

Estimating Compensating Wage Differentials with Endogenous Job Mobility*

Kurt Lavetti[†]
Ian M. Schmutte[‡]

January 17, 2025

Abstract

We show that the effect of job amenities on wages implied by Rosen's theory of equalizing differences can be identified under imperfect competition. Using administrative matched data from Brazil, we reconcile contradictory theoretical and empirical conclusions about compensating differentials by incorporating the role of firms in setting wages and amenities. This empirical framework approximates the Brazilian wage structure and job mobility patterns well. We develop a model with monopsonistic firms and search that rationalizes our empirical framework as an equilibrium outcome, and provides a mapping between exogeneity conditions and the interpretation of our empirical estimand relative to model primitives.

*We are grateful for helpful comments from John Abowd, David Blau, Molly Candon, David Card, Chris Cornwell, Andrew Friedson, Kaj Gittings, Joni Hersch, Henry Hyatt, Lauren Jones, Josh Kinsler, Pat Kline, Francis Kramarz, Mark Kutzbach, Rick Mansfield, David Neumark, Jesse Rothstein, Meghan Skira, Isaac Sorkin, Chris Taber, and seminar participants at Boston College, the U.S. Census Bureau, EALE, LERA, Ohio State, NYU, UC Berkeley, 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

The theory of equalizing differences implies that workers receive higher wages for accepting jobs with undesirable characteristics. Rosen’s model (Rosen 1974) proposes an equilibrium in which worker and firm decisions lead to a sorting function that characterizes the effect on wages of taking a job with a particular set of amenities. Despite a vast empirical and theoretical literature, there is little consensus on whether job amenities have the direct, causal, effect on wages implied by Rosen’s model. Understanding the tradeoffs between wages and amenities can shed light on many fundamental topics in labor economics. Compensating wage differentials may be an important source of variation in wages between otherwise identical workers, which has implications for the interpretation of measured earnings inequality. Compensating wage differentials are also used frequently to guide the design of public policies.¹ Finally, the pervasive difficulty in measuring compensating wage differentials has been attributed to labor market frictions.² Resolving these difficulties may contribute broadly to our understanding of the competitive structure of the labor market.

A core estimation challenge in separating implicit amenity prices from other wage components is that workers are nonrandomly assigned to jobs, potentially in imperfectly competitive labor markets. Rosen provides conditions for recovering compensating wage differentials and worker preferences in a perfectly competitive labor market (Rosen 1974). The search literature (Hwang et al. 1998; Lang and Majumdar 2004) has shown that the hedonic wage regressions commonly estimated in the empirical literature yield compensating wage differentials that are highly sensitive to assumptions about labor market imperfections. To incorporate systematic wage dispersion for identical workers, empirical search models have replaced Rosen’s equilibrium construct of a hedonic pricing function with either an exogenous stochastic offer function (Bonhomme and Jolivet 2009) or bilateral bargaining model (Dey and Flinn 2005). Alternatively, more recent studies shift focus from estimating compensating wage differentials as the slope of the hedonic pricing function, and instead estimate the share of wage variation that can be explained by a set of unobserved firm-level amenities (Sullivan and To 2014; Sorkin 2018; Taber and Vejlín 2016).

In this paper we show that the causal effect of amenities on wages implied by Rosen’s hedonic pricing function exists and can be identified even under imperfect competition. Our approach differs from the rest of the empirical literature in its focus on the role of firms in determining the equilibrium relationship between wages and amenities when firms offer

¹For example, wage differentials for the risk of fatal injury are used to measure the value of statistical life, which affects tens of billions of dollars in federal spending annually on public safety policies in the US.

²See Hwang et al. (1998); Lang and Majumdar (2004); Bonhomme and Jolivet (2009).

differentiated jobs.³ Consistent with many studies, we find that a substantial part of variation in wages is attributable to firms (Abowd et al. 1999; Card et al. 2013; 2016; Abowd et al. 2012), and that high-wage firms tend to pay high wages in all occupations. Incorporating this firm-level wage coordination across different jobs in the Rosen model allows the hedonic pricing function to be reinserted as an equilibrium construct despite imperfect competition.

We estimate compensating wage differentials for occupational fatality risk in administrative matched employer-employee data from Brazil. These data, among the largest and most detailed ever used to study this topic, provide a complete census of work-related deaths, yielding more precise measures of fatality risk than have previously appeared in the literature (Kniesner et al. 2012). They also allow us to control for selection on the basis of unobserved worker and establishment-level heterogeneity in the determinants of pay, and to evaluate assumptions on the process by which workers are assigned to jobs.

Our analyses provide three primary contributions. First, we reconcile a contradiction between theoretical and empirical studies regarding the implications of unobserved worker ability. Hwang, Reed and Hubbard (1992) predict that high ability workers will sort into jobs with both higher pay and better amenities. As a result, omitting ability should induce a positive bias in cross-sectional estimates of compensating wage differentials for amenities (or negative bias in the case of disamenities).⁴ Panel estimates, which intend to correct for unobserved ability by using within-worker variation in amenities across jobs reach the opposite conclusion (Brown 1980; Kniesner, Viscusi, Woock and Ziliak 2012). We show that this model restricts the identifying variation in amenities to a component driven by workers' endogenous movements between firms, inducing a net *increase* in total bias in our data. This suggests that extending the hedonic wage model to account for ability is an insufficient improvement, and the remaining misspecification centers on the unmodeled wage impact of imperfect job mobility of workers across firms.

Second, we show that this misspecification is primarily caused by omitting the role of firms from the model of equilibrium wages. After correcting this problem using a hedonic version of the two-way fixed effects wage model introduced by Abowd, Kramarz and Margolis (1999), which allows for both unobserved worker and employer heterogeneity, we no longer reject that mobility across jobs with different amenities is conditionally exogenous. Quantitatively, this correction increases our estimates by an order of magnitude. The pattern of results across specifications is consistent with the bias in within-worker estimates arising from search frictions, as noted by Hwang et al. (1998) and Lang and Majumdar (2004). In their models,

³As noted above, other related studies have incorporated the role of firms, but these studies do not estimate information about hedonic pricing functions or the effects of observed amenities on wages.

⁴Technically, this implication follows from the assumption that the consumption of amenities does not decrease as income rises.

most job changes also entail a change in utility, as when climbing a frictional job ladder. The within-worker estimate is confounded by the equilibrium correlation between amenities and unobserved employer pay. Adding establishment effects allows workers' sorting and job mobility choices to be arbitrarily correlated with unobserved wage, and observed or unobserved amenity differences, between origin and destination establishments.

Third, we introduce a new theoretical model that integrates elements of the hedonic search framework of Hwang et al. (1998) with the differentiated firms model proposed by Card et al. (2018). The model clarifies the conditions under which our empirical wage model is equivalent to the structural wage equation under profit-maximizing equilibrium firm behavior. These conditions provide a mapping between exogeneity assumptions in the empirical wage equation and the interpretation of our empirical estimates relative to structural primitives. A key exogeneity condition is that additively separable occupation-group, worker, and employer effects capture all of the wage components associated with worker sorting that are correlated with job safety. If this condition holds, our estimates have a preference-based interpretation; if not, they identify the direct effect on wages of moving between jobs with different amenities. Importantly, we show that the wage model need not be perfectly specified to recover compensating wage differentials—there are forms of worker-firm match effects that do not violate conditional exogeneity in our model.

To evaluate the exogeneity conditions that affect the interpretation of estimates relative to this theoretical model, we conduct four sets of analyses. First, we apply tests of the additive separability specification suggested by Card, Heining and Kline (2013). We show that in Brazil, like Germany, there is little evidence against the assumption that job mobility is exogenous conditional on worker and establishment effects. Second, to address the concern that unobserved learning about ability or job fit may be correlated with wages and separations, we show that the model yields very similar estimates when variation is restricted to job separations initiated by mass displacement events. Third, to assess whether residual match effects impact job assignment through separation decisions, we follow the approach of Abraham and Farber (1987) and use completed tenure in uncensored job spells to proxy for unobserved match-specific wage effects. We find that this control has no impact on our estimates. Finally, to alleviate concern that match quality may affect workers' subsequent choices of amenities, we estimate an IV model using the network structure of the data to instrument for changes in fatality rates across jobs using former co-workers' job changes. We cannot reject that the IV estimates are identical to those from our benchmark specification.

