazil between 2003–2010, among the largest and most detailed datasets ever used to study this topic. These detailed data also allow us to specify and assess whether the conditions needed for identification are satisfied. We demonstrate that attempts to alleviate bias by controlling for worker effects alone actually isolates variation arising from endogenous mobility across employers with different pay, causing a net *increase* in bias. Finally, we develop a model with job search across differentiated firms in which our wage equation represents optimal equilibrium firm behavior. The model clarifies the threats to our identification strategy, and provides empirical conditions under which our estimates can be interpreted as measures of worker preferences, which is required for consistency in the use of compensating differentials in many policy applications.

Our benchmark empirical model is an adaptation of the two-way fixed effects model introduced by Abowd, Kramarz and Margolis (1999), extended to include a restricted form of job-match effects. Despite the fact that controlling for worker effects alone seems to exacerbate bias, we find adding establishment effects is extremely effective in matching the predictions of the basic hedonic search model. If endogenous mobility bias arises from workers searching across employers with different compensation policies, within-worker estimates should be biased downward relative to estimates from our benchmark model that controls for worker and employer heterogeneity. The intuition behind this is straightforward: when workers change jobs to increase utility, they often receive simultaneous increases in wages and reductions in fatality rates, which introduces wrong-signed contamination.

We observe exactly this pattern in Brazil. Our estimate from the standard cross-sectional hedonic wage model indicates that moving from the 25th to the 75th percentile of the fatality rate distribution is associated with an increase in wages of about 2.7 percent. Adding worker effects, the corresponding estimate falls by an order of magnitude to 0.4 percent. However, when we add controls for establishment-level heterogeneity in wages, the estimated compensating wage differential increases back in the direction of the cross-sectional estimate, to 1.7 percent. This model relaxes several restrictive assumptions, allowing for the possibility that workers' sorting and job-mobility decisions may be correlated with unobserved characteristics of the origin and destination establishments.

While our basic results are consistent with theory, they beg the question of what economic objects of interest our parameter estimates actually identify. To illuminate the conditions for identification, we present a model that integrates hedonic search into the differentiated firms framework introduced by Card, Cardoso, Heining and Kline (2016). As in Hwang et al. (1998), our model shows that endogenous mobility bias can arise from a correlation between wages and an unobserved increment to utility offered by a job, even after controlling for employer heterogeneity. In our model, however, the extent of bias is an empirical question: if the omitted error component is separable in worker, employer, and occupational heterogeneity then our empirical specification can be interpreted as providing unbiased estimates of preferences. However, if match-specific utility premia are important determinants of job assignment in Brazil, then our empirical model will also yield biased estimates.

This theoretical concern relates directly to assumptions required by the empirical model. Our estimator is identified under an exogenous mobility assumption that requires the variation in fatality risk between jobs offered by the same employer to be exogenous with respect to residual variation in wages, holding worker and employer-specific factors fixed. We clearly reject the analogous exogeneity assumption in the worker effects model, and perform several residual diagnostic analyses to test whether mobility is conditionally exogenous in our preferred benchmark model. We apply the model specification tests suggested by Card, Heining and Kline (2013), and show that in Brazil, like Germany, there is little evidence against the assumption of exogenous mobility in our preferred model.

Our paper provides a bridge between the structural, theoretical, and reduced-form literatures on compensating differentials. Specifically, we show the statistical decomposition of wages originating with Abowd et al. (1999) does an extremely good job of matching the predictions of the basic hedonic search model, and in explaining the covariation between wages and job characteristics, in much the same way Card et al. (2013) did in the context of labor market inequality. Much of the recent empirical work on compensating differentials uses cross-sectional and panel data to estimate structural models of hedonic search (Bonhomme and Jolivet 2009; Dey and Flinn 2005; 2008; Villanueva 2007; Sullivan and To 2014) and Roy-style sorting (DeLeire et al. 2013). An emerging literature addresses models of compensating differentials using matched employer-employee data. In very innovative recent papers, Sorkin (2016) and Taber and Vejlin (2016) seek to explain how much variation in matching outcomes, job duration, and wages can be rationalized by compensating differentials. In these analyses, unlike our paper, job amenities are not measured; the presence of amenities is inferred from variation in outcomes, and the object of interest is a wage variance component rather than a wage differential for any specific amenity. Like us, Lalive (2003) and Tsai et al. (2011) estimate hedonic wage models using matched employer-employee data with observed firm-level amenities. However, the emphasis in both papers is on warning about the consequences of aggregation bias associated with measuring amenities using coarse industry averages, a flaw that does not affect our data. (Dale-Olsen 2006) also uses matched data to estimate worker willingness to pay from data on job durations, following an approach introduced by Gronberg and Reed (1994). Inspired by this work, we also consider duration data in a set of model extensions and sensitivity analyses.

We begin with an overview of the different hedonic wage models commonly estimated in the literature, and discuss how our empirical specification addresses the sources of bias that afflict them. Next, we briefly introduce our search model, derive the equilibrium wage equation, and connect the theoretical parameters to our empirical estimand. From there we introduce the data, present our main empirical results, and empirically evaluate the key identifying assumptions. The next section implements an IV strategy using the network structure of the data to establish that our main results are not sensitive to relaxing the key identifying assumptions. We consider several extensions to relax some of our parametric assumptions, show that our estimates are robust to different types of job mobility, including job changes caused by mass displacement events, and consider the relationship between job amenities, wages, and job durations. The paper concludes with a consideration of the broader implications of our findings for policy evaluation and labor market analysis.

2 Estimation Challenges and Empirical Model

Theory is a crucial guide for understanding the statistical assumptions underlying empirical hedonic wage models. We begin here by describing the two key challenges to identifying compensating differentials that our paper addresses: bias due to unmeasured ability, and bias due to unmeasured job quality. The basic analysis of these biases delivers clear intuition on how estimates of compensating wage differentials will change as we control for worker- and then employer-specific unobservables.

2.1 Ability Bias

Economists have long understood that unobserved worker ability can severely bias estimates of compensating wage differentials in cross-sectional data (Brown 1980; Thaler and Rosen 1976). Figure 1 depicts this basic idea graphically. Worker 1 has preferences represented by the indifference curve, u_1, and chooses a wage-risk combination (w_1, R_1) to maximize utility along 'offer curve 1,' the the upper envelope function of firms' wage-risk iso-profit functions. If workers have equal ability, variation in (w, R) pairs arises because workers with different preferences sort across jobs along the offer curve to maximize utility. In this simple case cross-sectional variation in wage-risk pairs identifies the hedonic pricing locus.[1]

A threat to this empirical strategy is that workers may differ in unmeasured ability. Suppose the representative firm can operate two offer curves that generate equal profit: 'offer curve 1' for workers of low ability and 'offer curve 2' for workers of relatively high ability (Hwang et al. 1992). In the figure, low ability workers choose (w_1, R_1) on indifference curve u_1 while high ability workers choose (w_2, R_2) on indifference curve u_2. Variation in wage-risk pairs now has two sources—variation along each offer curve, due to sorting, and

[1]In their comprehensive review of 32 studies that estimate compensating wage differentials for occupational fatality risk in the U.S., Viscusi and Aldy (2003) report that all but one relied upon this basic cross-sectional model for identification.

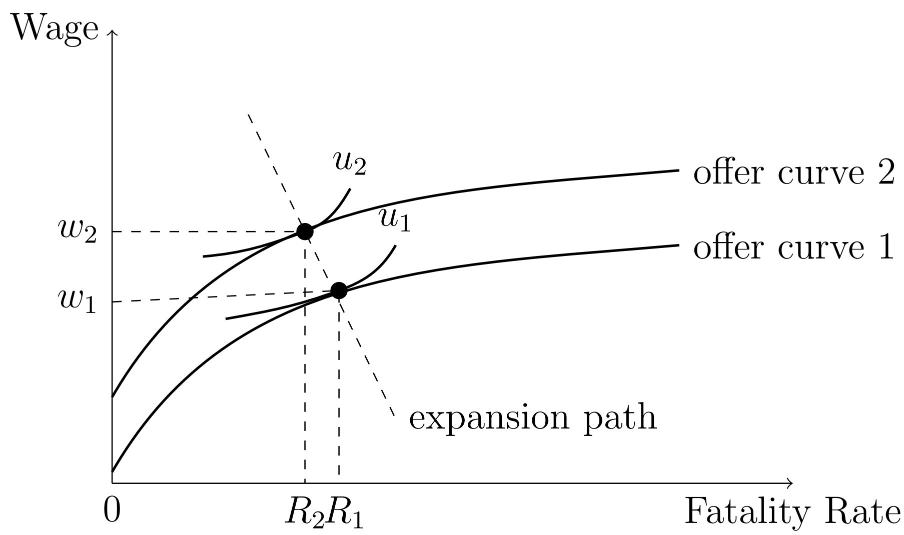

Figure 1: Equilibrium Wage-Risk Relationships

variation along the expansion path, driven by ability. If safety is a normal good, high ability workers trade off some of their higher earning power for reduced risk, causing the expansion path to slope downward even though both indifference curves slope upward. Variation along the offer curves, arising from differences in preferences for safety, is needed to identify the compensating wage differential, but the observed variation in accepted wage-risk pairs is contaminated by variation along the expansion path arising from differences in ability.[2]

Although this economic intuition is straightforward, it contradicts most empirical evidence. Brown (1980) and Kniesner et al. (2012) use panel data to estimate compensating wage differentials controlling for fixed worker effects and find much smaller estimates than comparable cross-sectional models produce, the opposite of what Figure 1 suggests. Hwang et al. (1992), in contrast, add proxies for unobserved ability to the cross-sectional model of Thaler and Rosen (1976) and find, consistent with theory, that doing so increases the estimated compensating differential for fatal occupational injury by a factor of 10 (the simulated bias is negative for disamenities and positive for amenities).

2.2 Endogenous Mobility Bias

This apparent contradiction can arise if the within-worker model exacerbates bias due to the endogeneity of workers decisions to move between firms offering different levels of total

[2] A closely related source of bias can arise if firms face imperfect competition for labor. If firms can attract workers by increasing utility, then they will choose to do so through a cost-minimizing combination of wages and job amenities. Hence, workers with high wages are also employed in less dangerous jobs because both job characteristics are correlated with the firm's demand for labor.

compensation (or utility). To see why, we can reinterpret Figure 1 as depicting the choices of one worker moving from job (w_1, R_1) to job (w_2, R_2). This type of variation is likely to occur in the presence of search frictions, for example if it takes time for a worker to find jobs offering higher utility as in Hwang et al. (1998), or if workers and firms learn about ability, match quality, or comparative advantage over time as in Gibbons and Katz (1992) and Gibbons, Katz, Lemieux and Parent (2005).

The movement from bundle (w_1, R_1) to (w_2, R_2) involves a simultaneous increase in the wage and a decrease in the fatality rate, and implies a wrong-signed compensating wage differential. If this type of variation has meaningful impact on equilibrium wages, it will be partially contained in the error term of panel data models like:

$$w_{it} = x_{it}\beta + \gamma R_{c(i,t),t} + \theta_i + \nu_{it}. \tag{1}$$

where w_{it} is the log wage of worker i at time t, x_{it} contains observable characteristics, $c(i,t)$ indicates the industry-occupation cell of the job at which worker i was employed in period t, $R_{c(i,t),t}$ is the fatality rate associated with that job at time t, and θ_i is a worker effect. Since heterogeneous workers are spread out along both offer curves, each of which is positively sloped, OLS estimates (omitting θ_i) tend to suggest a positive relationship between wages and fatality rates. However, including worker effects in the model causes variation along each offer curve to be absorbed by θ_i, leaving only variation along the expansion path to identify γ. As depicted in Figure 1, this wrong-signed component of variation is potentially even more problematic than its converse, increasing aggregate bias relative to a pooled OLS model.

This form of variation is not a hypothetical concern. Abowd, McKinney and Schmutte (2015) show empirically that workers tend to exit jobs at low paying establishments at a higher rate, even after conditioning on unobserved ability. Woodcock (2008) estimates that among workers in the US who experience job-to-job transitions, about 60% of their earnings growth is due to sorting into firms that pay higher average earnings to all workers for unobserved reasons.

2.3 Empirical Models

A natural progression of this panel approach to estimating compensating wage differentials in matched employer-employee data is to introduce amenities into the following AKM two-way fixed effects model (Abowd et al. 1999):

$$w_{it} = x_{it}\beta + \gamma R_{c(i,t),t} + \theta_i + \Psi_{J(i,t)} + \varepsilon_{it} \tag{2}$$

which includes an establishment effect, $\Psi_{J(i,t)}$, where the subscript indexes the establishment J employing worker i in period t. The inclusion of $\Psi_{J(i,t)}$ in the model allows for variation in pay across jobs in the same establishment and accommodates arbitrary sorting on the basis of risk and the worker effect. Another advantage of this specification is that $\Psi_{J(i,t)}$ absorbs the effect on wages of any unobserved firm-level amenities, substantially relaxing both data requirements and exogeneity assumptions relative to cross-sectional and within-worker models.[3] If this model is properly specified then γ can be interpreted as measuring the effect on wages of a change in fatality risk, holding other unobserved establishment-level amenities fixed.

Orthogonal Match Effects Model

In practice, we estimate an extension of the basic AKM specification that allows for correlation between the time-varying observables and arbitrary match effects. This extension, which we call which we call the orthogonal match effects (OME) model, is closely related to those introduced by Woodcock (2008) and Barth, Bryson, Davis and Freeman (2016). In discussing our results, and in Appendix Tables A.5 and A.6, we show that our findings are qualitatively and quantitatively robust to a wide range of alternative specifications, including estimating the more conventional AKM specification.

We fit the OME model in two steps. In the first step, we project wages onto observed time-varying controls (experience dummies, year dummies, and fatality risk), and an arbitrary match-specific effect.

$$w_{it} = x_{it}\beta + \tilde{\gamma} R_{c(i,t),t} + \Phi_{i,Jk(i,t)} + \epsilon_{it} \tag{3}$$

$\Phi_{i,Jk(i,t)}$ 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 step, we remove the effect of $x_{it}\widehat{\beta}$ from log wages and estimate the AKM model on the transformed dependent variable. The second stage model

$$P_{it} = \pi_{k(i,t)} + \gamma R_{c(i,t),t} + \tau_t + \theta_i + \Psi_{J(i,t)} + \xi_{it} \tag{4}$$

includes fatality risk, 1-digit occupation effects $\pi_{k(i,t)}$, year effects τ_t, as well as worker and establishment effects. $P_{it} \equiv w_{it} - x_{it}\widehat{\beta}$.[4]

[3] Recent working papers by Taber and Vejlin (2016) and Sorkin (2016) study the related, but distinct, problem of measuring the part of variation in establishment effects associated with unobserved amenities.

[4] The inclusion of year effects in both steps is important, because $\tilde{\tau}_t$ in the first step controls only for unobserved time-varying factors that influence the rate of growth of wages within a job over time, while τ_t accounts for additional unobserved intertemporal variation across jobs, such as the effect of beginning a job during an economic downturn.

The OME model has potential advantages over AKM. First, the estimated experience and year effects in the first stage are purged of any correlation with arbitrary match effects. Second, estimating the model in two steps helps isolate the variation used for identification. In the first-stage model, we estimate a coefficient on risk that is identified from the very slight within-match variation in fatality risk, which is only 3% of the total variation. We find that these small intertemporal changes are uncorrelated with wage changes. Our second-stage estimates of the compensating wage differential are therefore identified primarily from variation across, rather than within, jobs.

2.4 Corner Solutions

Conceptually, the expansion path in Figure 1 could hit the vertical axis, and a worker with sufficiently high ability may maximize utility by choosing a job at the corner solution, with zero probability of death. Likewise, firms that have a particularly high demand for labor will offer jobs that are as safe as possible. We therefore expect to find the highest-paid workers and the highest-paying firms concentrated in jobs with little to no fatality risk. Although it is unlikely that a worker's ex ante expectation of R is exactly zero, empirically about ten percent of jobs are in industry-occupation cells that experienced zero fatalities, since fatalities are rare events.

In estimation, we address corner solutions by including in every specification a dummy variable that equals 1 when $R_{c(i,t),t} = 0$, allowing for potential discontinuities in wages at zero risk. As shown by Caetano (2015), finding a discontinuous response at the corner is evidence of model mis-specification. Because a choice of fatality risk can never be exactly zero, we cannot apply her diagnostic test directly, nor the extensions developed in Caetano and Maheshri (2013). Nevertheless, the underlying economic and econometric intuition is helpful for interpreting our data, and suggests that the presence of endogeneity may be manifest as a non-monotonicity in the wage-risk profile in the vicinity of zero fatalities. We find that jobs with zero observed fatalities pay about seven percent higher wages, and show that this wage discontinuity is driven by disproportionately high-wage workers and high-wage firms.

2.5 Discussion

Intuitively, the AKM and OME models will be effective in correcting for endogenous mobility bias in the worker effects model if omitted firm-specific compensation is the primary source of un-modeled variation along the "expansion path" in Figure 1. This leads us to predict that the estimated compensating wage differential estimated from Equation (2) will be larger (more positive) than that estimated from the worker effects model. This prediction is supported in our data, as we show in Section 5.

The crucial question then becomes whether estimates of γ based on Equation (2) are more reliable than those based on the pooled or worker effects models. After all, we have argued that models controlling for unobserved worker characteristics may be more biased than models that do not. Why is adding firm effects to that model necessarily an improvement? This is really an econometric question: after controlling for establishment-specific heterogeneity, are the observed movements of workers between jobs with different levels of risk exogenous to wage residuals?

Moreover, it is not immediately obvious what relationship the estimated price of risk from the OME model bears to the model outlined in Figure 1: is it a measure of preferences, technology, market prices, or of variation along some other expansion path? Though we view this paper as being empirically driven, we now present a concise theoretical model that builds on this intuition and offers an example of a structural interpretation that connects the parameters in the OME model to these primitive theoretical objects.

3 A Search Model with Imperfect Competition and Amenities

We introduce a simple model of on-the-job search in which wages are affected by unobserved worker, firm, and job-level heterogeneity. Our goal is to demonstrate how matched longitudinal employer-employee data can reduce bias in the estimation of compensating wage differentials in frictional labor market settings. We derive an equilibrium structural expression for log wages, relate worker preferences to the derivative of wages with respect to fatal risk, and characterize a match-specific error term that explicates threats to identification and serves as a guide for our empirical diagnostic analyses.

Our model extends the "differentiated firms" framework of Card et al. (2016) by incorporating search frictions and endogenous choices over job-level safety across jobs in different firms and different occupations. Time is discrete, and workers and firms live forever. A fixed population of workers $i \in \{1, \ldots, N\}$ supply a single unit of labor inelastically and choose whether and where to work in each period after receiving job offers. Each worker has a fixed level of skill, $s(i) \in \{1, \ldots, S\}$. Workers receive offers with the same probability regardless of whether and where they are employed, and offers expire at the end of each period. Hence, workers will always choose whichever job offer provides the highest instantaneous utility. Worker i, when employed in occupation k by firm j receives utility $u_{ijkt} = \bar{u}_{sjkt} + \epsilon_{ijkt}$. The firm can control $\bar{u}_{sjkt}$, but ϵ_{ijkt} represents the worker's idiosyncratic taste for the job in period t, which is unknown to the firm. We assume ϵ_{ijkt} is distributed Type 1 Extreme

Value.

Job offers are made by firms, of which there is a large fixed population $j \in \{1, \ldots, J\}$. Firms are distinguished by industry, $b(j) \in \{1, \ldots, B\}$, and are exogenously endowed with firm-specific amenity a_j and productivity T_j, all of which may be arbitrarily distributed. Firms can offer employment in each of a fixed, finite set of occupations, $k \in \{1, \ldots, K\}$. Each occupation has an exogenous amenity d_k and an endogenous risk of death, R_{jkt}, that is chosen by the firm.

Firms choose wages and risk to attract workers, who receive indirect utility $\bar{u}_{sjkt} = f(w_{sjkt}, R_{jkt}) + g_s(a_j, d_k)$. Following the hedonic search literature, workers have common preferences over wages and risk. The function $f(w_{sjkt}, R_{jkt})$ is increasing and concave in w and decreasing and convex in the disamenity R. Regarding $g_s(a_j, d_k)$, which gives preferences over the exogenous firm- and occupation-specific amenities, we only assume it is increasing in both arguments.

The profits of firm j in period t are

$$L_{sjkt} \left[Q_{sjkt} - C_{bk}(w_{sjkt}, R_{jkt}) \right] \tag{5}$$

where L_{sjkt} is total employment of type s labor, Q_{sjkt} is revenue per worker, and $C_{bk}(w_{sjkt}, R_{jkt})$ is the unit cost of labor, which varies by industry and occupation. This allows for heterogeneity across firms in technology for providing safety. The unit cost function is increasing and convex in w and decreasing and concave in R.

3.1 The Labor Market

In each period four events take place: (1) firms choose wage and amenity offers (w_{sjkt}, R_{jkt}) to maximize expected steady-state profits; (2) firms make offers with certainty to all of their current (inside) workers, and with probability λ to each outside worker; (3) workers obtain a new draw from the idiosyncratic preference distribution ϵ; (4) workers accept the available offer that yields the highest utility.

Given the assumption of EV1 preferences, that the number of firms is large, and that each firm employs a negligible share of each type of worker, the probability that a firm's offer is accepted can be expressed by:

$$p_{sjkt} = K_s \exp(\bar{u}_{sjkt}), \tag{6}$$

where K_s is a normalizing constant. This follows from integrating the standard conditional logit choice probability over all possible consideration sets. The expression approximates the true probability with an error that scales with the firm's share of the labor market, as shown in Appendix B.1.

We now consider the firm's decisions about employment of a particular type of labor, s, in a particular occupation, k. We therefore drop subscripts except where needed for clarity. In steady-state, each firm will choose the same offer in each period, and the law of motion for the stock of employment is $L_{t+1} = pL_t + \lambda p(N - L_t)$ where N is the size of the workforce (of a specific skill-level). This is a bit different than the usual flow equation in utility posting models. The term pL_t is the expected number of current workers retained, and $\lambda p(N - L_t)$ is the expected number of offers made to, and accepted by, outside workers. The firm effectively faces two different upward-sloping labor supply curves in each period: one from its current workers, and a second from outside workers.

Imposing the steady-state condition $L_{t+1} = L_t \equiv L$ and substituting Equation (6) for p, steady-state employment as a function of the utility offer is:

$$H(\bar{u}) = \frac{\lambda K \exp(\bar{u}) N}{\Omega(\bar{u})} \tag{7}$$

where $\Omega(\bar{u}) \equiv [1 - (1 - \lambda) K \exp(\bar{u})]$ measures the increase in steady-state employment arising from the firm's advantage in making offers to its current employees. When $\lambda < 1$, the incumbent advantage is larger for jobs with particularly attractive exogenous characteristics, reducing the marginal cost of recruiting. These firms choose to grow larger, and increase the utility offer to do so. We discuss these implications in more detail below. When $\lambda = 1$, there is no incumbent-firm advantage because all workers get offers from all employers in every period. In this case, Equation (7) simplifies to the static labor supply equation in Card et al. (2016), but extended to include the endogenous amenity R. This is a useful benchmark case to which we return below.

3.2 The Firm's Choice of Wages and Risk

For each occupation, and each type of labor, the firm chooses an offer bundle (w, R) to maximize steady-state profits:

$$\max_{w,R} \left[Q - C(w, R)\right] H(\bar{u}). \tag{8}$$

where $\bar{u}$ depends on w and R as described above. The first-order condition with respect to w implies:

$$\left[Q - C(w, R)\right] \frac{\partial H(\bar{u})}{\partial w} = C_w(w, R) H(\bar{u}). \tag{9}$$

By the chain rule, the partial derivative of steady-state employment with respect to w is

$$\frac{\partial H}{\partial w} = f_w(w, R)\left(\frac{H(\bar{u})}{\Omega(\bar{u})}\right). \tag{10}$$

The first-order condition with respect to R is analogous.

Taking the ratio of the necessary first-order conditions with respect to w and R yields:

$$\frac{f_w(w, R)}{f_R(w, R)} = \frac{C_w(w, R)}{C_R(w, R)}. \tag{11}$$

Hence, as in the classical, static and frictionless hedonic wage model, the firm's optimal offer of wages and risk equates worker willingness-to-pay for safety with the marginal cost to the firm of providing it.

3.3 Equilibrium Wages and Compensating Wage Differentials

We now add assumptions about worker preferences and firm unit labor costs. Following Hwang et al. (1998), we assume indirect utility over wages and risk is given by $f(w, R) = \ln w - h(R)$, and that the logarithm of unit labor costs are given by $\ln C(w, R) = w - y_{bk}(R)$.[5] Note the bk subscript on y, which highlights that jobs differ in the marginal cost of providing safety. Finally, the revenue generated by a unit of type s labor when employed by firm j in occupation k is $Q_{sjk} = T_j \theta_s \pi_k$.

These assumptions yield several implications. First, from Equation (11), we find, just as in Hwang et al. (1998), $y'_{bk}(R) = h'(R)$. As a result, all firms in the same industry optimally choose the same level of risk for each occupation. This matches our empirical setting, since we measure fatality risk in detailed industry-occupation cells. Second, after making the relevant substitutions into the necessary first-order conditions, and taking logarithms, profit-maximizing equilibrium log wages are given by:

$$\ln w = \ln T_j + \ln \theta_s + \ln \pi_k + y_{bk}(R) + \ln\left(\frac{1}{1 + \Omega(\bar{u})}\right). \tag{12}$$

Differentiating Equation (12) with respect to R and substituting the equilibrium condition $y'(R_{jk}) = h'(R_{jk})$ gives a structural equation for the unconditional relationship between wages and risk:

$$\frac{d \ln w}{dR} = h'(R)\left[1 - \left(\frac{1 - \Omega(\bar{u})}{1 + \Omega(\bar{u})}\right)\right]. \tag{13}$$

[5] This specification of preferences also generalizes Card et al. (2016). They introduce a preference parameter on the log wage which is, in turn, a measure of the elasticity of labor supply with respect to the wage. We eliminate this parameter to simplify exposition. Adding it changes none of the implications of our model.

This implies the profit-maximizing relationship between wages and fatality risk is an attenuation of worker preferences for safety. The degree of attenuation varies based on the incumbency advantage $\Omega(\bar{u})$.

3.4 Connection to Empirical Model

To help guide and interpret our empirical analysis, we return to the benchmark empirical OME model from Equation (4) to connect the parameters to their theoretical analogs. The theoretical model yields an additively separable structural log wage equation as a profit-maximizing equilibrium result. Of primary interest is the connection between estimated parameter on risk, $\widehat{\gamma}$, from the OME model and the theoretical model. Equation (12) implies the OME model is mis-specified since it omits the incumbency advantage term, and Equation (13) suggests that the estimate from the linear conditional expectation function will be a weighted average of downward-biased estimates of worker preferences for safety.

It is instructive to consider the case where all workers receive offers from all jobs in every period: that is, where special $\lambda = 1$. In this case, there is no incumbency advantage, and the unmodeled utility component in Equation (12) is constant across jobs. The OME model is identified in this case, and $\gamma = \frac{d \ln w}{dR} = h'(R)$, an unbiased estimate of preferences. Not coincidentally, when $\lambda = 1$ worker mobility is exogenous with respect to wages. Interestingly, this result shows that the classic Rosen (1974) model can be relaxed to allow for firms to have power in the labor market.

In the more general case, $\lambda < 1$ and incumbent employers have a hiring advantage. Jobs offering greater utility ($\bar{u}$) have less need to hire on the outside market, and since labor supply curves are upward sloping this decreases the effective marginal cost of labor. This means, first, that risk affects wages through a direct channel, via its effect on the cost of labor, and indirectly, through its influence on recruiting. Second, exogenous amenities a_j enter the wage equation indirectly through Ω. Without making additional assumptions on the nature of preferences, the non-random assignment of workers to jobs on the basis of utility is reflected in wages. In the empirical model, this means that wage residuals will be correlated with risk, even after conditioning on firm and worker effects.

As a result, the compensating wage differential $\frac{\partial \ln w}{\partial R}$ still represents a downward-biased estimate of $h'(R)$, where the magnitude of the bias scales with the portion of the incumbent recruiting advantage that is not controlled by covariates. Just as in Hwang et al. (1998), the reduced-form compensating differential is a mixture of components of the underlying model. However, whereas they conclude that the resulting bias is extreme when considering firm-level amenities, our model is more optimistic, implying identification is still possible when amenities vary across jobs in the same firm.

Specifically, under the linearity assumption (which we relax Section 6) $\widehat{\gamma} = \frac{\partial \mathbb{E}[\ln w|x,\theta,\Psi]}{\partial R}$ is the estimated treatment effect on log wages of changing fatality rates, holding fixed covariates. Equation 13, in contrast, is the total derivative of log wages with respect to risk, which is equivalent to the estimand of a univariate OLS regression of log wages on risk without any controls, and is conceptually distinct from a 'compensating wage differential'. If $\Omega(\bar{u})$ were observable, the structural equation could be estimated directly to recover the partial derivative $\frac{\partial \ln w}{\partial R} = h'(R)$. However, since $\Omega(\bar{u})$ is unobserved our estimated γ potentially differs from $h'(R)$. From Equation 4,

$$\frac{\partial \mathbb{E}\left[\ln w|x,\theta,\Psi\right]}{\partial R} = \widehat{\gamma} + \mathbb{E}\left[\frac{\partial \ln w}{\partial \xi}\frac{\partial \xi}{\partial R}\bigg| x,\theta,\Psi\right] \tag{14}$$

γ is an unbiased estimator of the preference parameter $h'(R)$ if

$$\frac{\partial \ln w}{\partial R} = \frac{\partial \mathbb{E}\left[\ln w|x,\theta,\Psi\right]}{\partial R} \quad \Rightarrow \quad \mathbb{E}\left[\frac{\partial \ln w}{\partial \xi}\frac{\partial \xi}{\partial R}\bigg| x,\theta,\Psi\right] = 0 \tag{15}$$

which holds if the included variables in the OME model control for $\ln\left(\frac{1}{1+\Omega(\bar{u})}\right)$. In practice, the relationship between γ and $h'(R)$ therefore depends on whether the additive separability specification in the OME model is valid, or whether there is important residual match-level heterogeneity that affects wages.

To assess the magnitude of the resulting bias, we consider three factors. First, when firms have very small shares of the market $\Omega \approx 1$ and the bias is negligible. Second, the presence of firm and occupation effects in the wage decomposition will absorb variation arising from exogenous amenities. Hence, a bias can only arise from match-specific heterogeneity not picked up by those controls. We devote substantial attention in the empirical analyses to showing that this residual match-specific component is minimal. Third, to the extent that any large firms have non-negligible values of Ω, proxies for the worker retention probability can be used to construct a control function for the remaining portion of the structural error term. We show in Table 10 that controlling for completed tenure has no effect on the estimated compensating differential.

4 Data and Sample Descriptions

We use matched employer-employee data from Brazil's *Relação Anual de Informações Sociais* from 2003-2010. These data serve two purposes: first, as a source of information about workplace fatalities, and second as a source of information about jobs and earnings.