Our analysis bridges the structural, theoretical, and reduced-form literatures on compensating differentials by focusing on the role played by firms in setting wages and amenities. The dominant approach in the reduced-form literature uses panel data to estimate hedonic

wage models in the presence of ability bias (Brown 1980; Hwang et al. 1992; Kniesner et al. 2012; Garen 1988). Largely due to data availability, these studies do not explicitly address the role of firms in setting compensation. More fundamentally, these analyses are based on crucial assumptions that job-to-job mobility is driven by changes in preferences. Our estimates are consistent with job mobility being driven by search associated with labor market imperfections.

Relative to the structural literature, we offer an alternative view on the role of firms that admits a different approach to identification, and is supported by the data. Much of the structural literature follows Hwang et al. (1998) in assuming firms and jobs are indistinguishable. Under this assumption, it is impossible to recover compensating wage differentials using data on wages and amenities alone. Instead, these studies incorporate information on job durations and obtain identification from revealed preference restrictions (Bonhomme and Jolivet 2009; Dey and Flinn 2005; 2008; Villanueva 2007; Sullivan and To 2014). Our approach instead is to use variation across jobs within firms, while modeling the restrictions on wages implied by firms maximizing profits under imperfect competition.

The difficulty in obtaining quasi-experimental assignment to jobs has led some researchers to favor an experimental approach to estimating worker preferences. Mas and Pallais (2017) use a field experiment and Wiswall and Zafar (2018) use a hypothetical-choice survey, to measure willingness-to-pay for job amenities based on stated preferences of respondents. While these approaches can reveal the full distribution of worker preferences, they target one object of interest. Without a demand-side they are not informative about the market equilibrium price for job amenities. The implications of compensating wage differentials for policy and welfare depend on whether estimated differentials reflect equilibrium prices, or capture only workers' preferences for job characteristics (Card et al. 2018; Lamadon et al. 2017). We view this paper as providing a complementary approach to estimate equilibrium compensating wage differentials in the presence of non-random assignment.

We are not the first to use matched employer-employee data to study compensating differentials. Lalive (2003) and Tsai et al. (2011) also estimate hedonic wage models using matched employer-employee data. However, their analysis is focused on aggregation bias in measured amenities, a shortcoming that is not present in our data. Hotz et al. (2017) use Swedish data to estimate worker valuation of employment in family-friendly firms. In their model, firms play a role in affecting a set of job amenities, but do not coordinate wages across jobs. Dale-Olsen (2006) uses matched data to estimate worker willingness to pay from data on job durations, following the revealed preference approach introduced by Gronberg and Reed (1994). Motivated by this work, we also consider duration models in Section 6. Consistent with our identifying assumptions, we find that within firms, differences in fatality

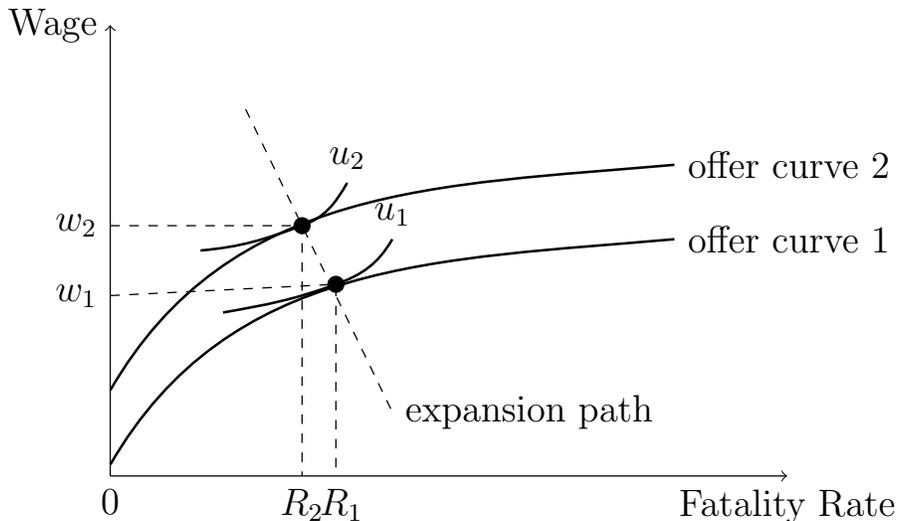


Figure 1: Equilibrium Wage-Risk Relationships

risk across jobs are not associated with the probability of voluntary separation.

2 Empirical Setting and Model

In this section we illustrate the challenges associated with unmeasured ability and nonrandom assignment by considering the Rosen equilibrium model through the lens of different data generating processes for wages and fatality risk. We then introduce our empirical model, and discuss how it addresses these estimation challenges.

2.1 Unobserved Ability

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). Consider Figure 1 as depicting a worker with preferences represented by indifference curve u_1 , who chooses a job characterized by a combination of wages and the risk of a fatal injury, (w_1, R_1) , to maximize utility along ‘offer curve 1.’ 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.⁵

The cross-sectional model is misspecified if workers differ in unmeasured ability. Suppose

⁵In 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.

the labor market supports two offer curves: ‘offer curve 1’ for workers of low ability and ‘offer curve 2’ for workers of high ability (Hwang et al. 1992). One can reinterpret Figure 1 as depicting low ability workers who choose (w_1, R_1) on indifference curve u_1 , and high ability workers who choose (w_2, R_2) on indifference curve u_2 . In this case, variation in wage-risk pairs has two sources—variation along each offer curve associated with differences in worker preferences, and variation along the expansion path, driven by differences in ability. If safety is a normal good, high ability workers trade off higher earning potential for reduced risk, causing the expansion path to slope downward even though indifference and offer curves slope upward. Variation along the offer curves 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.

Although this economic intuition is straightforward and well-known, it contradicts most empirical evidence. Brown (1980) and Kniesner et al. (2012) use panel data to estimate models of the form:

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

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. Interest centers on γ , which measures the relationship between wages and fatality rates. Estimates from models of this form tend to imply smaller (frequently zero) compensating wage differentials relative to cross-sectional estimates in the same empirical setting, contrary to theory. 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 ten.⁶

2.2 Endogenous Job Mobility

The theory of hedonic search provides a simple, but powerful, framework for reasoning about the counterintuitive estimates of compensating differentials in panel data models. In an imperfectly competitive labor market, workers’ decisions to change jobs may be associated with movements between firms offering different levels of total compensation (or utility). To see this, one can reinterpret Figure 1 as depicting the data generating process that arises from a single worker choosing to move 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 is costly for workers to

⁶The simulated bias is negative for disamenities and positive for amenities

find jobs offering higher utility as in Hwang et al. (1998) or Lang and Majumdar (2004), 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).

From this perspective, the degree of bias in the worker-effects specification depends in part on how much of the within-worker variation in wages and safety is associated with variation along an offer curve, as opposed to across offer curves as workers ascend the job ladder. If workers have heterogeneous preferences, and sort along offer curves but change jobs along expansion paths, then introducing θ_i in the model, in an attempt to reduce ability bias, may instead isolate the most problematic component of variation, potentially increasing total bias relative to a pooled OLS model. Note that even if safety is a normal good, the sign of the expansion path slope in this case is ambiguous—it can depend on how firm compensation levels impact the sorting of workers across firms, or whether high-compensation firms have a relative advantage in providing safety.⁷

There is good reason to believe that sorting across firms with different compensation practices is an important feature of labor markets. Abowd, McKinney and Schmutte (2017) show 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. This suggests that movements similar to that depicted by (w_1, R_1) and (w_2, R_2) are common. If job changes involve a simultaneous increase in the wage and decrease in the fatality rate, the compensating wage differential implied by Equation 1 is incorrect.