4.1 RAIS Data

RAIS is an annual census of all formal-sector jobs. Each year, the Brazilian Ministry of Labor and Employment (MTE) collects data on every job for the purpose of administering the *Abono Salarial* — a constitutionally mandated annual bonus equivalent to one month's earnings. The information in RAIS is provided to the MTE by a manager in each establishment. Compliance with reporting requirements is extremely high, as employers who fail to complete the survey face mandatory fines and also risk litigation from employees who have not received their *Abono Salarial*.[6] For each job, in each year, the employer reports characteristics of the worker, the job match, and the establishment. Worker characteristics include gender, race, age, and educational attainment.[7] Job characteristics relevant to this study include the wage, occupation, and whether the job ended because of a fatal injury. The establishment characteristics include the establishment's industry, location, and number of employees.[8]

4.2 Measuring Fatality Rates

When a job ends the employer reports the cause of separation, which determines any severance compensation to which the worker is entitled, from a list of 23 options, three of which cover work-related fatalities (see Appendix Table A.1). The RAIS data thus function as a census of fatal occupational injuries. We aggregate the job-level data to measure the average fatality rate in each of 11,440 two-digit industry by three-digit occupation cells as the number of fatal injuries per 1,000 full-time full-year-equivalent workers. This follows the method of reporting fatal injury rates used by the Bureau of Labor Statistics since 2007.[9] See Appendix B.2 for details of this calculation. We also pool fatality data in three-year windows. For example, the measured fatality rate for an industry-occupation cell in 2005 is constructed using fatality counts and hours across all jobs in that cell from 2003, 2004, and

[6] For details on labor market formality and wage setting institutions, see Appendix B.3

[7] Because individual characteristics are reported by the employer, they can change as workers move from job to job. Cornwell, Rivera and Schmutte (forthcoming) provide evidence that discrepancies in employers' reports of worker characteristics are associated with other unobserved determinants of earnings, so we leave these variables in as reported.

[8] That industry and occupation are reported by the employer is an advantage of RAIS. In many major U.S. surveys, occupations are measured with error, and are not consistently coded over time. Inconsistent measurement of occupation can badly bias panel data models, as has been illustrated using the CPS by Moscarini and Thomsson (2007), in the PSID by Kambourov and Manovskii (2008), and in the NLSY by Speer (2016). By contrast, Abraham and Spletzer (2010) find that businesses tend to report occupation more accurately and more consistently.

[9] See http://www.bls.gov/iif/oshnotice10.htm for a description of how and why the BLS constructs hours-based fatality rates. One relative advantage of our data is that we observe both the number of months a job lasted as well as the number of contracted weekly hours. By contrast, the BLS fatality rates are scaled by average hours at work from the CPS.

2005. We do so for comparability with previous literature, and to smooth out fluctuations in the annual fatality rates (Kniesner et al. 2012).

Our ability to disaggregate the data by detailed industry-occupation cells yields substantially more variation in fatality risk than has been available in previous studies. Lalive (2003) and Tsai, Liu and Hammitt (2011) find that coarse measures of fatality risk can produce substantial aggregation bias in estimates of compensating wage differentials. We therefore err on the side of using as disaggregate a measure as possible. Since fatal accidents are rare events, one concern is that the decreased bias from this disaggregation entails a large increase in variance of estimated cell-specific fatality rates. We address this trade-off by restricting our sample to cells with at least 10,000 full-time full-year-equivalent workers. We leave remaining measurement issues for future research.

In Table A.2, we report average fatality rates by aggregate industry and occupation as evidence that our measurements of fatality risk are sound. The overall fatality rate (including both men and women) is 0.049 fatalities per 1,000 full-time full-year equivalent workers. By comparison, the fatality rate in the U.S. was about 0.037 per 1,000 full-time full-year-equivalent workers over the same time period. In our data, fatality rates are highest in the Agriculture and Fishing, Mining, Construction, and Transportation industries. Among occupations, the fatality rate is highest among Production and Manufacturing I workers, and lowest among Professionals, Artists, and Scientists.

4.3 Analysis Sample and Variable Definitions

To prepare our data, we first define a population of interest, and then construct an analysis sample that we use throughout the empirical work. The population of interest consists of jobs held by male workers between the ages of 23 and 65. Like Abowd et al. (1999), Woodcock (2008), and Card et al. (2013), we restrict our sample to a single 'dominant' job for every worker in every year. For each worker, their dominant job in any year is the one with the highest expected earnings.[10] We further restrict analysis to jobs with at least 30 contracted hours per week in establishments with at least two workers. We also exclude government jobs and temporary jobs. As described in Section 4.2, we consider only jobs in 2-digit industry by 3-digit occupation cells that contain at least 10,000 full-time full-year-equivalent workers. Finally, we Winsorize the data at the 1st and 99th percentiles of the log wage distribution. After imposing these restrictions we have an analysis sample with about 83 million job-years.

The RAIS data report average monthly earnings. If the worker is in the job for less than 12 months during the year, the variable reported by RAIS represents one month's pay.

[10]We define expected earnings as the product of the average monthly wage rate with the number of months the worker was employed.

Table 1: Summary Statistics

	Population	Analysis Sample
Age	36.98	36.23
Race *branco* (White)	0.56	0.58
Elementary or Less	0.40	0.40
Some High School	0.09	0.10
High School	0.36	0.39
Some College	0.04	0.04
College or More	0.11	0.07
Contracted Weekly Hours	42.19	43.34
Hourly Wage	6.10	5.10
Log Hourly Wage	1.47	1.37
Total Experience (Years)	20.58	19.86
Job Tenure (Months)	58.70	44.28
Fatality Rate (per 1,000)	0.071	0.083
Zero Fatality Rate (Percent)	0.14	0.09
Number of Observations	158,254,802	83,418,032

Notes: The population includes all dominant jobs held by men between ages 23 and 65. 'Analysis Sample' restricts to jobs 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, and with hourly earnings between the 1st and 99th percentiles of the Analysis Sample earnings distribution.

In practice, this variable measures the monthly wage rate, which is a common institutional arrangement in Brazil. For consistency with prior research, we convert to an hourly wage rate measured in 2003 Brazilian Reais.[11]

For each job, the data report the date of hire. Hence, even for the first in-sample job-year, we have an accurate measure of tenure on that job. Using tenure we impute labor market experience as the maximum of tenure in the first observed job or potential experience, whichever is largest, plus observed accumulated experience from jobs held during the years in which we have data.

Table 1 reports descriptive statistics for the male population and analysis sample. Relative to the population, observations in the analysis sample include workers that are slightly younger, less educated, less experienced, and in riskier jobs. This is due primarily to selection on jobs with more than 10,000 full-time full-year-equivalent workers. The average monthly wage in the analysis sample is 682 Reais, and the average fatality rate is 0.083 deaths per

[11] First we calculate a weekly wage rate as the monthly wage rate divided by 4.17. We then calculate the hourly wage rate as the weekly wage rate divided by the contracted weekly hours, which are also reported for every job. Conveniently, one Brazilian Real in 2003 is worth approximately 1.5 Brazilian Reais in 2010. Likewise, in 2010, one U.S. dollar was worth 1.66 Brazilian Reais. Hence, one can loosely interpret our results in 2010 dollars.

1,000 full-time full-year workers. Finally, 9 percent of sample observations are associated with jobs that have a measured fatality rate of zero.

5 Results

We begin our empirical analysis by showing that, consistent with our predictions, the estimate from the within-worker model is biased down relative to the estimates from the pooled and the OME models. A residual diagnostic shows the within-worker model is misspecified, and that the OME model largely resolves this misspecification. We then provide supporting evidence that the OME model is not contaminated by omitted match effects. The section ends with a full discussion of the variation that identifies the effect of risk in the OME model, and study the relationship between fatality risk and omitted worker and establishment wage effects. As expected, high-wage workers are employed in much less risky jobs, and high-wage establishments offer safer employment.

5.1 Baseline Results

Table 2 compares estimates of the compensating wage differentials and implied VSLs using each of the empirical models from Section 2. Column (1) includes the estimates from the pooled cross-sectional model. The estimated compensating wage differential, $\hat{\gamma} = 0.279$, suggests that an increase in the average fatality rate of one death per 1,000 full-time equivalent worker-years is associated with an approximately 28 percent increase in the wage. In estimating this model, our control variables include dummies for each year (up to 30) of labor market experience, worker race, education, plant size, year, state of employment, 1-digit industry and 1-digit occupation effects.[12] Re-scaling this coefficient implies an estimated value of statistical life (VSL) of 2.84 million Brazilian Reais (in 2003 Reais) with 95 percent confidence interval $[2.83, 2.86]$.[13]

[12] Our set of control variables is fairly standard, and we maintain the same control set in all subsequent models, with adjustments as needed to account for collinearity between worker-, firm-, and match-specific characteristics in our panel data models. We control using dummies for each year of experience (up to a maximum of 30 years of experience) for two reasons. First, the magnitude of our data facilitate a flexible specification of experience profiles. Second, as Card et al. (2016) illustrate, person effects are not identified relative to year and experience effects without some normalization. Our experience profile is such a normalization, albeit a specific one. As we show in Table A.5, our main results are not sensitive to alternative specifications of the experience profile.

[13] Following the original treatment in Rosen (1974) and subsequent literature, we calculate the VSL as: $VSL = \frac{\partial w}{\partial a} * 1000 * 2000$. Since wages are measured in Reais per hour, while the fatality rate is measured in deaths per 1,000 full-time equivalent worker years, the derivative is scaled by 1,000 FTE worker-years at 2,000 hours worked per FTE. Because of our log-linear specification, $\frac{1}{w}\frac{\partial w}{\partial a} = \widehat{\gamma}$, so $VSL = \bar{w}\widehat{\gamma} * 2,000,000$. As in Kniesner et al. (2012), we evaluate this VSL function and associated confidence interval at the population mean wage, treating $\bar{w}$ as a constant.

Table 2: Compensating Wage Differentials for Full-Time Prime-Age Men

	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) OME
Fatality Rate (3-Yr MA)	0.279*	0.037*	–0.006*	0.170*
	(0.001)	(0.001)	(0.001)	(0.001)
Zero Fatality Rate	0.073*	0.008*	–0.006*	0.014*
	(0.000)	(0.000)	(0.000)	(0.000)
N	83,411,371	83,418,032	83,418,032	83,418,032
R-Sq	0.458	0.913	0.978	0.930
VSL (millions of reais)	2.84	0.37	-0.06	1.73
95% CI	[2.83, 2.86]	[0.35, 0.39]	[-0.09, -0.03]	[1.72, 1.75]

Notes: Model 1 also includes 1-digit industry effects, 1-digit occupation effects, year effects, state effects, race effects, years of experience effects (censored at 30), indicators for small and medium-sized establishments, and education effects. Model 2 includes worker effects and the same controls as Model 1 except for race and education. Model 3 includes job-match effects, years of experience effects, and year effects. Model 4 includes worker effects, establishment effects, 1-digit occupation effects, and year effects. The analysis sample includes dominant jobs of men 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. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage, and reported in millions of reais. * Indicates significance at the 0.01 level.

The coefficient on the indicator for zero fatality rates is also notable in this specification, implying that workers employed in the very safest jobs are paid roughly seven percent higher wages than otherwise equivalent workers. Failure to account for this discontinuity in the wage-risk profile substantially attenuates the estimated compensating wage differential in this pooled specification.

Column (2) presents estimates from the worker effects model, Equation 1. Relative to the pooled model in column (1), the estimated compensating wage differential falls by about 87% to 0.037, comparable to patterns in the literature using US data (see in particular Brown (1980), Kniesner et al. (2012), and Lavetti (2015)). The estimated effect of being employed at a job with zero fatalities also declines an order of magnitude, to 0.008.

Finally, estimates from each of the two steps of our preferred orthogonal match effects model are presented in columns (3) and (4). Column (3) reports results from estimating Equation (3), the match effects specification. In this model γ is identified only from the time-series variation in fatality risk within jobs, the nature of which is affected by our decision, following Kniesner et al. (2012), to construct measured fatality rates using a three-year moving average. We find that wages do not vary in response to changes in fatality risk that occur within the same match. Specifically, the estimated compensating wage differential using within-match variation is -0.006, which is not economically meaningful.

Column (4) presents our preferred estimates from Equation (4), which controls for all unobserved worker- and establishment-specific effects on wages. The estimated compensating wage differential, $\widehat{\gamma} = 0.170$ (SE 0.001), is 350 percent higher than the estimate from the worker effects model in column (2). Rather than attenuating the estimated wage differential further, controlling for unobserved establishment heterogeneity restores the differential to a level between the pooled and within-worker estimates. The pattern is exactly consistent with the type of bias that is predicted to arise theoretically in a model of hedonic search.[14]

One might be concerned that our preferred estimates from the OME model are sensitive to features of the model specification. In Table A.5 and Table A.6 we show that our results are extremely robust to a wide range of alternative specifications for the time-varying controls, whether estimated by OME or AKM. Finally, regarding inference, Table A.7 reports alternative standard error estimates allowing for correlation in errors either within establishments or within occupation-industry cells. The OME estimates remain statistically significant at the 0.01 level under all of the alternative clustering specifications.

[14] Table A.3 shows these results are robust to dropping industry and occupation controls. Table A.4 shows the same bias pattern holds when we estimate the model separately by region.

5.2 Model Diagnostics and Endogenous Mobility

The results in Table 2 suggest that the worker effects model is biased due to the omission of job-specific characteristics that affect mobility and wages. In this section, we present residual diagnostic evidence that the worker effects model is mis-specified. We then report model diagnostics introduced by Card et al. (2013) to show that most of the remaining variation in log wages is well-explained by the inclusion of establishment effects.

5.2.1 Are Changes in Fatality Risk Correlated with Changes in the Residual?

If the worker effects model in Equation (1) were properly specified, then for workers who change jobs the expected change in the wage residual should be zero conditional on the change in risk. Figure 2a shows this is clearly not true. The figure displays a binned scatterplot of the average change in residuals against the change in fatality risk using observations corresponding to a job-to-job change.

First, the average change in the residual is positive. This is consistent with the 'job ladder' inherent in search behavior observed across many studies, including Schmutte (2015). Furthermore, Figure 2a also suggests that workers with large decreases in risk experience larger increases in the wage residual, which is consistent with the form of endogenous mobility suggested by a model in which job changes involve movements to jobs that are more attractive on both wage and safety dimensions.

The data thus clearly reject the exogenous mobility assumption required for the worker effects model to be identified. A key question is whether controlling for establishment heterogeneity is sufficient to eliminate endogenous mobility bias. Figure 2b presents the analogue of Figure 2a using residuals from the OME model. There are two important results. First, the average change in residuals is centered around to zero. Second, over most of the domain there is not a strong systematic relationship between the change in risk and average change in residuals. This suggests limited scope for potential endogenous mobility bias in the OME model.

5.2.2 Do Match Effects Matter?

Our theoretical model implies wages depend on a term that is a product of individual and employer-specific unobservables. However, it is not clear whether or how badly omitting this term will bias estimates of the compensating wage differential, which depends on how much of the unobserved source of variation is absorbed by worker and employer effects. We apply the diagnostic tools developed by Card et al. (2013) to evaluate the separability and exogenous mobility assumptions of the OME model. Figure 3 displays the mean residual within cells defined by deciles of the estimated worker and establishment effects from our

Figure 2: Binned Scatterplot of Average Change in Residual by Change in Fatality Risk for Job Changers

(a) Worker Effects Model

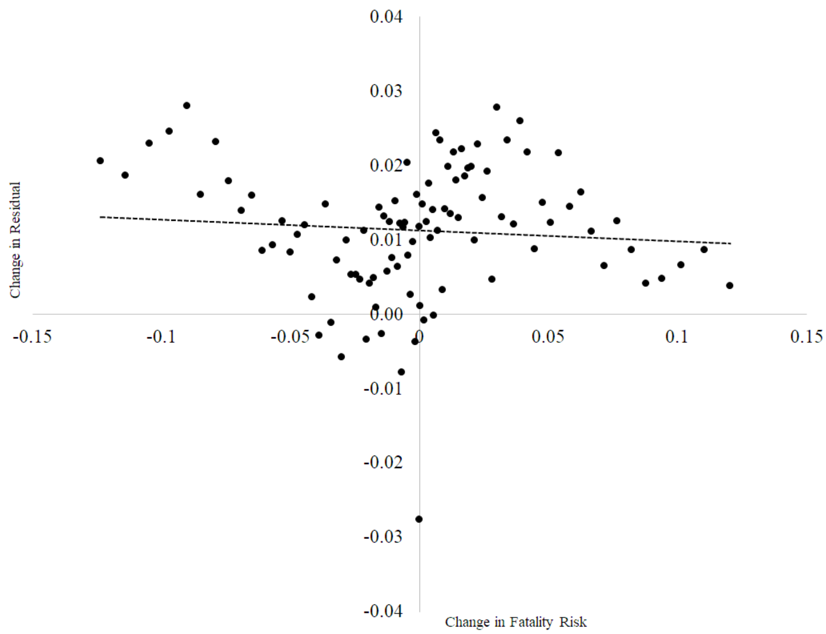

(b) Orthogonal Match Effects Model

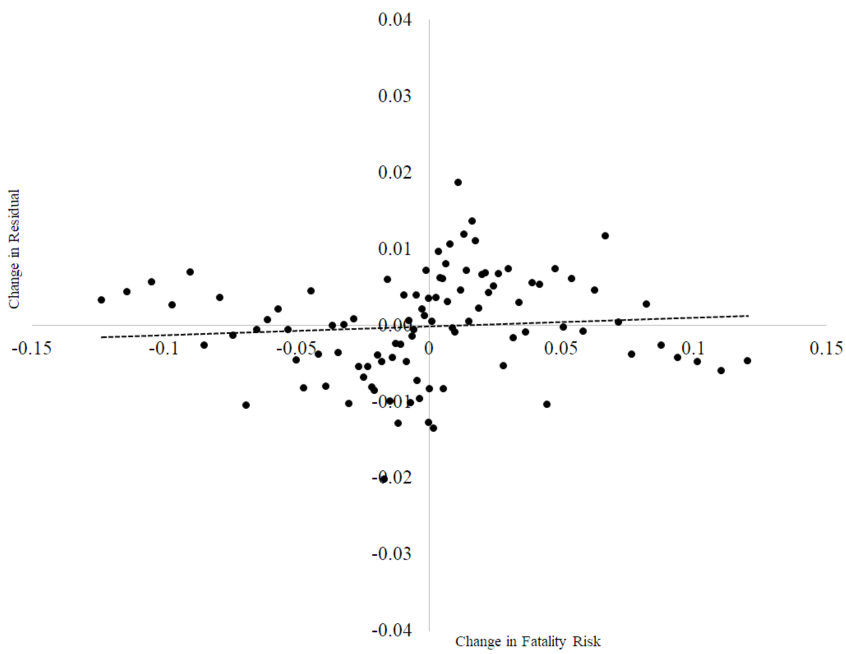

Notes: Figures plot of the average change in residuals for workers who change jobs year-over-year within each percentile of the distribution of change in the fatality rate. The residuals are from the worker effects and orthogonal match effects models, respectively. Fatality rates are measured in deaths per 1,000 full-time full-year equivalent workers.

Figure 3: Mean Residuals by Decile of Establishment/Person Effect, 2005–2010

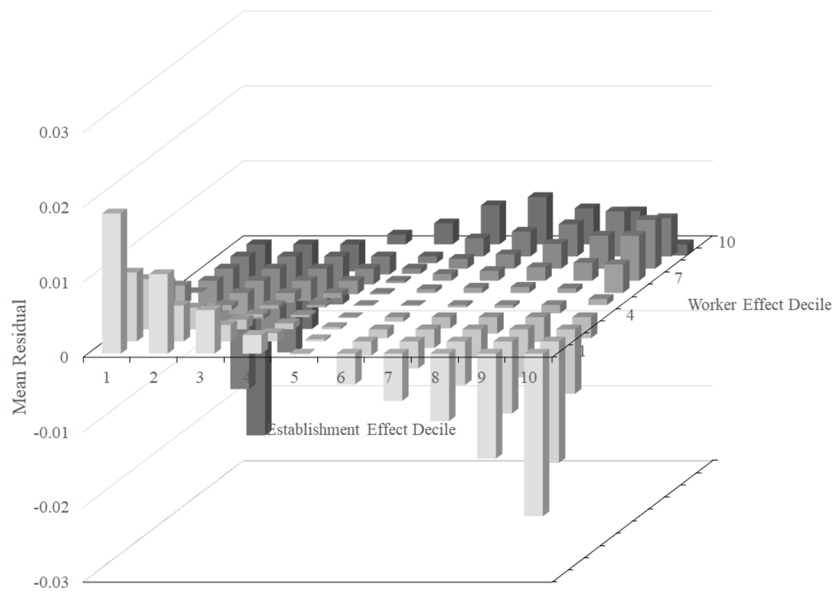

Notes: Figure displays the mean residual from the OME model within cells defined by the estimated establishment effect interacted with the decile of estimated worker effect.

benchmark model. Except for the lowest-paid workers, and for workers employed in the lowest-paying establishments, these errors are all less than 0.01 in magnitude. This suggests the separability assumption is a good approximation to the true data generating process, except perhaps at the bottom of the wage distribution, where minimum wages and other institutional constraints on pay are more likely to bind. We return to discussing this figure in Section 6.3, in which we conduct sensitivity analyses to restricting identifying variation to observations in the middle of the distributions of worker and establishment effects, where the residuals are closest to zero.

If match effects play an important role in job assignment, one would expect to observe that when workers switch jobs into establishments that pay lower average wages (lower Ψ), many workers would still experience wage increases due to improvements in match quality. This implies that the importance of match effects can be assessed by estimating whether wage gains associated with transitioning into higher Ψ establishments are asymmetric to wage losses associated with a transition of the same magnitude down the Ψ distribution.

Table 3 reports the average wage change associated with a move from each decile of the establishment wage effects distribution to each other decile. The average wage changes are highly symmetric—a move from the fifth decile to the first decile, for example, is associated

Table 3: Mean Wage Change of Movers by Decile of Origin and Destination Establishment Effect, 2005–2010

		Destination Estalishment Effect Decile									
		1	2	3	4	5	6	7	8	9	10
	1	-0.001	0.123	0.230	0.319	0.406	0.489	0.580	0.705	0.867	1.190
	2	-0.123	0.000	0.075	0.150	0.224	0.300	0.383	0.483	0.621	0.909
	3	-0.233	-0.074	-0.001	0.062	0.136	0.210	0.291	0.390	0.525	0.793
	4	-0.320	-0.150	-0.063	0.000	0.063	0.132	0.207	0.303	0.436	0.701
Origin	5	-0.403	-0.226	-0.135	-0.061	0.000	0.062	0.137	0.235	0.367	0.623
Decile	6	-0.491	-0.300	-0.206	-0.131	-0.064	0.005	0.066	0.160	0.287	0.543
	7	-0.589	-0.382	-0.288	-0.212	-0.141	-0.067	0.000	0.082	0.203	0.457
	8	-0.706	-0.483	-0.387	-0.305	-0.238	-0.158	-0.078	-0.001	0.110	0.352
	9	-0.864	-0.623	-0.522	-0.437	-0.366	-0.284	-0.200	-0.108	0.001	0.193
	10	-1.192	-0.906	-0.790	-0.705	-0.624	-0.548	-0.454	-0.356	-0.189	-0.002

Notes: Table entries are mean differences between wages on the origin and destination job for workers who change jobs. Each job is classified into deciles based on the estimated establishment effect from the OME Model, Equation 4.

with a 40.3 percent reduction in wages, while a move in the opposite direction from the first to the fifth decile is associated with an increase in wages of 40.6 percent. This close symmetry holds for every origin-destination decile pair—there are no pairs with an asymmetry greater than 0.005. Second, job transitions within any decile of the distribution (along the diagonal of the table) are associated with no average change in wages. This suggests that on average when workers change jobs they only experience wage increases if the destination job has a higher Ψ, leaving very little role for the potential influence of improvements in match quality or interaction effects between establishment-occupation pairs.

5.3 Sources of Identifying Variation

In this section we present a decomposition of the components of variation in wages and fatality rates, and show that unobserved worker effects and firm effects are negatively correlated with fatality risk. We then present evidence relaxing the assumptions in our benchmark empirical model that the effect of risk on wages is the same for moves within or between establishments, industries, or occupations. Finally, we aggregate wages to the industry-occupation level, the level at which fatality risk varies, to graphically depict the basic source of identifying variation and the nature of our bias correction.

Table 4: Variance Decomposition of OME Model Components

	Component
Std. Dev. of Log Wage w_{it}	0.650
Std. Dev. of P_{it}	0.648
Std. Dev. of θ_i (Worker Effect)	0.456
Std. Dev. of $\Psi_{J(i,t)}$ (Estab. Effect)	0.298
Std. Dev. of $\gamma R_{c(i,t)}$	0.014
Std. Dev. of Residual	0.172
Correlation between $(\theta_i, \Psi_{J(i,t)})$	0.280
Correlation between $(R_{c(i,t)}, \theta_i)$	−0.091
Correlation between $(R_{c(i,t)}, \Psi_{J(i,t)})$	−0.108
Std. Dev. of $\Phi_{i,J(i,t)}$ (Match Effect)	0.133

Notes: Variance components estimated from the orthogonal match effects model described in Equations 3 and 4. Standard deviation of match effects is estimated by the square root of the difference between the AKM mean squared error and the mean squared error from Equation 3.

5.3.1 Variance Decomposition

In Table 4, we report the contributions to total wage variance of the components of the OME wage decomposition. The estimated worker and establishment effects explain 70% and 46%, respectively, of the variation in the dependent variable in Equation 4. The standard deviation of match effects is 45 percent of the standard deviation of establishment effects.[15] We also find that fatality risk is strongly negatively correlated with estimated worker and establishment effects, with estimated correlation coefficients of –0.091 and –0.108, respectively. These correlations support the central intuition underlying both ability and endogenous mobility bias: higher paid workers tend to be employed in less risky jobs, and higher paying establishments also offer jobs with lower fatality risk.

The correlations between fatality rates, worker effects, and establishment effects are also visually apparent in Figure 4. Panel (a) presents a binned scatterplot, in which each point along the horizontal axis represents a percentile in the distribution of fatality rates, and the vertical axis reports the average estimated worker effect θ_i in the bin containing all jobs in the corresponding percentile. The figure shows that high wage workers tend to be employed at less risky jobs, consistent with the long-hypothesized view idea that workers with higher earnings potential choose to consume lower risk. Panel (b) is the equivalent figure for establishment effects, $\Psi_{J(i,t)}$, and demonstrates that in firms where the average fatality rate is very low, estimated firm effects are above average. This is consistent with

[15] Card et al. (2013) find this share is somewhat smaller, 35 percent, in data from West Germany.

Figure 4: Binned Scatterplots of Worker and Establishment Effects versus Fatality Rates

(a) Worker Effects (b) Establishment Effects

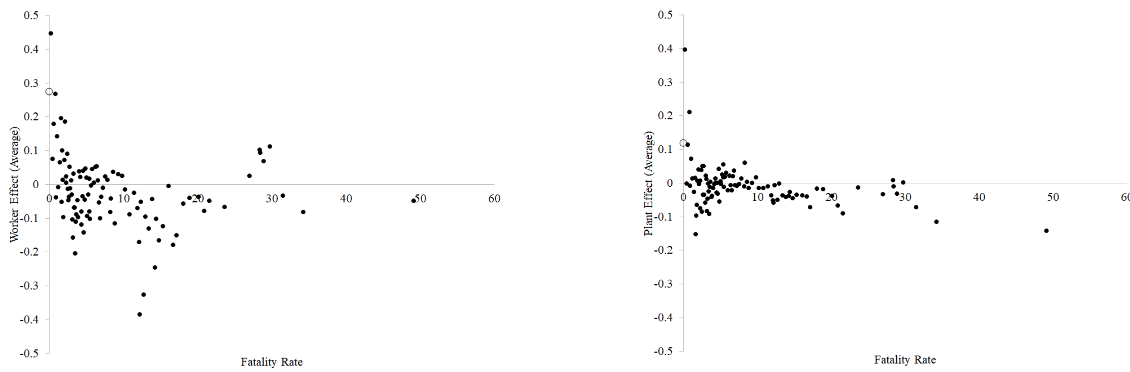

Notes: The figures plot the average worker and establishment effects estimated from the model in Equation 3 at each percentile of the distribution of the fatality rate. Fatality rates are measured in deaths per 100,000 full-time full-year equivalent workers.

the implication of the job search model that firms improve wages and job amenities when trying to attract or retain workers. For fatality rates away from zero, the distribution of establishment effects is also quite flat, suggesting that the choice of establishment wage effects is not strongly related to the level of risk.[16] This pattern would be unlikely to arise if establishments paid occupation-specific wage premia, since most of the variation in fatality rates occurs across occupations.

In our analysis sample, over 97% of the total variance in log wages occurs across jobs. Of this variation, 95% can be explained by a two-way fixed effects model with worker and establishment effects alone. These facts suggest that any residual unexplained wage variation is extremely small. A decomposition of the estimated establishment effects reveals that 17% of the variation can be explained by variation within establishments across 3-digit occupations. However, a two-way fixed effects model with worker effects and establishment-by-3-digit occupation effects explains less than 2% more of the variation in wages relative to a two-way model with only worker and establishment effects. These patterns suggest that, although there is variation in wages across occupations within establishments, the variation looks very different than a systematic wage premium, consistent with the patterns in Table 3 and the consensus interpretation of establishment wage effects in the literature.

[16] Doubling the mean fatality rate is associated with about a 0.05 standard deviation decrease in Ψ.

Table 5: Sensitivity of γ to Type of Job Change

	(1)	(2)
Fatality Rate	0.157*	0.157*
	(0.001)	(0.001)
Fatality Rate*Change Occupation	0.007*	0.001
	(0.001)	(0.001)
Fatality Rate*Change Establishment	0.009*	–0.014*
	(0.001)	(0.001)
Fatality Rate*Change Industry		0.041*
		(0.001)
N	83,418,032	83,418,032
R-Sq	0.930	0.930

Notes: Estimates are from the OME specification. 'Change Occupation' and 'Change Industry' equal 1 for all years of an origin and destination job that differ in 3-digit occupation code and 2-digit industry code, respectively. 'Change Establishment' equals 1 if the job change is across establishments, and equals 0 if the worker changes occupation within an establishment. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. * Indicates significance at the 0.01 level.

5.3.2 Variation in the Nature of Job Change

There is substantial variation in fatality rates across occupations within establishments.[17] Of the total variation in risk across jobs, 69% occurs across 3-digit occupations, 33% occurs across 2-digit industries, and 77% occurs across either 3-digit occupation or across 2-digit industry. Since the OME specification includes controls for one-digit occupation effects, the identifying variation in fatality rates is across 3-digit occupations conditional on establishment and 1-digit occupation effects, which is 33% of the total across-job variation in fatality rates.

Table 5 presents results from a model in which we allow γ to vary for different types of job changes. The estimates show that conditional on the OME controls, the compensating differential per unit of fatal risk is approximately the same for job changes within or across occupations, and for job changes within or across establishments. γ is slightly higher for within occupation and across industry job changes (0.184) than it is for within occupation and industry job changes (0.165). Nevertheless, the results are broadly consistent with our

[17] Of the total variation in fatality rates, only about 3% occurs within job matches. The primary source of this variation is a general downward trend in fatality rates throughout Brazil between 2003–2010. If job search is imperfect, one may not expect these decreases in fatality risk to be fully reflected in wage changes during the match. It is also possible that such small movements in fatality rates within jobs are not salient to workers. For these reasons, we do not rely upon this variation as a primary source of identification. Indeed, our estimate of the CWD using within-match variation is effectively zero, consistent with a job search model without renegotiation.

model, which assumes that after conditioning on worker and establishment heterogeneity the compensating wage differential does not depend on the type of job change.