A natural progression to estimating compensating wage differentials in matched employer-employee data is to introduce amenities into the AKM wage model (Abowd et al. 1999):

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

The subscript $J(i, t)$ is an index function mapping to the establishment j employing worker i in period t . The inclusion of $\Psi_{J(i,t)}$ allows for variation in pay across jobs in the same establishment and accommodates arbitrary sorting on the basis of risk and the worker effect. The term $\Psi_{J(i,t)}$ also absorbs the effect on wages of all unobserved firm-level amenities. This model allows for the possibility that workers are nonrandomly assigned to firms that offer unobserved wage premia, or differ in unobserved amenities that are common across jobs at the firm. If the 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

⁷See Section 5.5 and Appendix C.7 for more on this point.

fixed.

In principle, the conditional exogeneity assumptions in this model could be relaxed further by including a job match effect, as in Lavetti (2018). However, in our empirical setting only 3% of the total variation in fatality rates occurs over time within jobs, suggesting that a match effect would absorb nearly all of the variation in $R_{c(i,t),t}$. We find that the remaining, small intertemporal changes in fatality rates are uncorrelated with wage changes. The objective of our empirical model is to instead identify γ using across-job variation in $R_{c(i,t),t}$, while controlling for the factors that shift the offer curve.

2.3 Orthogonal Match Effects (OME) Model

Our benchmark model is a variation of the AKM wage model, which we estimate in two steps. First, we project wages onto observed time-varying controls (experience effects, year effects, and fatality risk), and an unrestricted worker-occupation-establishment match effect.

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

$\Phi_{i,Jk(i,t)}$ denotes the match effect between worker i and the Jk establishment-occupation pair at which worker i is employed in period t . Next, we remove the effect of $x_{it}\hat{\beta}$ from log wages and estimate the AKM model on the transformed dependent variable:

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

where $P_{it} \equiv w_{it} - x_{it}\hat{\beta}$.⁸ Our preferred specification includes fatality risk, 1-digit occupation effects $\pi_{k(i,t)}$, year effects τ_t , in addition to worker and establishment effects. We show in Appendix Tables A.1 and A.2, that our findings are qualitatively and quantitatively robust to a wide range of alternative specifications, including interacting experience and race or education, controlling for tenure, year of hire effects, excluding occupation controls, and estimating each of these specifications using a single equation AKM model.⁹

⁸Woodcock (2008) and Barth, Bryson, Davis and Freeman (2016) adopt a similar approach to estimation in settings where interest centers on a match-specific characteristic, like fatality risk. Doing so has two advantages. First, the experience and year effects estimated in the first stage are purged of any correlation with arbitrary match effects. Second, the coefficient on risk estimated in equation (3) is identified from within-match variation in fatality risk, which accounts for only 3% of total variation. 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.

⁹The 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.

2.4 Key Assumptions and Their Empirical Counterparts

If adding worker effects to the wage equation potentially increases bias, it is natural to ask why introducing establishment effects improves the model. The answer hinges on the empirical data generating process for wages and fatality rates, and on the statistical object of interest. The typical object of interest in hedonic wage models is the *hedonic pricing function*, which measures the expected change in wages for a given change in amenities, holding all other job characteristics fixed. Identifying the slope of the hedonic pricing function using Equations 1 or 4 requires that there is no component of the wage residual that is associated with job mobility, which could potentially be correlated with changes in amenities across jobs. This condition is related to the exogenous mobility assumption required for identification of establishment effects in the conventional AKM model; the validity of this assumption is an empirical question.

We show that our reduced-form analysis based on the hedonic AKM model provides a striking amount of economic insight given its simplicity. If most of the sorting associated with the utility ladder takes place across firms, then $\hat{\gamma}$ should be larger in the OME model than in the within-worker model, which is exactly what we find. However, it remains to show empirically that the movement of workers across jobs with different fatality rates is independent of unmodeled wage innovations. We devote considerable attention to diagnosing whether this assumption is consistent with the data generating process in Brazil. An important tradeoff is that weakening the exogenous mobility assumption generally leaves less variation with which to identify the implicit price of amenities.

For purposes of exposition, it is also worth clarifying the distinction between several statistical and theoretical objects of potential interest. Our primary object of interest is the *hedonic pricing function*, which is identified under the conditionally exogenous mobility assumption. In the context of fatality risk, the slope of the hedonic pricing function is the key parameter for inferring the value of statistical life (Rosen 1988), though this interpretation generally relies on some additional assumptions about preferences. Other theoretical objects of interest in the literature include workers' *willingness to pay for (accept)* amenities (disamenities), or firms' iso-profit functions. These objects are model-based concepts that generally differ from the hedonic pricing function. After discussing our empirical estimates, we return in Section 5 to present a theoretical model that clarifies how these objects of interest relate to the exogeneity conditions in the OME model.

We note that one limitation of our setting is that we do not observe a large set of non-wage amenities. However, since the model controls for unobserved establishment and occupation effects, the only amenities that may affect the interpretation of γ are those that vary across jobs within establishments and occupation groups. The presence of such amenities, if they are

correlated with fatality rates, means that when estimating the model using our administrative data from Brazil, γ identifies the variation in wages associated with the bundle of amenities correlated with risk. Although our objective in this paper relates primarily to the conceptual framework and methodology for estimating compensating wage differentials, which can be adapted to consider a vector of amenities, the quantitative application of the model is still informative of the expected change in wages associated with a change in fatal risk in Brazil.

3 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.

3.1 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*.¹⁰ For each job, in each year, the employer reports characteristics of the worker, the job, and the establishment. Worker characteristics include gender, race, age, and educational attainment.¹¹ Job characteristics relevant to this study include the monthly wage, weekly contracted hours, occupation, and the cause of job separations (which includes whether the job ended because of a fatal work-related injury.) The establishment characteristics include the establishment’s industry, location, and number of employees.¹²

¹⁰For details on labor market formality and wage setting institutions, see Appendix C.1

¹¹Because individual characteristics are reported by the employer, they can change as workers move from job to job. Cornwell, Rivera and Schmutte (2016) 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.

¹²That 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.

3.2 Measuring Fatality Rates

The RAIS data also function as a census of fatal occupational injuries. 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.3). 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 2007.¹³ See Appendix C.3 for details of this calculation. For comparability with previous literature, and to smooth out fluctuations in annual rates, we construct a three-year moving average fatality rate for each industry-occupation pair (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 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. Although many interesting questions remain about the measurement of occupational safety, we leave these questions (some of which we consider in Lavetti and Schmutte (2017)) for future research.

In Appendix Table A.4, we report average fatality rates by aggregate industry and occupation as evidence that our measurements of fatality risk are sound. The overall fatality rate is 0.049 deaths 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.

¹³See <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.

3.3 Analysis Sample and Variable Definitions

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 65.¹⁴ 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.¹⁵ We further restrict the analysis sample 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. Finally, we Winsorize the data at the 1st and 99th percentiles of the log wage distribution. These restrictions yield an analysis sample with about 83 million job-years.

The RAIS data report average monthly earnings. If a worker holds a job for less than 12 months during the year, the variable reported by RAIS represents one month’s pay. In practice, this variable measures the monthly wage rate, which is a common institutional arrangement in Brazil. For consistency with prior research, we convert monthly earnings to an hourly wage rate measured in 2003 Brazilian reais.¹⁶

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 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.

¹⁴See Lavetti and Schmutte (2017) for a detailed discussion of the gender-specific aspects of this topic.

¹⁵We define expected earnings as the product of the average monthly wage rate with the number of months the worker was employed.

¹⁶First 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. In 2010, one U.S. dollar was worth 1.66 Brazilian reais. Hence, one can approximately interpret our estimates as 2010 dollars.

Table 1: Summary Statistics

	Population	Analysis Sample
Age	36.98	36.23
Race <i>branco</i> (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 (Reais)	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.

4 Estimates and Specification Diagnostics

We first present estimates from Equations 1, 3, and 4. We then present a variety of diagnostic evidence to clarify the nature of the endogeneity bias in Equation 1 and connect the empirical data generating process in Brazil to the conceptual framework discussed in Section 2. Since our fatality rates are measured as industry-occupation averages, we conduct complementary analyses to show the identifying aggregate variation at the industry-occupation level, and assess the stability of estimates to different types of job changes, including changes in occupations within establishments, and job changes within occupations across industries.