5.3.3 Analysis of Industry-Occupation Aggregates

Figure 5a presents a more transparent illustration of the variation identifying the relationship between fatality risk and wages. We first fit the AKM model, but with establishment-occupation effects rather than just establishment effects. In our data, fatality risk is measured in industry-occupation cells, and we identify the compensating wage differential from movements of workers across industry-occupation cells with differing levels of risk. Any variation in compensation associated with variation in fatality risk across different jobs should load onto the estimated establishment-occupation effects, denoted $\hat{\psi}_{J(i,t),k(i,t)}$.

We construct an industry-occupation level dataset whose entries, $(\bar{R}_{k,n}, \bar{\hat{\psi}}_{k,n})$, are the average risk and average establishment-occupation effect of all jobs in a given occupation-industry pair, where k indexes occupations and n indexes industries. Figure 5a presents a binned scatterplot of all pairwise differences of average fatality risk:

$$\left(\bar{R}_{k,n} - \bar{R}_{k',n}\right) - \left(\bar{R}_{k,n'} - \bar{R}_{k',n'}\right) \tag{16}$$

against the corresponding difference-in-differences for industry average plant-occupation wage effects:

$$\left(\bar{\hat{\psi}}_{k,n} - \bar{\hat{\psi}}_{k',n}\right) - \left(\bar{\hat{\psi}}_{k,n'} - \bar{\hat{\psi}}_{k',n'}\right). \tag{17}$$

The quantity in (16) is the excess change in fatality risk associated with moving to a job in occupation k from a job in occupation k' when that job is offered in industry n rather than industry n'. The quantity in (17) measures the change in job-specific compensation associated with moving to a job in occupation k from a job in occupation k' when that job is offered in industry n rather than industry n'. The latter measure is free of individual characteristics by construction, and is also purged of establishment-specific components of compensation by comparing jobs offered in the same establishment, but different occupations.

Figure 5a shows that the residual variation in compensation is strongly associated with the residual variation in risk across these jobs. This is a graphical display of the variation that our OME model exploits to identify the compensating wage differential. It is instructive to compare Figure 5a with its first-differenced counterpart. Figure 5b is a binned scatterplot of the difference in average establishment-occupation effects across occupations $\left(\bar{\hat{\psi}}_{k,n} - \bar{\hat{\psi}}_{k',n}\right)$ for each industry against the difference in risk across occupations. In contrast to the difference-in-difference plot, Figure 5b shows that, within industries, jobs in occupations with higher average wages are also less risky. This highlights the central identification

Figure 5: Identifying variation: The Relationship Between Wages and Fatality Risk

(a) Difference-in-Difference (b) First-Difference across Industries

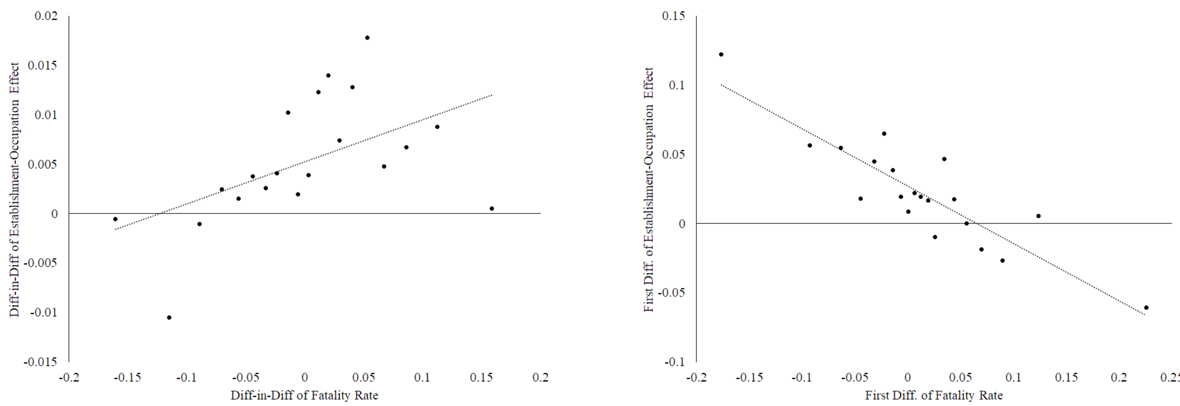

Notes: Figure 5a shows a binned scatterplot of the average difference-in-difference of establishment-specific occupation effects described in (17) against the difference-in-difference of fatality risk in (16). Figure 5b shows a binned scatterplot of the average within-industry difference across occupations of establishment-specific occupation effects against the within-industry difference across occupations in fatality risk in (16).

problem in our study – jobs that offer more desirable working conditions also offer better wages, as predicted by Hwang et al. (1998). Failure to correct for job-specific variation in total compensation will lead to an downward bias, and possibly sign reversal, in the estimated compensating wage differential for disamenities.

6 Robustness

In this section we consider the robustness of our basic results to modified specifications that isolate the influence of endogenous mobility and related mis-specification. Where appropriate, we make explicit reference to the model of Section 3 to help motivate and interpret our diagnostic tests, though we also consider forms of mis-specification beyond those envisioned by our model.

6.1 Instrumental Variables Estimates Using Realized Mobility Network

We propose an IV estimator based on the employment histories of coworkers to address any remaining endogeneity from omitted match effects. To develop the intuition behind the model, we begin with the second step of the OME specification, Equation (4): $P_{it} = \pi_{k(i,t)} + \gamma R_{c(i,t),t} + \tau_t + \theta_i + \Psi_{J(i,t)} + \xi_{it}$ Our concern is that the error may include a match effect plus a statistical residual $\xi_{it} = \mu_{i,J(i,t)} + \varepsilon_{it}$. In first differences, the second-stage model is: $\Delta P_{it} = \Delta\pi_{k(i,t)} + \gamma\Delta R_{c(i,t),t} + \Delta\tau_t + \Delta\Psi_{J(i,t)} + \left(\Delta\mu_{i,J(i,t)} + \Delta\varepsilon_{it}\right)$ where $\Delta\Psi_{J(i,t)}$ denotes the change in establishment wage effects between period $t-1$ and t. An unbiased estimate requires the exogenous mobility assumption $E(\Delta R_{c(i,t),t}\Delta\mu_{i,J(i,t)}|\Delta\Psi_{J(i,t)}) = 0$.

Our goal is to construct an instrument that is correlated with the change in accepted risk, $\Delta R_{c(i,t),t}$, but uncorrelated with the change in unobserved match effects, $\Delta\mu_{i,J(i,t)}$. We exploit the relational structure of the data to construct such an instrument as follows. First, we restrict attention to observations across pairs of years in which a worker changed dominant jobs. That is, to observations for which $J(i,t) \neq J(i,t+1)$. For each such observation in the data, indexed by (i,t), we define its 'neighbors,' denoted $N(i,t)$, to be the observations (i',τ) for $\tau \in \{t-1, t-2\}$ satisfying (i) $J(i',\tau) = J(i,t)$, (ii) $c(i',\tau) = c(i,t)$, and (iii) $J(i',\tau) \neq J(i',t)$. In words, the neighbor set contains observations from workers employed at the same establishment as worker i, who had the same occupation at that establishment, and who separated from that job in the two years preceding, $t-1$ and $t-2$.

Our proposed instrument is $\Delta\tilde{R}_{it} = \frac{1}{|N(i,t)|}\sum_{\ell\in N(i,t)}\Delta R_\ell$, the average change in risk on accepted jobs for observations in $N(i,t)$.[18] The intuition behind this instrument is that since workers in $N(i,t)$ sorted into the same job as worker i, they are likely to have similar preferences, skills, and outside opportunities. Therefore, the characteristics of their destination jobs on separation are informative of the set of outside opportunities for i. The instrument is valid as long as worker i's idiosyncratic draw from the distribution of match effects is uncorrelated with his former co-worker's subsequent change in occupational risk. This assumption holds if the residual variation in $\Delta\tilde{R}_{it}$ within plants is uncorrelated with $\Delta\mu_{i,J(i,t)}$, which requires that the expected change in match quality be zero within $N(i,t)$. The omitted match effect on accepted destination jobs reflects a predictable component, which is common across similar workers who exited the same establishment under similar circumstances, and an idiosyncratic component. The average change in risk within $N(i,t)$ is correlated with the risk accepted by the focal worker, i, but independent of the idiosyncratic component of the realized match effect.

6.1.1 Estimation Sample

We implement the IV strategy in a sample restricted to years in which workers move from one dominant job to another. For each worker, we measure the observed change in fatality rates between the origin and the destination job. We then construct instruments for each worker as the average change in fatality rates experienced by workers who departed the same origin job (establishment-occupation) in the preceding two years. The requirements for the instrument mean that the analysis is ultimately restricted to 2008–2010, with the observations that contribute to the instrument being drawn from job changes in 2006–2009. After these restrictions, the analysis sample for the IV model uses 4,599,345 workers who

[18]Note that for observation $\ell = (i',\tau) \in N(i,t)$, $\Delta R_\ell = R_{c(i',\tau),\tau} - R_{c(i',\tau-1),\tau-1}$.

Table 6: Instrumental Variable Estimates

	(1) First-Differenced	(2) Establishment Effects	(3) IV First Stage	(4) IV	(5) OME on IV Sample
ΔFatality Rate	-0.048	0.236*		0.210*	
	(0.003)	(0.000)		(0.011)	
Avg. Δ Fat. Rate			0.338*		
in $N(i.t)$			(0.001)		
Fatality Rate					0.203*
					(0.009)
N	5,653,428	5,403,738	5,403,738	5,403,738	5,403,738
VSL (million reais)	-0.39	1.94		1.72	1.68
95% CI	[-0.44, -0.35]	[1.89, 1.99]		[1.55, 1.90]	[1.53, 1.82]

Notes: The dependent variable is the change in log wages (net of observed time-varying characteristics) between the dominant job in the prior year and the new dominant job this year. All models control for experience through the first-stage match effects model. In addition, all models control for major occupation and year. Fatality rates are measured in deaths per 1,000 full-time full-year equivalent workers. VSLs calculated at mean hourly wage. * Indicates significance at the 0.01 level.

changed jobs between 2008-2010. We describe this sample in Table A.8. The sample is slightly younger, and slightly less-educated, but is otherwise similar to the formal workforce covered by RAIS.

6.1.2 IV Results

Table 6 compares the IV estimates with estimates in simple first-differences and first-differences controlling for both origin and destination establishment effects. For consistency with the earlier estimation, we fit the model in two stages. We fit the first stage of the orthogonal match effects model for the full sample, and then estimate the remaining models using the dependent variable for the second stage of the OME model. Column (1) reports a basic first-differenced estimate of the compensating wage differential of -0.048. The specification is comparable to the worker-effects model from Table 2. Column (2) adds origin and destination establishment effects. The resulting estimate of 0.236 is consistent with our benchmark finding that controlling for establishment effects eliminates attenuation bias.

The instrumental variable estimates in Column (3) and (4) control for origin and destination establishment effects, while also instrumenting for the change in fatality rates. In the first-stage model the point estimate on the instrument is 0.338 with an F-statistic of 1.5×10^5, indicating the instrument is strongly correlated with the change in risk.

The IV estimate of the compensating wage differential in Column (4), $\hat{\gamma} = 0.210$ is slightly smaller than the effect estimated in the model controlling for establishment effects, and we reject the null hypothesis that the estimates are equal. The endogenous mobility bias that is corrected by the instrument relative to the establishment effects model appears to be modestly positive. However, when we use the IV sample to estimate the orthogonal match effects model, as shown in column (5), we estimate $\hat{\gamma} = 0.203$, which is statistically indistinguishable from the IV result. The instrumental variable results thus corroborate that to the extent the exogenous mobility assumption does not hold in the orthogonal match effects model, the impact of any associated endogeneity bias is not economically meaningful.

6.2 Separations Due to Mass Displacement Events

An alternative approach to evaluating whether voluntary movements driven by improvements in match quality are likely to bias estimates from the OME model is to consider job transitions associated with mass displacement events. This strategy was initially proposed by Gibbons and Katz (1992) to eliminate bias in estimated industry wage premia that could arise if workers learn about, and sort by, ability or comparative advantage over time.

The approach can help address the problem, discussed by Solon (1988), that panel estimates of compensating wage differentials may be biased due to the self-selection of job movers. When choosing their next job, workers moving due to layoff should not be influenced by the pay or working conditions on their previous job. The mass displacement sample should have a disproportionate number of workers laid-off for reasons unrelated to their productivity. We therefore expect selection associated with time-varying worker or match effects that are not addressed in the benchmark OME model to be substantially reduced in the mass displacement sample.

Table 7 reports estimates of our hedonic wage model when we restrict the sample to job spells on either side of a job-to-job transition. Among these direct job-to-job transitions, we allow the compensating wage differential to differ if the job transition was initiated by a mass displacement event.[19] The structure of Table 7 is otherwise identical to the benchmark models in Table 2.

[19]We assemble the data for this analysis as follows. First, we define mass displacement events. Following the literature (Jacobson, LaLonde and Sullivan 1993; Abowd, McKinney and Vilhuber 2009; Couch and Placzek 2010; David and von Wachter 2011), we restrict attention to establishments with at least fifty FTE employees, and say a mass displacement occurred if FTE employment decreased by at least thirty percent. Next, we merge the mass displacement indicator to the complete set of longitudinal work histories in the analysis data. For each worker, we take only observations that are within two years of a job-to-job transition. Out of a total sample of 44,224,540 observations associated with job-to-job transition, 3,808,443 are job-years at firms experiencing mass displacements, and 9,302,630 occur within a 2-year window surrounding a mass displacement event. Our goal is to contrast the estimated compensating wage differential that uses variation from all job-to-job transitions with estimates restricted to mass displacement events.

Table 7: Mass Displacement Estimates

	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) OME	(5) TWFE
Fatality Rate (3-Yr MA)	0.475*	0.079*	–0.011*	0.205*	0.193*
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)
Fatality Rate × Mass Disp.	0.209*	0.003		–0.014*	–0.012*
	(0.002)	(0.002)		(0.002)	(0.002)
Zero Fatality Rate	0.089*	0.013*	–0.004*	0.016*	0.016*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Zero Fatality Rate × Mass Disp.	–0.006*	0.004*		0.005*	0.004*
	(0.001)	(0.001)		(0.000)	(0.000)
Mass Disp. Origin	–0.023*	0.016*		0.009*	0.009*
	(0.000)	(0.000)		(0.000)	(0.000)
Mass Disp. Destination	–0.031*	0.002*		0.001	–0.000
	(0.000)	(0.000)		(0.000)	(0.000)
N	44,220,194	44,224,540	44,224,540	44,224,540	44,224,540
R-Sq	0.448	0.914	0.976	0.925	0.925
VSL (millions of reais)	5.12	0.86	-0.12	2.21	2.08
95% CI	[5.09, 5.14]	[0.82, 0.89]	[-0.17, -0.07]	[2.18, 2.24]	[2.05, 2.11]

Notes: Models 1 to 4 correspond to the specifications reported in Table 2, and model 5 is an AKM two-way fixed effects model that includes worker effects, establishment effects, 1-digit occupation effects, experience effects, and year effects. The sample is restricted to observations within two years of a job-to-job transition at establishments with at least 50 FTE workers. ' Mass Disp.' indicates that the observation is associated with a job-to-job move in which the worker separated from an establishment experiencing a mass displacement episode. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage. * Indicates significance at the 0.01 level.

Column (2) shows the price of risk identified by mass displacement events (0.082) is very similar to the base coefficient (0.079). Likewise, Column (4) shows that in the OME model, the estimated compensating wage differential for displaced workers (0.191) is slightly smaller than the estimate (0.205) for non-displaced workers in this sample of job-to-job transitions in firms with at least fifty FTE employees. Both results suggest that the scope for bias arising from drift in unobserved ability or match effects is economically negligible.

6.3 Sensitivity to Violations of the Separability Assumption

In our evaluation of modeling assumptions in Section 5.2.2, we follow Card et al. (2013) in interpreting Figure 3 as supportive of the assumption that log wages are separable in unobserved worker and firm heterogeneity. However, the mean log wage residual is deviates from zero in the tails of the joint distribution of (θ, Ψ). While the magnitude of this deviation is small relative to estimated establishment effects, it is not small relative to the wage variation associated with fatality risk. Our model suggests these unmodeled interaction effects may bias the estimated compensating wage differential downward.

Table 8: Sensitivity of $\widehat{\gamma}$ to Excluding Tails of the (θ, Ψ) Joint Distribution

Sample	Pooled	Worker Effects	OME
5th to 95th Percentiles	0.308*	0.037*	0.170*
	(0.001)	(0.001)	(0.001)
10th to 90th Percentiles	0.282*	0.035*	0.170*
	(0.001)	(0.001)	(0.001)
15th to 85th Percentiles	0.261*	0.035*	0.171*
	(0.001)	(0.001)	(0.001)
20th to 80th Percentiles	0.244*	0.039*	0.174*
	(0.001)	(0.001)	(0.001)
25th to 75th Percentiles	0.223*	0.043*	0.180*
	(0.001)	(0.001)	(0.001)
30th to 70th Percentiles	0.201*	0.048*	0.187*
	(0.001)	(0.001)	(0.001)
35th to 65th Percentiles	0.175*	0.051*	0.196*
	(0.001)	(0.001)	(0.001)
40th to 60th Percentiles	0.154*	0.054*	0.204*
	(0.001)	(0.001)	(0.001)
45th to 55th Percentiles	0.138*	0.053*	0.207*
	(0.001)	(0.001)	(0.002)

Notes: Coefficients are estimated values of γ from the corresponding pooled, worker effects, and OME models using the main analysis sample, keeping only person-year observations in which either the estimated person effect (θ) or establishment effect (Ψ) falls within the percentile ranges indicated in each row. Log wages are Winsorized at the 1st and 99th percentiles. * Indicates significance at the 0.01 level.

To address this concern, we re-estimate each of our benchmark specifications allowing for different estimates of fatality risk in regions where the separability assumption is supported by the diagnostic evidence in Figure 3. Table 8 presents these estimates. In the first row, the reported coefficients are from an interaction between the fatality rate and an indicator variable that equals 1 if the observation is between the 5th and 95th percentiles of either the $\widehat{\theta}$ distribution or the $\widehat{\Psi}$ distribution. It is identified discarding variation coming from jobs involving either low-wage workers (below the 5th percentile worker effect) or low-wage establishments (below the 5th percentile establishment effect).[20] Going down the rows, the estimates are based on increasingly restricted sets of jobs with values of θ or Ψ closest to the median values, where Figure 3 is most supportive of the separability assumption.

[20] A separate coefficient (not shown) is estimated for the interaction between the fatality rate and the remaining observations.

There are several patterns of interest in these results. First, excluding the corners of the (θ, Ψ) distribution has relatively little impact on the baseline estimates, corroborating the interpretation by Card et al. (2013) that a very similar distribution of residuals in West Germany was consistent with only minimal evidence of match effects. For example, keeping only observations in the interquartile ranges of θ and Ψ decreases the pooled estimate to 0.22, but has little effect on the worker effects estimate (0.043) or the OME estimate (0.18). Second, going down column 1, γ monotonically decreases from 0.308 to 0.138 as the identifying variation is restricted to observations in which worker effects only vary across establishments in the middle of the Ψ distribution, and establishment effects only vary across workers in the middle of θ distribution. The bottom row of column 1 has similar properties to the worker effects model, in which variation is driven by workers with median ability moving to establishments with higher Ψ, the pooled model yields a γ below the OME estimate.

However, going down column 3, the OME estimates are relatively insensitive to differences in the average wage residual. Consistent with our model, γ increases slightly when the potentially problematic portions of the (θ, Ψ) distribution are not used for identification. The increase in γ is both gradual and monotonic, from 0.170 to 0.207, when the observations used for identification are limited to those in the middle of the distribution of unobservables, observations in which the evidence is most supportive of the separability assumption.

6.4 Relaxing the Linearity Assumption: Nonparametric Wage Profiles

To this point, we have followed the hedonic search literature in assuming the relationship between risk and wages is approximately log-linear. This assumption is not needed for our empirical approach nor by our conceptual framework, but once we abandon the linear approximation we do not have a strong prior on the correct functional form for the conditional expectation of wages given risk. Therefore, in Figure 6 we consider non-parametric estimates. These estimates suggest the linear approximation is quite good.

To estimate the conditional relationship between risk and wage non-parametrically, we discretize fatality risk into 75 equally-spaced bins for risk levels between 0 and 0.3 deaths per 1,000 FTFY worker-years. We then re-estimate our models including a full set of dummies for each bin along with a linear control for fatality rates above 0.3. Figure 6 plots the point estimates from the worker effects model (Panel a) and OME model (Panel b), along with smoothed spline functions fitted to the parameter estimates. In both models, the relationship between wages and risk is not severely mis-represented by the linear approximation. In the OME model, the results suggest that, if anything, willingness-to-pay for safety is gradually increasing in fatality risk, consistent with many models of preferences (Pratt and Zeckhauser 1996).

Figure 6: Non-parametric Estimates of the Wage-Fatality Rate Profile

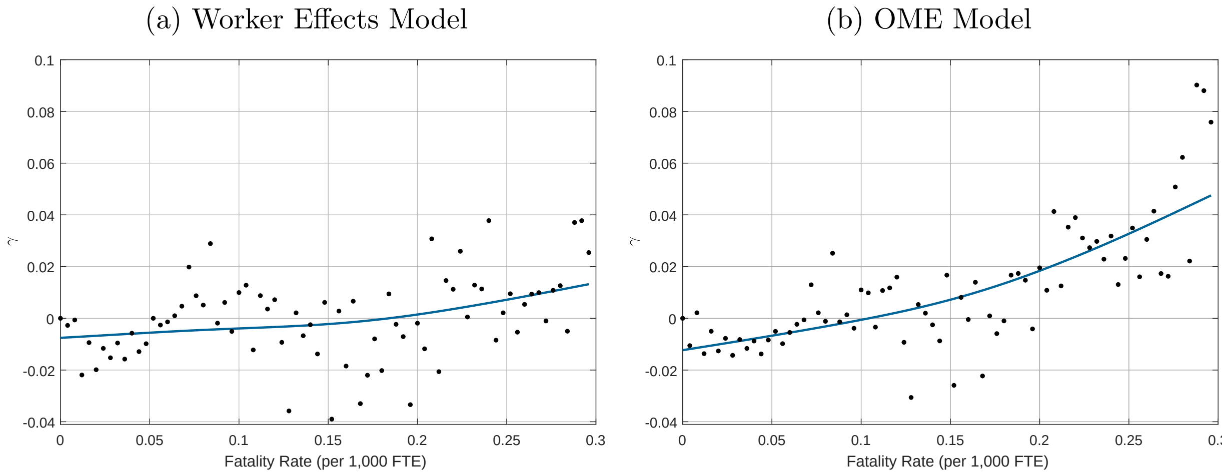

Notes: The vertical axis measures the estimated coefficients from a regression of log wages on 75 binary indicators for the fatality rate level, each representing a bin of width 0.004 deaths per 1,000 FTEs, and a continuous control for fatality rates above 0.3. Fitted lines are smoothed spline functions.

In the worker effects model, we see some evidence of non-monotonicity between 0.05 and 0.1 deaths per 1,000 FTE jobs, as well as in the vicinity of zero. While a non-monotonic relationship between wages and fatality risk is not technically impossible, it is not consistent with most conventional models of hedonic pricing. These patterns support our assertion that this model is mis-specified, particularly in light of the auxiliary evidence in Figure 4 and the associated Caetano (2015) diagnostic, as discussion in Section 2.4. The misspecification from the worker effects model appears to be largely remedied in the OME model according to this diagnostic.

6.5 Job Tenure, the Probability of Separation, and Wages

Our model from Section 3 implies a direct connection between job durations, fatality risk, and wages. First, it suggests that estimating worker willingness-to-pay from duration data can be misleading. Second, it suggests that we can use completed tenure as a proxy for the utility term in the structural log wage equation, (12). We consider these implications in turn.

Table 9 reports models that predict whether a worker voluntarily separates from their job as a function of fatality risk, completed tenure, and wages. It is common in the empirical literature on hedonic search to use estimates from job duration models to back out estimates of worker willingness to pay for safety. As Gronberg and Reed (1994) initially showed, in conventional McCall-style models of job search, with appropriate parametric assumptions, willingness-to-pay can be identified as minus one times the ratio of the coefficient on fatality

Table 9: Probability of Job Separation

	Dependent Variable: Voluntary Separation			
	(1)	(2)	(3)	(4)
Log Wage	–0.019*	–0.013*	–0.014*	–0.009*
	(0.000)	(0.000)	(0.000)	(0.000)
Fatality Rate	0.020*	0.012*	0.012*	–0.001
	(0.000)	(0.000)	(0.000)	(0.000)
Zero Fatality Rate	–0.001*	–0.001*	–0.001*	–0.000*
	(0.000)	(0.000)	(0.000)	(0.000)
Tenure (years)		–0.002*	–0.002*	–0.002*
		(0.000)	(0.000)	(0.000)
Establishment Size			Y	
Establishment Effects				Y
N	83,411,371	83,411,371	83,411,371	83,411,371
R-Sq	0.016	0.018	0.019	0.074

Notes: Estimates are from linear probability models in which the dependent variable is an indicator for whether the worker voluntarily separates from their dominant job in the current year. In addition to those reported, the models include the same controls as the pooled specification in Table 2, except for establishment size controls, which are introduced in column (3). Column (4) includes establishment effects. Fatality rates are measured in deaths per 1,000 full-time full-year equivalent workers. Robust standard errors are reported in parenethese. * Indicates significance at the 0.01 level.

risk to the coefficient on log wages, rescaled by the wage. This approach generally results in much larger estimates of willingness to pay than the corresponding estimates from hedonic wage models suggest (Dale-Olsen 2006; Bonhomme and Jolivet 2009). Our data also conform to this pattern. Table 9 presents estimates from linear separation models that control for log wages, fatality rates, and the observed covariates included in the pooled wage regressions. One useful feature of the data is that it includes the cause of any job separations. Since the predictions from our model relate to separations initiated by workers, the dependent variable equals 1 if the worker voluntarily resigns from his job, as opposed to a job ending due to employer-initiated termination, retirement, death, or the expiration of a contract, among other causes.

Estimates in column (1) suggest that a one log-point decrease in wages reduces the probability of voluntarily separating by 1.9%, while an increase in the fatality rate of one death per 1,000 worker-years increases the probability of separation by 2.0%. Applying the Gronberg and Reed (1994) approach to the results in Column (1) implies $\frac{\partial \ln w}{\partial R} = 1.05$, several times larger than the direct estimates from our hedonic wage models, which range between 0.16 and 0.20.

The discrepancy between estimates based on duration and wage data is generally interpreted through the lens of hedonic utility posting models. In those models, the (cross-sectional) hedonic wage equation is badly misspecified, while the model for job separations is less so. Our conceptual framework and preceding empirical work suggest the mis-specification of the hedonic wage equation can be corrected. However, the model for job separation is more challenging to correct.

Table 10: Compensating Wage Differentials for Full-Time Prime-Age Men, Completed Jobs Sample

	Pooled		Worker Effects		OME	
	(1)	(2)	(3)	(4)	(5)	(6)
Fatality Rate (3-Yr MA)	0.373*	0.407*	0.037*	0.043*	0.199*	0.200*
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Zero Fatality Rate	0.064*	0.061*	0.009*	0.010*	0.018*	0.018*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Completed Job Tenure		0.003*		0.001*		0.001*
		(0.000)		(0.000)		(0.000)
N	23,518,979	23,518,979	23,520,871	23,520,871	23,520,871	23,520,871
R-Sq	0.441	0.464	0.902	0.903	0.924	0.924
VSL (millions of reais)	3.61	3.95	0.36	0.42	1.93	1.94
95% CI	[3.58, 3.64]	[3.92, 3.97]	[0.32, 0.40]	[0.38, 0.46]	[1.89, 1.97]	[1.90, 1.98]

Notes: All models are the same as the corresponding benchmark specifications in Table 2. The analysis sample includes only completed dominant jobs of men 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. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage. * Indicates significance at the 0.01 level.

The standard separation model omits unobserved firm- and occupation-specific amenities, which we expect to be correlated with wages and risk. Theory suggests high wage firms are likely to offer better amenities, and if so the effect of wages on job separation could be overstated without controlling for these unobserved amenities. This pattern is evident moving from column (1) to column (4) in Table 9. Column (2) controls for tenure, and each of the coefficients of interest is attenuated somewhat. Column (3) adds establishment size controls, which has little additional effect.

Extending this model to introduce fixed establishment effects, or to control for predicted establishment wage heterogeneity, is challenging because as we show in Table 3 the majority of wage gains associated with job separations occur when workers move into higher paying establishments. Consistent with this conclusion, column (4) includes fixed establishment effects, and indicates that within establishments there is no economically meaningful relationship between either log wages (0.009) or fatality rates (0.001) and separation rates. However, the inclusion of establishment effects is necessary to alleviate bias due to unobserved establishment-heterogeneity in amenities.[21]

Overall, this analysis suggests that, at the very least, estimates of willingness-to-pay from duration data are sensitive to attempts to control for the presence of unobserved amenities. By contrast, the TWFE and OME models easily accommodate variation in unobserved firm- and occupation-level amenities that are arbitrarily correlated with wages and fatality risk.

Finally, our model also implies we can use completed tenure as a proxy for the utility component in Equation (12). As in Abraham and Farber (1987), completed tenure varies with the unobserved match-specific component of utility that we expect is correlated with wages. Table 10 re-estimates all of our models restricted to the sample of jobs for which we observe completed tenure. We report each specification estimated on this selected sample, and then report the model including completed tenure. In the pooled model and in the worker effects model, including completed tenure leads to non-trivial increases in the estimated effect of fatality risk on wages. By contrast, there is no significant effect on the estimated coefficient on fatality risk in the OME model from adding completed tenure to the model.[22]

7 Conclusion

Our objectives have been to demonstrate the advantages of using matched employer-employee data to correct for the effects of endogenous mobility bias in estimating compensating wage differentials. Controlling for employer heterogeneity in a relatively straightforward way yields results that are strikingly consistent with the implications of basic hedonic search models. We can furthermore clearly articulate, and provide empirical support for, the conditions under which the estimated compensating wage differential identifies workers' marginal willingness to accept fatal risk.

Our paper complements a growing body of work addressing the effects of search frictions and endogeneity bias when estimating the effects of non-wage amenities on labor market outcomes. Much of the recent work uses cross-sectional and panel data to estimate structural models of hedonic search (Bonhomme and Jolivet 2009; Dey and Flinn 2005; 2008; Villanueva 2007; Sullivan and To 2014) and Roy-style sorting (DeLeire et al. 2013). An emerging literature addresses models of compensating differentials using matched employer-employee data.

[21] Although these estimates are from linear probability models, we find very similar patterns from logit specifications.

[22] Another approach would be to estimate the tenure and wage equation jointly, as in Bonhomme and Jolivet (2009). To do so would require strong assumptions on the nature of the joint distribution of individual and establishment heterogeneity and distract from our main objective of highlighting the bias introduced in hedonic wage models by worker sorting across jobs with different compensation practices.