4.1 Benchmark Results

Table 2 compares estimates of the compensating wage differentials from each of the empirical models discussed in Section 2. Column (1) reports estimates from the pooled cross-sectional model, $w_{it} = x_{it}\beta + \gamma R_{c(i,t),t} + \epsilon_{it}$. The estimated $\hat{\gamma} = 0.279$ 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 wages. In estimating this model, our control variables include dummies for each year (up to 30) of labor market experience, worker race,

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

	(1)	(2)	(3)	(4)
	Pooled	Worker Effects	Match Effects	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.

education, plant size, year, state of employment, 1-digit industry and 1-digit occupation effects.¹⁷ 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].¹⁸

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. This attenuation of the within-worker estimate relative to the cross-sectional estimate is consistent with evidence from US data (see in particular Brown (1980), Kniesner et al. (2012), and Lavetti (2018)). The estimated effect of being employed at a job with zero fatalities also declines by 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 Equation (3), the match effects specification. In this model γ is identified only from the time-series variation in fatality risk within jobs, which includes only 3% of the total variation in fatality rates. The estimated compensating wage differential is -0.006 , which suggests that wages do not vary in an economically meaningful way in response to the relatively small changes in risk within jobs.

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, $\hat{\gamma} = 0.170$ (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

¹⁷Our 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. (2018) illustrate, person effects are not identified relative to year and experience effects without some normalization. As we show in Table A.1, our main results are not sensitive to alternative specifications of the experience profile.

¹⁸Following 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, $\frac{1}{w} \frac{\partial w}{\partial a} = \hat{\gamma}$, so $VSL = \bar{w} \hat{\gamma} * 2,000,000$. 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 known population statistic.

with the type of bias that is predicted to arise theoretically in a model of hedonic search.

We evaluate a wide range of alternative specifications in Appendix Tables A.1 and A.2, including interactions between experience and race or education, controlling for tenure, year of hire effects, excluding occupation controls, and estimating each of these specifications using a single equation AKM model. The estimated values of $\hat{\gamma}$ are quite stable across these alternative model choices, ranging from 0.152 to 0.190. Table A.5 shows these results are robust to dropping industry and occupation controls. Table A.6 shows the same bias pattern holds when we estimate the model separately by region. We also present results from a collapsed aggregate specification in Appendix Section B.

4.2 The Worker-Effects Model is Misspecified

The results in Table 2 suggest that the worker effects model is biased due to the omission of firm characteristics that affect mobility and wages. Figure 2a confirms this intuition. The figure shows that, when workers switch jobs, changes in wage residuals from the within-worker model (Column 2 of Table 2) are correlated with changes in fatality risk. Note, first, that 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 shows that when a job change involves a large decrease in risk, there is a larger increase in the wage residual. This finding is consistent with a model in which job changes involve movements to jobs that are more attractive on both wage and safety dimensions. The data clearly contradict the exogenous mobility assumption required for the worker effects model to be identified.

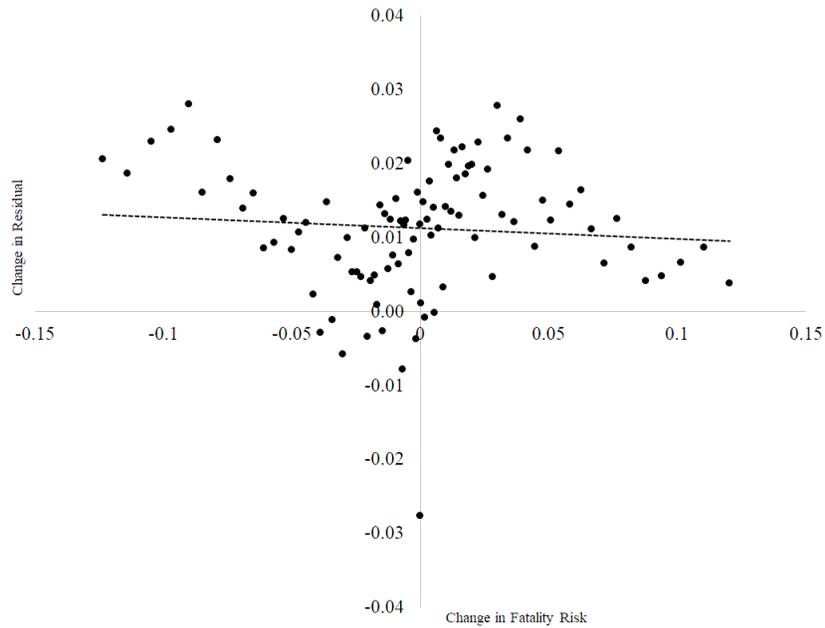
4.3 OME Model Diagnostics

Figure 2b presents the same diagnostic using residuals from the OME model. The inclusion of establishment effects in the model causes the average change in residuals associated with a job change to be centered at zero. In addition, there is no strong systematic relationship between the change in risk and average change in residuals. This diagnostic is suggestive that there is limited scope for potential endogenous mobility bias in the OME model.

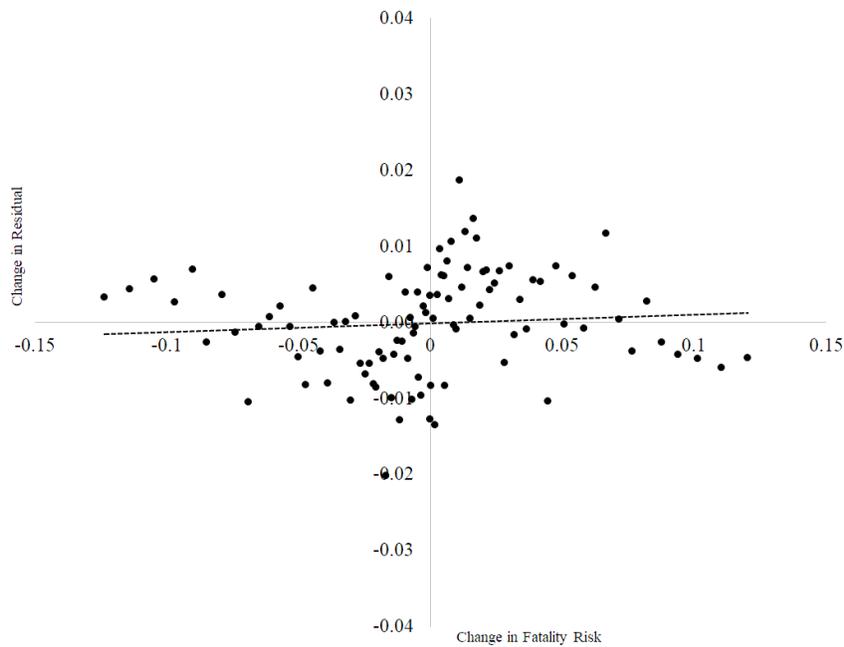
Table 3 presents the estimated components from the OME model, as well as correlations between components. The table shows that fatality risk exhibits a significant negative correlation with the estimated worker (-0.09) and establishment (-0.11) effects. These correlations support the conclusion that omitting worker or establishment effects is likely to cause endogeneity bias in a hedonic wage model. Table 3 also reports the contribution of each component to the total variation in wages. The compensating differential for fatal risk

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.

Table 3: 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
Average Establishment Size	17.4
Number of Workers in Mover Sample	19,646,048
Average Number of Jobs per Worker	1.9

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.

contributes a relatively small share of the variation, representing about 2% of the standard deviation of wages. The estimated worker and establishment effects explain 70% and 46%, respectively, of the variation in wages. These shares are approximately the same as Card et al. (2013) and Alvarez et al. (2018) find in recent years of West German and Brazilian data respectively.