In very innovative recent papers, Sorkin (2016) and Taber and Vejlin (2016) seek to explain how much variation in matching outcomes, job duration and wages can be rationalized by compensating differentials. In these analyses, unlike our paper, job amenities are not measured; the presence of amenities is inferred from variation in outcomes. Also, we are not the first to estimate compensating wage differentials using matched employer-employee data Lalive (2003) and Tsai et al. (2011) estimate hedonic wage models using matched employer-employee data with observed firm-level amenities. However, the emphasis in both papers is limited to studying the effects of aggregation bias associated with measuring amenities using industry averages.

Our paper is the first to use matched employer-employee data to directly illustrate how endogenous mobility arising from job search can bias estimates of compensating wage differentials. In doing so, we provide a bridge between the structural, theoretical, and reduced-form literatures. Specifically, this paper shows the statistical decomposition of wages originating with Abowd et al. (1999) does an extremely good job of matching the predictions of the basic hedonic search model, and in explaining the covariation between wages and job characteristics.

The analysis of hedonic wage models is fraught with challenges for applied work, and no study can resolve them all. Future work must address key measurement issues that were beyond the scope of this study. One trade-off associated with using administrative data, rather than survey data, is that we do not observe information on other job amenities. However, another advantage of our empirical model is that by controlling for establishment and occupation effects we actually reduce omitted variable bias associated with unobserved amenities, like health insurance, that are often employer-specific. On the other hand, if, for example, fatal and non-fatal risk tend to be bundled together in the same way across jobs within establishments, then our model estimates the compensating differential for changes in this composite bundle. This interpretive issue is common to all studies that use observational data to study the determinants of compensation in the labor market.

There are reasons to suspect that the endogenous mobility problem we highlight is not unique to Brazil. Our analysis is motivated in part by the contrast between cross-sectional estimates of the compensating differential for fatal injury and the much smaller estimates from U.S. panel data. This pattern is consistent with hedonic search. There is a good chance that employer and match-specific variation in wages could explain the U.S. data. Woodcock (2008) estimates that among workers in the US who experience job-to-job transitions, about 60% of their earnings growth is due to sorting into firms that pay higher average earnings to all workers for unobserved reasons.

While the focus of our paper is not on estimating the VSL *per se*, it is useful to characterize the magnitude of our estimates in the context of this widely-used policy parameter. For comparison, we follow the the benefit transfer method to rescale VSL estimates from the US into Brazilian reais, as recommended by Hammitt and Robinson (2011) and as implemented by Miller and Façanha (2016). The benefit transfer method adjusts VSL estimates for differences in per-capita income, which has large effects on the scale of willingness to pay.[23] Applying this method to US estimates implies a VSL in Brazil of between 2.14 and 3.10 million 2003 reais. The VSL implied by our pooled model, 2.84 million reais (Column 1 of Table 2,) is right in the middle of this range. However, our benchmark OME model (Column 4 of Table 2) implies a VSL of 1.73 million reais, suggesting this range may be too high. Our analysis also suggests that VSL estimates based on within-worker models are much too low. This information should be taken into account in cost-benefit calculations.

Our results suggest that models of compensating differentials with costly search in the spirit of Hwang et al. (1998) can provide a useful guide to further empirical work. The approach we have developed is relevant for other non-wage amenities for which the literature suggests endogenous mobility bias may be present (Brown 1980; Garen 1988; Hersch 1998). There would also be considerable value in efforts to develop and estimate a structural model in the spirit of Bonhomme and Jolivet (2009) or Lavetti (2015) that can address the simultaneous determination of wages and job tenure given workers' forward-looking behavior.

[23] 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 dollars, 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:

$$VSL_{Brazil} = VSL_{US} \times \frac{\text{GNI per capita in Brazil}}{\text{GNI per capita in U.S.}}. \tag{18}$$

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 (18) 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.

References

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

Abowd, J. M., McKinney, K. L. and Schmutte, I. M. (2015). Modeling endogenous mobility in earnings determination, *Journal of Business & Economic Statistics* **0**(ja): 0–0.

Abowd, J. M., McKinney, K. L. and Vilhuber, L. (2009). The link between human capital, mass layoffs, and firm deaths, *Producer Dynamics: New Evidence from Micro Data*, NBER Chapters, National Bureau of Economic Research, Inc, pp. 447–472.

Abraham, K. G. and Farber, H. S. (1987). Job duration, seniority, and earnings, *The American Economic Review* **77**(3): 278–297.

Abraham, K. G. and Spletzer, J. R. (2010). *Are the New Jobs Good Jobs?*, University of Chicago Press, pp. 101–143.

Barth, E., Bryson, A., Davis, J. C. and Freeman, R. (2016). Its where you work: Increases in the dispersion of earnings across establishments and individuals in the united states, *Journal of Labor Economics* **34**(S2): S67–S97.

Bonhomme, S. and Jolivet, G. (2009). The pervasive absence of compensating differentials, *Journal of Applied Econometrics* **24**(5): 763–795.

Brown, C. (1980). Equalizing differences in the labor market, *The Quarterly Journal of Economics* **94**(1): 113–34.

Caetano, C. (2015). A test of exogeneity without instrumental variables in models with bunching, *Econometrica* **83**(4): 1581–1600.

Caetano, G. and Maheshri, V. (2013). Do 'broken windows' matter? identifying dynamic spillovers in criminal behavior, *Working Papers 2013-252-22*, Department of Economics, University of Houston.

Card, D., Cardoso, A. R., Heining, J. and Kline, P. (2016). Firms and labor market inequality: Evidence and some theory, *Working Paper 22850*, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w22850*

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.

Cornwell, C., Rivera, J. M. and Schmutte, I. M. (forthcoming). Wage discrimination when identity is subjective: Evidence from changes in employer-reported race, *Journal of Human Resources* .

Couch, K. A. and Placzek, D. W. (2010). Earnings losses of displaced workers revisited, *American Economic Review* **100**(1): 572–89.

Dale-Olsen, H. (2006). Estimating workers' marginal willingness to pay for safety using linked employer-employee data, *Economica* **73**(289): 99–127.

David, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss, *Brookings Papers on Economic Activity* pp. 1–72.

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.

Dey, M. and Flinn, C. (2008). Household search and health insurance coverage, *Journal of Econometrics* **145**(1-2): 43–63.

Dey, M. S. and Flinn, C. J. (2005). An equilibrium model of health insurance provision and wage determination, *Econometrica* **73**(2): 571–627.

Garen, J. (1988). Compensating wage differentials and the endogeneity of job riskiness, *The Review of Economics and Statistics* **70**(1): 9–16.

Gasparini, L. and Tornarolli, L. (2009). Labor informality in Latin America and the Caribbean: Patterns and trends from household survey microdata, *Desarrollo Y Sociedad* **63**(1): 13–80.

Gibbons, R. and Katz, L. (1992). Does unmeasured ability explain inter-industry wage differentials?, *The Review of Economic Studies* **59**(3): 515–535.

Gibbons, R., Katz, L. F., Lemieux, T. and Parent, D. (2005). Comparative advantage, learning and sectoral wage determination, *Journal of Labor Economics* **23**: 681–723.

Gronberg, T. J. and Reed, W. R. (1994). Estimating workers' marginal willingness to pay for job attributes using duration data, *The Journal of Human Resources* **29**(3): 911–931.

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.

Hwang, H.-s., Mortensen, D. T. and Reed, W. R. (1998). Hedonic wages and labor market search, *Journal of Labor Economics* **16**(4): 815–47.

Hwang, H.-s., Reed, W. R. and Hubbard, C. (1992). Compensating wage differentials and unobserved productivity, *Journal of Political Economy* **100**(4): pp. 835–858.

Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993). Earnings losses of displaced workers, *The American Economic Review* **83**(4): 685–709.

Kambourov, G. and Manovskii, I. (2008). Rising occupational and industry mobility in the United States: 1968-97, *International Economic Review* **49**(1): 41–79.

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.

Lang, K. and Majumdar, S. (2004). The pricing of job characteristics when markets do not clear: Theory and policy implications, *International Economic Review* **45**(4): 1111–1128.

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

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*

Moscarini, G. and Thomsson, K. (2007). Occupational and job mobility in the US, *Scandinavian Journal of Economics* **109**(4): 807–836.

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

Pratt, J. W. and Zeckhauser, R. J. (1996). Willingness to pay and the distribution of risk and wealth,, *Journal of Political Economy* **104**: 747–763.

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

Schmutte, I. M. (2015). Job referral networks and the determination of earnings in local labor markets, *Journal of Labor Economics* **33**(1): 1–32.

Solon, G. (1988). Self-selection bias in longitudinal estimation of wage gaps, *Economics Letters* **28**(3): 285 – 290.

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

Speer, J. D. (2016). How bad is occupational coding error? A task-based approach, *Economics Letters* **141**: 166 – 168.

Sullivan, P. and To, T. (2014). Search and nonwage job characteristics, *Journal of Human Resources*

49(2): 472–507.

Taber, C. and Vejlin, R. (2016). Estimation of a Roy/search/compensating differential model of the labor market, *IZA discussion paper*, IZA.

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.

Villanueva, E. (2007). Estimating compensating wage differentials using voluntary job changes: Evidence from Germany, *Industrial and Labor Relations Review* **60**(4): 544–561.

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.

Viscusi, W. K. and Aldy, J. E. (2003). The value of a statistical life: A critical review of market estimates throughout the world, *Journal of Risk and Uncertainty* **27**(1): 5–76.

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

A Additional Tables and Figures – For Web Publication Only

Table A.1: Causes of Separation Reported in RAIS

Value	Label Portuguese	Label English
0	nao desl ano	no separation this year
10	dem com jc	terminated with just cause
11	dem sem jc	terminated without just cause
12	term contr	end of contract
20	desl com jc	resigned with just cause
21	desl sem jc	resigned without just cause
30	trans c/onus	xfer with cost to firm
31	trans s/onus	xfer with cost to worker
40	mud. regime	Change of labor regime
50	reforma	military reform - paid reserves
60	falecimento	demise, death
62	falec ac trb	death - at work accident
63	falec ac tip	death - at work accident corp
64	falec d prof	death - work related illness
70	apos ts cres	retirement - length of service with contract termination
71	apos ts sres	retirement - length of service without contract termination
72	apos id cres	retirement - age with contract termination
73	apos in acid	retirement - disability from work accident
74	apos in doen	retirement - disability from work illness
75	apos compuls	retirement - mandatory
76	apos in outr	retirement - other disability
78	apos id sres	retirement - age without contract termination
79	apos esp cre	retirement - special with contract termination
80	apos esp sre	retirement - special without contract termination

Table A.2: Average Fatality Rates By Industry and Occupation

Industry	Average Fatality Rate	Number of Job-Years
Agriculture and Fishing	10.25	22,762,420
Mining	10.48	1,814,957
Manufacturing	5.24	76,712,576
Utilities	4.19	2,023,931
Construction	13.77	26,098,278
Trade and Repair	6.04	82,004,063
Food, Lodging, and Hospitality	4.99	15,589,304
Transportation, Storage, and Communication	14.53	20,941,098
Financial and Intermediary Services	1.01	6,947,728
Real Estate, Renting, and Services	4.59	57,447,503
Public Administration, Defense, and Public Security	0.84	72,055,976
Education	1.58	12,418,485
Health and Social Services	1.67	14,089,834
Other Social and Personal Services	3.98	15,469,519
Domestic Services	5.76	116,086
Occupation		
Public Administration and Management	2.63	18,035,409
Professionals, Artists, and Scientists	1.09	39,178,629
Mid-Level Technicians	2.50	40,972,375
Administrative Workers	1.87	78,792,943
Service Workers and Vendors	4.40	98,796,568
Agriculture Workers, Fishermen, Forestry Workers	9.26	25,417,204
Production and Manufacturing I	11.65	94,955,794
Production and Manufacturing II	5.28	15,947,072
Repair and Maintenence Workers	7.39	13,871,753

Notes: Average fatality rates are calculated as deaths per 100,000 full-time full-year-equivalent workers using the 100% Brazilian RAIS data from 2003-2010.

Table A.3: Estimated Compensating Wage Differentials for Full-Time Prime-Age Men, Excluding Industry and Occupation Effects

	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) OME
Fatality Rate	0.343*	0.068*	–0.006*	0.156*
	(0.001)	(0.001)	(0.001)	(0.001)
Zero Fatality Rate	0.211*	0.022*	–0.006*	0.018*
	(0.000)	(0.000)	(0.000)	(0.000)
N	83,411,371	83,418,032	83,418,032	83,418,032
R-Sq	0.377	0.912	0.978	0.930
VSL (millions of reais)	3.49	0.70	-0.06	1.59
95% CI	[3.48, 3.51]	[0.68, 0.71]	[-0.09, -0.03]	[1.57, 1.61]

Notes: Model 1 also includes year effects, state effects, race effects, years of experience effects (censored at 30), indicators for small and medium-sized establishments, and education effects. Model 2 includes worker effects and the same controls as Model 1 except for race and education. Model 3 includes job-match effects, years of experience effects, and year effects. Model 4 includes worker effects, establishment effects, and year effects. The analysis sample includes dominant jobs of men 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. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage, and reported in millions of reais. * Indicates significance at the 0.01 level.

Table A.4: Heterogeneity by Region in Brazil

	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) OME	Avg $\ln(Wage)$	Avg Fatality	Var Ψ
Northern Region	0.366*	0.046*	0.026*	0.154*	1.227	0.094	0.081
	(0.003)	(0.004)	(0.008)	(0.004)			
Northeast Region	0.682*	0.094*	–0.007	0.165*	1.076	0.080	0.073
	(0.002)	(0.002)	(0.004)	(0.002)			
Southeast Region	0.208*	0.025*	–0.016*	0.170*	1.486	0.081	0.089
	(0.001)	(0.001)	(0.002)	(0.001)			
South Region	0.265*	0.065*	0.010*	0.173*	1.420	0.086	0.070
	(0.002)	(0.002)	(0.003)	(0.002)			
Central West Region	0.242*	0.006	0.031*	0.139*	1.306	0.089	0.079
	(0.003)	(0.004)	(0.006)	(0.003)			

Notes: All models and sample selection criteria are identical to those in Table 2, except that they are estimated separately by region. Average fatality rates are measured in deaths per 1,000 full-time full-year equivalent workers. * Indicates significance at the 0.01 level.

Table A.5: Sensitivity of OME Estimates to Model Specification

	(1)	(2)	(3)	(4)	(5)
Fatality Rate	0.168*	0.190*	0.165*	0.172*	0.152*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Zero Fatality Rate	0.013*	0.014*	0.012*	0.013*	0.007*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
1st Stage Exp. by Educ. Effects	Y	N	N	N	N
1st Stage Replace Exp. with Tenure Effects	N	Y	Y	N	N
2nd Stage Include Exp. Effects	N	N	Y	N	N
2nd Stage Include Hiring Year by Year Effects	N	N	N	Y	N
1st Stage Cubic in Exp. Interacted with Race	N	N	N	N	Y
N	83,411,371	83,418,032	83,418,032	83,418,032	83,418,032
R-Sq	0.914	0.935	0.936	0.931	0.967
VSL (millions R$)	1.71	1.93	1.69	1.75	1.55
95% CI	[1.70, 1.73]	[1.92, 1.95]	[1.67, 1.70]	[1.74, 1.77]	[1.53, 1.58]

Notes: All models are similar to the OME specification reported in Table 2 except for the indicated differences in control variables. The analysis sample includes dominant jobs of men 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. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage, and reported in millions of reais. * Indicates significance at the 0.01 level.