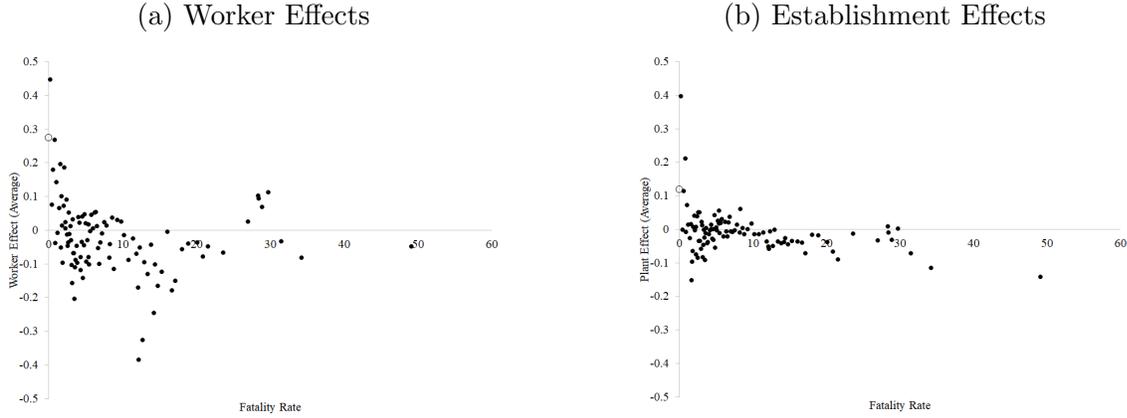
4.4 Bias Decomposition

To quantify the sources of relative bias, we decompose the raw relationship between fatality risk and wages into the sum of three components:

$$\begin{aligned}
 \hat{\gamma}^{raw} &= \underbrace{\hat{\gamma}^{OME}}_{\text{OME estimate}} + \underbrace{\frac{\text{cov}(\theta, R)}{\text{var}(R)}}_{\text{bias from worker eff.}} + \underbrace{\frac{\text{cov}(\psi, R)}{\text{var}(R)}}_{\text{bias from estab. eff.}} + \underbrace{\sum_k \frac{\text{cov}(x_k, R)}{\text{var} R}}_{\text{bias from controls}} \\
 -0.181 &= 0.170 \quad -0.212 \quad -0.272 \quad +0.134
 \end{aligned}$$

As theory suggests, the omission of worker and establishment effects both contribute strong negative bias of similar magnitude. Neglecting to account for sorting on observed control variables induces a partially offsetting positive bias, potentially arising from more experienced workers being employed in riskier jobs.

Figure 3: Binned Scatterplots of Worker and Establishment Effects versus Fatality Rates



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.

To conduct this decomposition, we follow the approach of Gelbach (2014) and Abowd et al. (1999). First, we obtain the estimate $\hat{\gamma}^{raw} = -0.181$ from a model regressing log wage onto fatality risk, a dummy for jobs with zero fatality risk, and a set of year dummies. We estimate the relative bias components by projecting the estimated worker and firm effects from the OME model onto the same set of variables, and likewise for the additional controls in the OME model (experience and occupation effects). Additional details on the decomposition are included in Appendix C.5.

To visualize the nature of ability and employer sorting bias, Figures 3a and 3b present binned scatterplots of estimated worker and establishment effects, respectively, against fatality risk. Both figures show that the highest-paying establishments and highest-earning workers tend to be concentrated in very safe jobs. This pattern is consistent with economic theory—as the expansion path in Figure 1 approaches the vertical axis, a worker with sufficiently high ability, or firm with high labor demand, may choose a level of risk at the corner solution, with zero probability of death. As a result, omitting worker or establishment effects from the model leads to a large positive wage residual in jobs with low fatality rates. Similar intuition is formalized by Caetano (2015), who shows that a discontinuity at the corner of a function that should theoretically be smooth can be used to construct a test of model mis-specification.¹⁹

Among jobs with fatality rates away from zero, the distribution of estimated establish-

¹⁹Because the risk of a fatal accident 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; it suggests that the presence of endogeneity may be manifest as a non-monotonicity in the wage-risk profile in the vicinity near zero fatalities.

ment effects is less strongly correlated with fatality rates.²⁰ This pattern is consistent with a model in which most firms offer jobs with a wide range of fatality rates, and compensation for fatality risk is not establishment-specific.

4.5 Identifying Variation

In this section, we clarify the sources of identifying variation in the OME model in two ways. First, we present estimates from a version of the OME model that allows γ to differ for job changes within establishments, within occupations, and within industries. Isolating portions of the identifying variation, the results suggest that all three types of job changes contribute similarly to the average γ estimated in the benchmark OME model. Second, we graphically depict the identifying variation by showing that, across industry-occupation cells, average differences-in-differences of establishment-occupation effects are correlated with differences-in-differences in fatality rates.

4.5.1 Heterogeneity in γ by Type of Job Change

The OME model identifies compensating differentials from movements of workers across jobs in different industries, occupations, and establishments. If the model is correctly specified, then the estimated effect should be the same regardless of the type of job-to-job move. Table 4 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 similar for job changes within or across occupations, and for job changes within or across establishments. $\hat{\gamma}$ is slightly larger when estimated using variation within occupations. Specifically, the within-occupation and across-industry estimate (0.191) exceeds the across-occupation and within-establishment estimate (0.159) and the across-occupation, within-industry, across-establishment estimate (0.148).

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.²¹

²⁰Doubling the mean fatality rate is associated with about a 0.05 standard deviation decrease in Ψ .

²¹Of 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,

Table 4: Sensitivity of γ to Type of Job Change

	(1)	(2)
Fatality Rate	0.178*	0.190*
	(0.001)	(0.001)
Fatality Rate*Within Occupation	-0.006*	0.001
	(0.001)	(0.001)
Fatality Rate*Within Establishment	-0.013*	0.011*
	(0.001)	(0.001)
Fatality Rate*Within Industry		-0.042*
		(0.001)
N	83,418,032	83,418,032
R-Sq	0.930	0.930

Notes: Estimates are from the OME specification. ‘Within Occupation’ and ‘Within Industry’ equal 1 for all years of an origin and destination job that have the same 3-digit occupation code or 2-digit industry code, respectively. ‘Within Establishment’ equals 1 if a worker changes occupation within an establishment, and equals 0 if the job change is across establishments. ‘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.

4.5.2 Analysis of Industry-Occupation Aggregate Measures of Fatality Risk

While our measures of fatality rates are more detailed than those used in previous work, in this section we provide evidence that using industry-occupation aggregate measures yields similar conclusions. To do this, we construct an industry-occupation level dataset whose entries, $(\bar{R}_{k,n}, \bar{\psi}_{k,n})$, 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. We calculate a difference-in-differences measure that captures 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' :

$$(\bar{R}_{k,n} - \bar{R}_{k',n}) - (\bar{R}_{k,n'} - \bar{R}_{k',n'}) \quad (5)$$

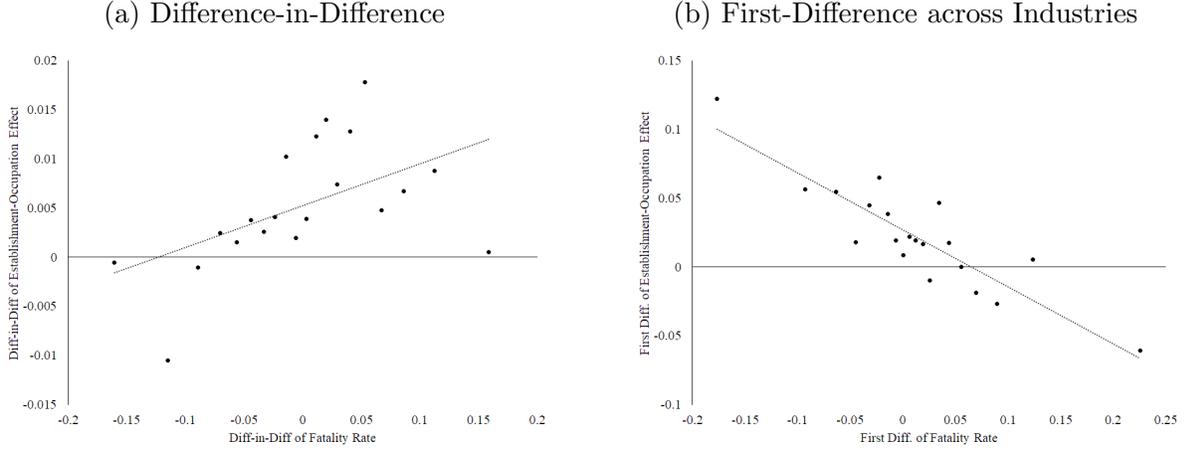
We then repeat a similar calculation using establishment effects,

$$(\bar{\psi}_{k,n} - \bar{\psi}_{k',n}) - (\bar{\psi}_{k,n'} - \bar{\psi}_{k',n'}), \quad (6)$$

that measures the change in average establishment-occupation wage effects associated with moving from a job in occupation k to a job in occupation k' when that job is offered in

our estimate of the CWD using within-match variation is effectively zero, consistent with a job search model without renegotiation.

Figure 4: Identifying variation: The Relationship Between Wages and Fatality Risk



Notes: Figure 4a shows a binned scatterplot of the average difference-in-difference of establishment-specific occupation effects described in (6) against the difference-in-difference of fatality risk in (5). Figure 4b 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 (5).

industry n relative to industry n' .

Figure 4a presents a binned scatterplot of pairwise differences-in-differences in average establishment-occupation wage effects versus differences-in-differences in average fatality risk. The figure shows that if moving between two occupations is associated with a greater increase in risk in one industry than the same occupation change would cause in another industry, workers are clearly compensated for this additional change in fatality risk. Since the vertical axis is constructed using establishment-occupation wage effects, the entire compensating wage differential for fatal risk must be absorbed in this wage component by construction. The positive slope in Figure 4a indicates that this unobserved wage component containing the CWD increases as excess fatality risk increases.

It is instructive to compare Figure 4a with its first-differenced counterpart. Figure 4b is a binned scatterplot of the difference in average establishment-occupation effects across occupations $(\hat{\psi}_{k,n} - \hat{\psi}_{k',n})$ for each industry against the difference in risk across occupations. In contrast to the difference-in-difference plot, Figure 4b shows that, within industries, excess fatality risk is associated with lower wages. This is consistent with establishment wage effects differing on average across industries, creating an aggregate form of endogenous mobility bias comparable to the problem discussed in Section 2, which negatively biases estimates. Indeed, labor economists documented and studied industry-level violations of the law of one wage (Krueger and Summers 1988) long before matched longitudinal data made possible the more refined analysis of establishment-level earnings effects.

Appendix B also presents a multi-level version of the OME model that aggregates wage

outcomes to the same level of variation as fatality risk. In this model, we first estimate an industry-occupation-year effect, and then regress the estimated wage component on fatality risk. The estimated γ is slightly larger, though not statistically significantly different, than our OME estimate.²²

5 A Search Model with Imperfect Competition and Amenities

Although our estimates show the model of equilibrium wage determination is improved by accounting for firm compensation heterogeneity, the interpretation of parameters from this model may not be obvious. In this section 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 derive an equilibrium expression for log wages that can relate the structural primitives in Rosen’s model (the hedonic pricing function, the marginal willingness to accept fatal risk, and the slope of firms’ isoprofit functions with respect to risk) to parameters in our empirical wage model. The solution to the model relates worker preferences to the derivative of wages with respect to fatal risk, and characterizes a match-specific error term that explicates threats to identification and serves as a guide for further empirical diagnostic analyses in Section 6 that clarify the assumptions behind potential interpretations of our estimates.

Our model extends the “differentiated firms” framework of Card et al. (2018) 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 $i \in \{1, \dots, N\}$ 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, $s(i) \in \{1, \dots, S\}$. Workers receive offers with the same probability regardless of whether and where they are employed, and offers expire at the end of each period. Therefore, workers always choose whichever job offer provides the highest instantaneous utility. Worker i , when employed in occupation k by firm j in period t 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 $j \in \{1, \dots, J\}$. Firms are distinguished by industry, $b(j) \in \{1, \dots, B\}$, 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, $k \in \{1, \dots, K\}$.

²²One drawback of this approach is that some industry-level differences in average establishment effects are included in the industry-occupation-year effect, so an endogenous mobility problem is still present, and is more challenging to address in this model specification.

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

$$L_{sjkt} [Q_{sjkt} - C_{bk}(w_{sjkt}, R_{jkt})] \quad (7)$$

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 .

5.1 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:

$$p_{sjkt} = K_s \exp(\bar{u}_{sjkt}), \quad (8)$$

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 C.6.

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 (8) for p , steady-state employment as a function of the utility offer is:

$$H(\bar{u}) = \frac{\lambda K \exp(\bar{u}) N}{\Omega(\bar{u})} \quad (9)$$

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 (9) simplifies to the static labor supply equation in Card et al. (2018), but extended to include the endogenous amenity R . This is a useful benchmark case to which we return below.

5.2 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:²³

$$\max_{w, R} [Q - C(w, R)] H(\bar{u}). \quad (10)$$

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

$$[Q - C(w, R)] \frac{\partial H(\bar{u})}{\partial w} = C_w(w, R) H(\bar{u}). \quad (11)$$

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

$$\frac{\partial H}{\partial w} = f_w(w, R) \left(\frac{H(\bar{u})}{\Omega(\bar{u})} \right). \quad (12)$$

²³We abstract from initial conditions and transition dynamics. Once in steady-state, the present value of discounted expected profits is equal to the expected per-period steady-state profit divided by the interest rate, so the two objective functions yield the same choices.

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:

$$\frac{f_w(w, R)}{f_R(w, R)} = \frac{C_w(w, R)}{C_R(w, R)}. \quad (13)$$

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.

5.3 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 that indirect utility is additively separable in (log) wages²⁴ and risk, given by $f(w, R) = \ln w - h(R)$, and that the logarithm of unit labor costs are given by $\ln C(w, R) = \ln w - y_{bk}(R)$.²⁵ This unit labor cost function has the property that the marginal cost of increasing safety is greater for high-wage workers.²⁶ Note the bk subscript on y , which highlights that the cost of providing safety may depend on the industry and occupation of the job. 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 (13), 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:

$$\ln w = \ln T_j + \ln \theta_s + \ln \pi_k + y_{bk}(R) + \ln \left(\frac{1}{1 + \Omega(\bar{u})} \right). \quad (14)$$

Differentiating Equation (14) with respect to R and substituting the equilibrium condition $y'(R_{jk}) = h'(R_{jk})$ gives a structural equation for the unconditional relationship between

²⁴Note that Hwang et al. (1998) consider linear wage regressions and assume a corresponding indirect utility function that is linear, instead of logarithmic, in wages. This eliminates the possibility of an income effect.

²⁵This specification of preferences also generalizes Card et al. (2018). 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.

²⁶We use this feature of the cost function to incorporate tradeoffs faced by firms between the speed at which workers perform tasks and worker safety. An alternative approach is to incorporate this tradeoff through the production function.

wages and risk:

$$\frac{d \ln w}{dR} = h'(R) \left[1 - \left(\frac{1 - \Omega(\bar{u})}{1 + \Omega(\bar{u})} \right) \right]. \quad (15)$$

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 $\Omega(\bar{u})$.

5.4 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 from the OME model, $\hat{\gamma}$, and the theoretical relationship between wages and risk. Compared to Equation (14), the OME model omits the incumbency advantage term. Equation (15) suggests that the estimate from the linear conditional expectation function will be a weighted average of downward-biased estimates of worker preferences for safety.

Given this theoretical source of potential bias, what is the appropriate interpretation of $\hat{\gamma}$? It is instructive to first consider the case where all workers receive offers from all jobs in every period—that is, where $\lambda = 1$. In this case, there is no incumbency advantage, and the unmodeled utility component in Equation (14) 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 conditionally exogenous with respect to wages in the OME model. This result is noteworthy because it shows that the classic Rosen (1974) model can be extended to accommodate imperfect competition in labor markets.

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 may be correlated with risk, even after conditioning on firm and worker effects.