Table A.6: Alternative AKM TWFE Model Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Fatality Rate	0.165*	0.168*	0.165*	0.165*	0.169*	0.153*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Zero Fatality Rate	0.014*	0.013*	0.014*	0.014*	0.013*	0.018*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
1-Digit Occ. Effects	Y	Y	Y	Y	Y	N
Linear Tenure Control	Y	N	Y	Y	Y	Y
Tenure Effects	N	Y	N	N	N	N
Experience by Education Effects	N	N	Y	N	N	N
Hiring Year Effects	N	N	N	Y	Y	N
Year-by-Hiring Year Effects	N	N	N	N	Y	N
N	83,418,032	83,418,032	83,411,371	83,418,032	83,418,032	83,418,032
R-Sq	0.931	0.931	0.931	0.931	0.931	0.930
VSL (millions R$)	1.68	1.72	1.68	1.68	1.72	1.56
95% CI	[1.67, 1.70]	[1.70, 1.73]	[1.67, 1.70]	[1.67, 1.70]	[1.71, 1.74]	[1.55, 1.58]

Notes: All specifications are AKM two-way fixed effects models, and include worker effects, establishment effects, year effects, and experience effects (censored at 30). The analysis sample includes dominant jobs of men 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. Fatality Rate is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage. * Indicates significance at the 0.01 level.

Table A.7: Estimates with Clustered Standard Errors

	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) OME
Fatality Rate	0.279	0.037	–0.006	0.170
Unclustered SE	(0.001)*	(0.001)*	(0.001)*	(0.001)*
Clustered by Establishment	(0.018)*	(0.004)*	(0.009)	(0.003)*
Clustered by Occupation*Industry	(0.163)	(0.033)	(0.029)	(0.032)*
Zero Fatality Rate	0.073	0.008	–0.006	0.014
Unclustered SE	(0.000)*	(0.000)*	(0.000)*	(0.000)*
Clustered by Establishment	(0.004)*	(0.001)*	(0.001)*	(0.001)*
Clustered by Occupation*Industry	(0.022)*	(0.006)	(0.009)	(0.006)
N	83,411,371	83,418,032	83,418,032	83,418,032
N Establishment Clusters	1,634,452	1,634,464	1,634,464	1,634,464
N Occupation-Industry Clusters	624	624	624	624
R-Sq	0.458	0.913	0.978	0.930

Notes: Models and the analysis sample are the same as those in Table 2. 'Fatality Rate' is measured in deaths per 1,000 full-time full-year equivalent workers. Occupation-Industry clusters use 3-digit occupation codes and 2-digit industry codes. * Indicates significance at the 0.01 level.

Table A.8: Descriptive Statistics: IV Sample

	IV Sample
Race *pardo* or *preto*	0.42
Elementary or less	0.44
Some High School	0.08
High School	0.39
Some College	0.03
College or More	0.06
Log Hourly Wage	1.41
Total Experience (Years)	18.83
Fatality Rate (per 100,000)	8.10
Zero Fatality Rate (Percent)	7.77
Number of Observations	$5,652,917$

NOTE–Means of key variables for the sample used to estimate IV models. See text for a complete description of the sample restrictions

Table A.9: Compensating Wage Differentials for Full-Time Prime-Age Men: IV Sample

Dependent Variable:		$\ln(Wage)$		$\ln(Wage) - X\widehat{\beta}$
	(1)	(2)	(3)	(4)
	Pooled	Worker Effects	Match Effects	Orth. Match Effects
Fatality Rate (3-Yr MA)	0.365*	0.011	–0.006	0.203*
	(0.003)	(0.014)	(2.079)	(0.009)
Zero Fatality Rate	0.079*	0.014*	0.001	0.021*
	(0.001)	(0.003)	(0.188)	(0.001)
N	5,403,258	5,403,738	5,403,738	1,522,702
R-Sq	0.458	0.977	1.000	0.958
VSL (millions of reais)	3.02	0.09	-0.05	1.68
95% CI	[2.97, 3.08]	[-0.14, 0.32]	[-33.79, 33.69]	[1.53, 1.82]

Notes: Estimates of benchmark specifications restricted to the IV sample. Model 1 also includes 1-digit industry effects, 1-digit occupation effects, year effects, state effects, race effects, years of experience effects (censored at 30), indicators for small and medium-sized establishments, and education effects. Model 2 includes worker effects and the same controls as Model 1 except for race and education. Model 3 includes job-match effects, years of experience effects, and year effects. Model 4 includes worker effects, establishment effects, 1-digit occupation effects, and year effects. The analysis sample includes dominant jobs of men 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. is measured in deaths per 1,000 full-time full-year equivalent workers. Log wages are Winsorized at the 1st and 99th percentiles. VSLs calculated at mean hourly wage. * Indicates significance at the 0.01 level.

B Appendix – For Web Publication Only

B.1 Model Derivations Omitted from Main Text

In this appendix we derive the probability a firm's offer is accepted and show that it is approximately exponential in the posted utility. It is convenient to define the probability that a firm's offer is accepted conditional on being made. Importantly, we will not condition on whether the offer is made to the firm's inside or outside workers.

For simplicity, in this appendix we elimiate the distinction between firms and occupations and use the notation m to differentiate jobs. The key stochastic elements in the model are

- The variable ω_{im} is an indicator equal to one if the worker i receives an offer from job m. We assume offers are made independently and with equal probability, λ. We say $\Pr[\omega_{im} = 1] = \lambda$. For any worker i, the vector ω_i to be the $JK \times 1$ vector whose mth entry is ω_{im}.

- The variable Z_{im} is an indicator for the event that an offer from m to i is accepted, conditional on the offer having been made.

- The variable $M_{im'}$ is an indicator for the event that i is employed at job m' at the beginning of the period.

- The $JK \times 1$ vector, V_S, records the exponential in the utility offer to type-s workers from each job m, $\exp(\bar{u}_{sm})$.

- The variable L_{sm} denotes the number of type s workers employed in job m.

The probability that an offer is accepted conditional on being made is $\Pr[Z_{im} = 1|\omega_{im} = 1]$. We can rewrite this as the expectation of winning over all possible consideration sets. We assume each worker always has the opportunity of returning to unemployment and receiving indirect utility $\bar{u}_{s0} + \epsilon_{i0}$. Appealing to the law of iterated expectations,

$$\Pr[Z_{im} = 1|\omega_{im} = 1] = \sum_{\omega_i|\omega_{im}=1} \Pr[Z_{im} = 1|\omega_i, \omega_{im} = 1] \Pr[\omega_i|\omega_{im} = 1]. \tag{19}$$

The inner term is, based on our assumptions about idiosyncratic utility and worker decision-making, a standard conditional choice probability

$$\Pr[Z_{im} = 1|\omega_i, \omega_{im} = 1] = \frac{\exp(\bar{u}_{s(i)m})}{\exp\left(\bar{u}_{s(i)0}\right) + \sum_{m'} \omega_{im} \exp\left(\bar{u}_{s(i)m'}\right)}. \tag{20}$$

So, the probability an offer is accepted the expectation of the conditional choice probabilities over the full set of feasible choice sets.

To compute this expectation, we need to derive $\Pr[\omega_i|\omega_{im} = 1]$. If all offers were made at random, this would be a simple product of JK trials of probability λ. Because firms make non-random offers to their current employees, we have

$$\Pr[\omega_{im} = 1] = \frac{L_{sm}}{N_s} + \lambda\left(1 - \frac{L_{sm}}{N_s}\right). \tag{21}$$

We assume all firms are extremely small relative to the market, so $\Pr[\omega_{im} = 1] \approx \lambda$.

The probability of worker i holding offer vector ω_i varies depending on the identity of his current employer. We therefore have

$$\Pr[\omega_i|\omega_{im} = 1] = \sum_{m'} \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \Pr[M_{im'} = 1|\omega_{im} = 1]. \tag{22}$$

The first term inside the summation on the right-hand side, when $m' \neq m$ is zero when $\omega_{im'} = 0$. Therefore, we can condition on $\omega_{im'} = 1$. When $m \neq m'$

$$\begin{aligned} &\Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \\ &= \tfrac{\Pi_\ell \Pr[\omega_{i\ell}=1]}{\Pr[\omega_{im}=1]\Pr[\omega_{im'}=1]} \\ &= \tfrac{\Pr[\omega_i]}{\Pr[\omega_{im}=1]\Pr[\omega_{im'}=1]}. \end{aligned} \tag{23}$$

The preceding derivation invokes the independence of offers across workers. When $m = m'$, we lose a piece of information, and so

$$\Pr[\omega_i|\omega_{im} = 1, M_{im} = 1] = \frac{\Pi_\ell \Pr[\omega_{i\ell} = 1]}{\Pr[\omega_{im} = 1]}. \tag{24}$$

So we have

$$\Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] = \frac{\Pr[\omega_i|\omega_{im} = 1, M_{im} = 1]}{\Pr[\omega_{im'} = 1]} \tag{25}$$

We calculate $\Pr[M_{im'} = 1|\omega_{im} = 1]$ using Bayes' Rule:

$$\Pr[M_{im'} = 1|\omega_{im} = 1] = \frac{\Pr[\omega_{im} = 1|M_{im'} = 1] \Pr[M_{im'} = 1]}{\Pr[\omega_{im} = 1]}. \tag{26}$$

When $m \neq m'$ the right-hand side is

$$\frac{\Pr[\omega_{im} = 1|M_{im'} = 1] \Pr[M_{im'} = 1]}{\Pr[\omega_{im} = 1]} = \frac{\lambda \frac{L_{sm'}}{N_s}}{\Pr[\omega_{im} = 1]} = \frac{\lambda L_{sm'}}{L_{sm}(1-\lambda) + \lambda N_s} \tag{27}$$

When $m = m'$ the right-hand side is

$$\frac{\Pr[\omega_{im} = 1|M_{im} = 1] \Pr[M_{im} = 1]}{\Pr[\omega_{im} = 1]} = \frac{\frac{L_{sm}}{N_s}}{\Pr[\omega_{im} = 1]} = \frac{L_{sm}}{L_{sm}(1-\lambda) + \lambda N_s} \tag{28}$$

Finally, we can rewrite the RHS of (22)

$$\begin{aligned} &\textstyle\sum_{\{m':\omega_{im'}=1\}} \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \Pr[M_{im'} = 1|\omega_{im} = 1] \\ &= \Pr(\omega_i|\omega_{im} = 1) \left[\Pr[M_{im} = 1|\omega_{im} = 1] + \textstyle\sum_{\{m':m'\neq m;\omega_{im'}=1\}} \tfrac{\Pr[M_{im'}=1|\omega_{im}=1]}{\Pr[\omega_{im'}=1]}\right] \end{aligned} \tag{29}$$

Making substitutions for the probability of current employment conditional on an observed offer, we get

$$\begin{aligned}&\textstyle\sum_{\{m':\omega_{im'}=1\}} \Pr\left[\omega_i|\omega_{im}=1, M_{im'}=1\right] \Pr\left[M_{im'}=1|\omega_{im}=1\right] \\ &= \Pr(\omega_i|\omega_{im}=1) \tfrac{1}{L_{sm}(1-\lambda)+\lambda N_s} \left\{ L_{sm} + \textstyle\sum_{m' \neq m} \tfrac{\lambda L_{sm'}}{\Pr[\omega_{im'}=1]} \right\}\end{aligned} \tag{30}$$

B.2 Details of Fatality Rate Calculations

Within a cell, c, we construct the fatality rate R_c as

$$R_c = \frac{F_c}{(H_c/2,000)} \times (100,000). \tag{31}$$

The numerator, F_c is the number of fatal injuries in cell c. The denominator is the number of full-time full-year-equivalent jobs, assuming a baseline 40 hour work week and a 50 week work year. H_c is the total number of contracted hours worked over the year.[24] For each job, j, in the cell c, we count the number of hours worked as $H_i = (MonthsWorked/12) * 50 * (Hours/Week)$. H_c is the sum of H_i over all i in cell c. Finally, we inflate the count by 100,000 for consistency with the BLS measure. In most of the paper, we re-scale the fatality rate to deaths per 1,000 workers for ease of presentation of results.

B.3 Brazil's Labor Market Institutions

B.3.1 Formal Employment

In Brazil a worker is formally employed if he or she has a registered identification number with one of two social security programs: the *Programa de Integração Social* (PIS), or Social Integration Program, or the *Programa de Formação do Patrimônio do Servidor Público* (PASEP), or Civil Servants Equity Formation Program, depending on whether the worker is employed in the private sector or the public sector. PIS/PASEP numbers are consistent across workers and follow a worker for life. For firms, formal employment means that the employer contributes the *Abono Salarial* along with other social security payments to a bank account administered by either *Caixa Econômica Federal* if registered with PIS, or *Banco do Brasil* for PASEP workers. Formal employers must also have employment contracts for all employees. The most common contract type is the *Consolidação das Leis de Trabalho* (CLT), or Labor Law Consolidation. Other contract types include internships, independent contractors, directorships and government contractors. The Brazilian government defines formal employment with these criteria, and this definition is consistent with definitions used by researchers when studying other Latin American economies (Gasparini and Tornarolli 2009). Formal employment grew steadily in Brazil during our sample period, from nearly 42 million jobs in 2003 to over 65 million jobs in 2010. Unemployment decreased from eleven percent to five percent, and real wages grew over the period as well. Our sample therefore covers a period of growth and tightening labor-market conditions.

[24]Changes in the definition of full-year work will only affect the scale of our fatality rates. We chose a definition close to the BLS definition, although in Brazil full-year work may be closer to 48 weeks.

B.3.2 Wage Regulations

The formal sector of Brazil's labor market is governed by several overlapping institutions, some understanding of which is relevant to the interpretation of our results. Our data record the total monetary compensation that the employer is contracted to pay the worker. The data do not report non-monetary compensation, including employer-provided health and life insurance. As in the U.S., in Brazil, life and health insurance are frequently provided by one's employer. The value of such insurance is another amenity whose provision may be associated with that of occupational safety and earnings. We note that this shortcoming of the data is common to almost the entire literature. Nevertheless, any structural interpretation of our results depends on standard assumptions that unobserved workplace amenities are conditionally uncorrelated with observed amenities.

Additionally, in Brazil, wages are tied to safety formally through health and safety regulations known as *Norma Regulamentarora de Seguranca e Saude no Trabalho* (NR). The NRs stipulate a schedule of wage premia to be paid in association with work activities deemed to be unpleasant or dangerous. If these wage setting institutions were strong, we would still expect to find evidence of compensating wage differentials, but their presence would complicate our interpretation of the estimates as measuring individual preferences. A complete accounting of this complex institutional environment would require richer data on the NRs and enforcement activity. However, a couple of factors suggest these institutions have a small effect on our data. First, the statutory premia are generally 10-20 percent of the Federal minimum wage, which is quite low in absolute terms, so likely to be non-binding. Second, and relatedly, compliance with NRs are not a focus of the enforcement activities of the labor ministry, as they have very little influence on health and safety outcomes. We therefore proceed under the assumption that these institutions do not substantially alter the behavior of workers and firms.

In Brazil the NRs are norms elaborated and enforced by the MTE. They seek to promote health and safety in the workplace in compliance with constitutional (art. 7, XXII) and statutory (CLT arts. 60, 189, 200) obligations, as well as with international agreements and standards. The NRs affect all employers of labor in the formal sector, both public and private. The NRs stipulate a schedule of wage premia to be paid in association with work activities deemed to be unpleasant or dangerous.

In practice, each establishment is required to produce, in consultation with health and safety specialists, a document classifying the degree of exposure to harm for all jobs (occupations) within the establishment (known in most sectors as a PPRA). According to the regulations set forth in the NRs and CLT, the resulting premium for the specific plant-occupation pair is set as a percentage between zero and forty percent of the Federal minimum wage. The employer can reduce the wage premium in two ways: first, by investing in collective risk mitigation mechanisms, which reduce risk exposure for the all workers, and second by investing in individual protection mechanisms, which reduce risk exposure for a specific worker.