In this case, the compensating wage differential $\hat{\gamma} = \frac{\partial \mathbb{E}[\ln w | x, \theta, \Psi]}{\partial R}$ is the estimated treatment effect on log wages of changing fatality rates, holding fixed covariates. The interpre-

tation of estimates relative to preferences changes, because $\frac{\partial \ln w}{\partial R}$, the total derivative of log wages with respect to risk, 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. 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 $h'(R)$. Just as in Hwang et al. (1998), the 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 of either the treatment effect on wages or of preferences is still possible when amenities vary across jobs in the same firm.

From Equation 4,

$$\frac{\partial \mathbb{E}[\ln w | x, \theta, \Psi]}{\partial R} = \hat{\gamma} + \mathbb{E} \left[\frac{\partial \ln w}{\partial \xi} \frac{\partial \xi}{\partial R} \middle| x, \theta, \Psi \right] \quad (16)$$

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

$$\frac{\partial \ln w}{\partial R} = \frac{\partial \mathbb{E}[\ln w | x, \theta, \Psi]}{\partial R} \Rightarrow \mathbb{E} \left[\frac{\partial \ln w}{\partial \xi} \frac{\partial \xi}{\partial R} \middle| x, \theta, \Psi \right] = 0 \quad (17)$$

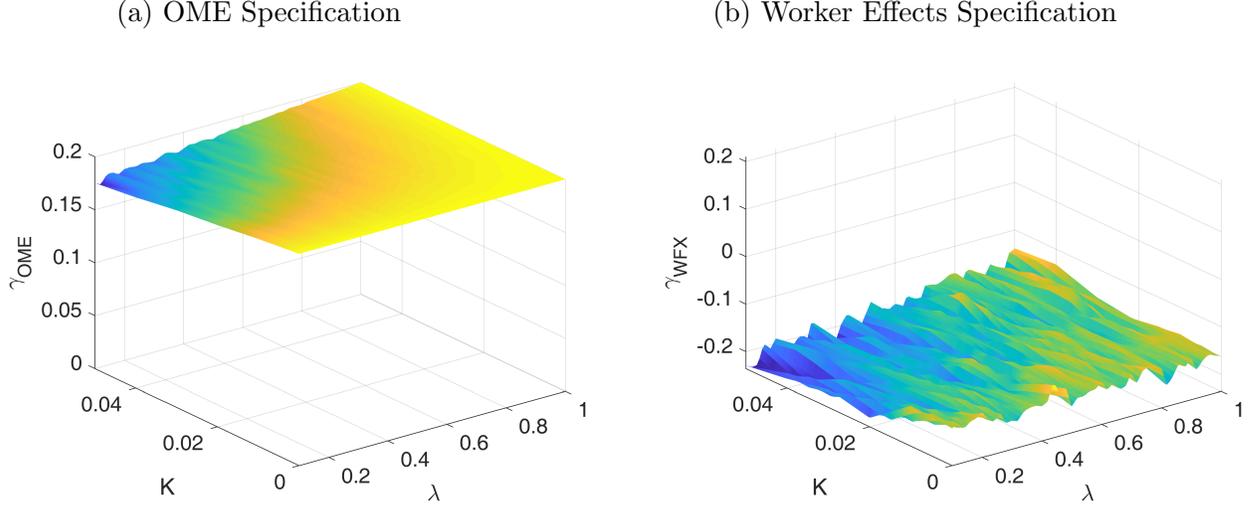
which holds if the included variables in the OME model control for $\ln \left(\frac{1}{1+\Omega(\bar{u})} \right)$. In practice, the relationship between γ and $h'(R)$ therefore depends on whether the additive separability specification in the OME model is valid, or whether there is important residual 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. In the next section we focus attention on evaluating this residual match-specific wage component. 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 8 that controlling for completed tenure has no effect on the estimated compensating differential.

5.5 Monte Carlo Simulation

We estimate a Monte Carlo simulation of the model, and then attempt to recover the true marginal willingness to pay for safety from the simulated data using the OME and worker

Figure 5: Monte Carlo Estimates of $\hat{\gamma}$ when True $\gamma = 0.2$



Notes: Estimates are based on 25000 simulated workers over 30 periods for each (λ, K) pair. See Appendix for additional simulation details.

effects specifications. For simplicity, we assume homogeneous firm technology that satisfies the first order conditions and the profit maximizing log wage equation (14). In the simulation, workers receive three outside offers and an inside offer in each period. Wage offers follow the equilibrium wage equation where (Ψ, R) are jointly normally distributed with the empirical means, variances, and covariance. θ is normally distributed with mean zero the empirical standard deviation (0.456).

Figure 5a depicts estimates of $\hat{\gamma}$ from the OME model as λ and K vary. As we explain above, when $\lambda = 1$ there is no incumbent advantage, and the bias in the OME specification is exactly zero. Similarly, when K approaches zero, firms expect the probability of their offer being accepted to decrease towards zero. As the labor market approaches this competitive case, the classic Rosen result holds, and the OME specification again has zero bias. For a wide range of λ , the OME model yields relatively stable estimates of $\hat{\gamma}$ when the number of firms is large (and therefore K is small). The correlation between included variables and the error component in Equation 14 $\left[\ln \left(\frac{1}{1+\Omega(\bar{u})} \right) \right]$ biases $\hat{\gamma}$ modestly downward from the true value of 0.2 when $\lambda < 1$ and $K > 0$.

In stark contrast to the OME estimates, Figure 5b depicts estimates from the worker effects specification based on the same simulated data. The estimates are biased downward in excess of 150% of the true γ over the full range of λ and K . When $1 \geq \lambda \geq 0.5$ and $0.05 \geq K \geq 0$, the maximum relative bias in the OME specification is 7.3%, compared

to 202.8% for the worker effects model. Additional details on the Monte Carlo model are provided in the appendix.

Notice that in our model, even though firms' choices of R are independent of T_j , omitting firm effects from the wage equation yields a biased estimate $\hat{\gamma}$. This bias is caused by nonrandom selection into jobs. Workers' decisions to switch jobs when offered an increase in utility creates a relationship between observed changes in R and changes in T_j . For example, an observed increase in R can only be rationalized by an increase in T_j , whereas a decrease in R could be rationalized by either an increase or decrease in T_j . Therefore, for a given worker, ΔR is on average correlated with ΔT_j , even if the offer function does not impose a correlation. This can be seen directly by estimating the Monte Carlo model while forcing the covariance between R and Ψ in the offer function to be zero. As shown in the appendix, the bias in the worker effects model remains several times larger than the OME bias in this case.²⁷

6 Evaluation of Model Restrictions

Our theoretical model implies that the interpretation of $\hat{\gamma}$ from the OME wage model depends on whether a residual match effect in the wage model is correlated with amenities. In this section we conduct a set of tests to assess whether such a match effect exists. We begin with diagnostic tests introduced by Card et al. (2013), which support the conditional exogeneity assumption and suggest that, although worker-establishment match effects contribute modestly to the aggregate variation in log wages, there is little evidence that match effects are correlated with workers' job mobility decisions, in which case the exogeneity condition in Equation 17 is satisfied. We then re-estimate each of the benchmark wage specifications but allow the compensating differential to differ for job-to-job changes initiated by a mass displacement event. The intuition is that mass displacements are less likely to be contaminated by any one worker's idiosyncratic match-specific wage effect. We then estimate the models including a control for completed job tenure (in the non-censored job spell sample) as a proxy for worker retention probabilities. Finally, we propose an IV model that uses the network structure of the data to instrument for changes in fatality rates across jobs using former co-workers' subsequent job changes. The intuition is that one's coworkers are likely to face similar job consideration sets, but any changes in amenities are unlikely to be correlated with the focal worker's own idiosyncratic match effect, which is mean zero by construction for each establishment, each occupation, and each worker. IV estimates are very similar to estimates from the OME model.

²⁷Therefore alternative model assumptions that impose a connection between T_j and R would only reinforce this need to control for T_j in the wage equation.

We also present two additional analyses that are informative of the sensitivity of our estimates to other model assumptions. To evaluate the linearity assumption, we present semi-parametric OME estimates that impose fewer restrictions on $\frac{\partial \ln w}{\partial R}$. We also extend the Gronberg and Reed (1994) framework for estimating the willingness to accept fatal risk using job separation choices, and show that omitting employer effects can significantly inflate estimates. This result helps clarify a pervasive discrepancy in the literature between estimates from job separation models and hedonic wage models.

6.1 Do Match Effects Matter?

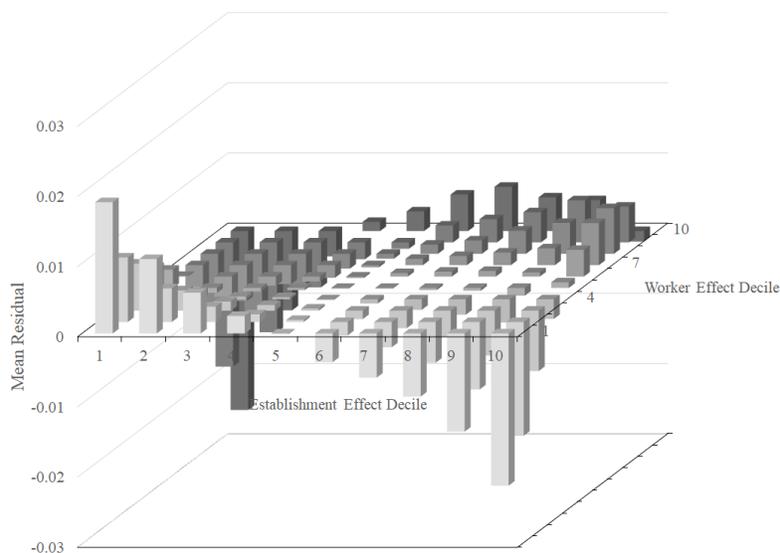
In our analysis sample, over 97% of the total variance in log wages occurs across jobs. Of this variation, 95% can be explained by a two-way fixed effects model with worker and establishment effects alone. These facts suggest that any residual unexplained wage variation is extremely small. A decomposition 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. The combination of these patterns is suggestive that, although there is variation in wages across occupations within establishments, the variation looks quite different than a systematic wage premium.

In this section, we evaluate the separability and exogenous mobility assumptions of the OME model more thoroughly by applying diagnostic tools developed by Card et al. (2013).

Figure 6 displays the mean residual within cells defined by deciles of the estimated worker and establishment effects from our 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 log points 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 6.2, 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.

Next, if match effects play an important role in job assignment, one would expect to see 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.

Figure 6: 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.

Table 5 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 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 log points. 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.

6.2 Sensitivity to Violations of the Separability Assumption

Figure 6 is generally supportive of the additive separability of log wages in unobserved worker and firm heterogeneity. However, to the extent that mean log wage residuals differ from zero, these deviations occur primarily 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. These interaction effects

Table 5: Mean Wage Change of Movers by Decile of Origin and Destination Establishment Effect, 2005–2010

		Destination Establishment Effect Decile									
		1	2	3	4	5	6	7	8	9	10
Origin Decile	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
	5	-0.403	-0.226	-0.135	-0.061	0.000	0.062	0.137	0.235	0.367	0.623
	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.

may bias the estimated compensating wage differential downward.

To assess how sensitive our estimates are to this potential violation of additive separability, Table 6 presents estimates from each of our benchmark specifications using only identifying variation that Figure 6 suggests is most likely to support the separability assumption. 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 $\hat{\theta}$ distribution or the $\hat{\Psi}$ distribution. It is identified discarding variation from jobs involving either low-wage workers (below the 5th percentile worker effect) or low-wage establishments (below the 5th percentile establishment effect).²⁸ 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, which Figure 6 suggests are most likely to satisfy the separability assumption.

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

²⁸A separate coefficient (not shown) is estimated for the interaction between the fatality rate and the remaining observations.

Table 6: Sensitivity of $\hat{\gamma}$ to Excluding Tails of the (θ, Ψ) Joint Distribution

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

(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, the pooled OLS model, has similar properties to the worker effects model, in which variation is driven by workers with median ability moving to establishments with higher Ψ , and 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. As expected, 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.

6.3 Separations Due to Mass Displacement Events

Gibbons and Katz (1992) suggested that longitudinal estimates of job characteristics, like compensating differentials, may be biased if workers learn about, and sort by, ability or comparative advantage over time. They proposed using estimates based on job transitions associated with mass displacement events, as these are less likely to be driven by worker and employer learning. The mass displacement sample should have a disproportionate number of job separations that occur for reasons unrelated to idiosyncratic match quality at either the origin or destination job. We therefore expect any remaining selection associated with time-varying worker or match effects that is not addressed in the OME model to be substantially reduced in the mass displacement sample.

Table 7 reports estimates from each model specification using an analysis sample that is restricted to job spells within two years 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.²⁹

Column (2) shows that in the worker effects model, the estimated $\hat{\gamma}$ from 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 (0.191) is very similar to the estimate for non-displaced workers (0.205). These results suggest that the scope for bias arising from drift in unobserved ability or match effects is economically negligible.

6.4 Completed Tenure as a Proxy for Match Quality

Abraham and Farber (1987) propose that completed tenure can serve as a proxy for the unobserved match-specific component of utility. In Table 8, we re-estimate 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 modest increases in the estimated effect of fatality risk on wages. In contrast, there is no significant change in $\hat{\gamma}$ in the OME model when completed tenure is added as a

²⁹Following 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.

Table 7: Mass Displacement Estimates

	(1)	(2)	(3)	(4)	(5)
	Pooled	Worker Effects	Match Effects	OME	TWFE
Fatality Rate (3-Yr MA)	0.475* (0.001)	0.079* (0.002)	-0.011* (0.002)	0.205* (0.001)	0.193* (0.001)
Fatality Rate \times Mass Disp.	0.209* (0.002)	0.003 (0.002)		-0.014* (0.002)	-0.012* (0.002)
Zero Fatality Rate	0.089* (0.000)	0.013* (0.000)	-0.004* (0.000)	0.016* (0.000)	0.016* (0.000)
Zero Fatality Rate \times Mass Disp.	-0.006* (0.001)	0.004* (0.001)		0.005* (0.000)	0.004* (0.000)
Mass Disp. Origin	-0.023* (0.000)	0.016* (0.000)		0.009* (0.000)	0.009* (0.000)
Mass Disp. Destination	-0.031* (0.000)	0.002* (0.000)		0.001 (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

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. * Indicates significance at the 0.01 level.

control.³⁰

6.5 Instrumental Variables Estimates Based on Co-worker Histories

In this section 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: $\Delta P_{it} = \Delta \pi_{k(i,t)} + \gamma \Delta R_{c(i,t),t} + \Delta \tau_t + \Delta \Psi_{J(i,t)} + (\Delta \mu_{i,J(i,t)} + \Delta \varepsilon_{it})$ 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 $E(\Delta R_{c(i,t),t} \Delta \mu_{i,J(i,t)} | \Delta \Psi_{J(i,t)}) = 0$.

³⁰Another 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.

Table 8: 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.

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 a potential change in unobserved match effects, $\Delta \mu_{i,J(i,t)}$. To do this, we exploit the relational network structure of the data. 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 $N(i,t)$.³¹ 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 upon 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

³¹Note that for observation $\ell = (i',\tau) \in N(i,t)$, $\Delta R_{\ell} = R_{c(i',\tau),\tau} - R_{c(i',\tau-1),\tau-1}$.

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 establishments 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.

6.5.1 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 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.

6.5.2 IV Results

Table 9 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 OME 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 of the IV 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.

Table 9: Instrumental Variable Estimates

	(1) First- Differenced	(2) Establishment Effects	(3) IV First Stage	(4) IV	(5) OME on IV Sample
Δ Fatality Rate	-0.048 (0.003)	0.236* (0.000)		0.210* (0.011)	
Avg. Δ Fat. Rate in $N(i.t)$ Fatality Rate			0.338* (0.001)		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.

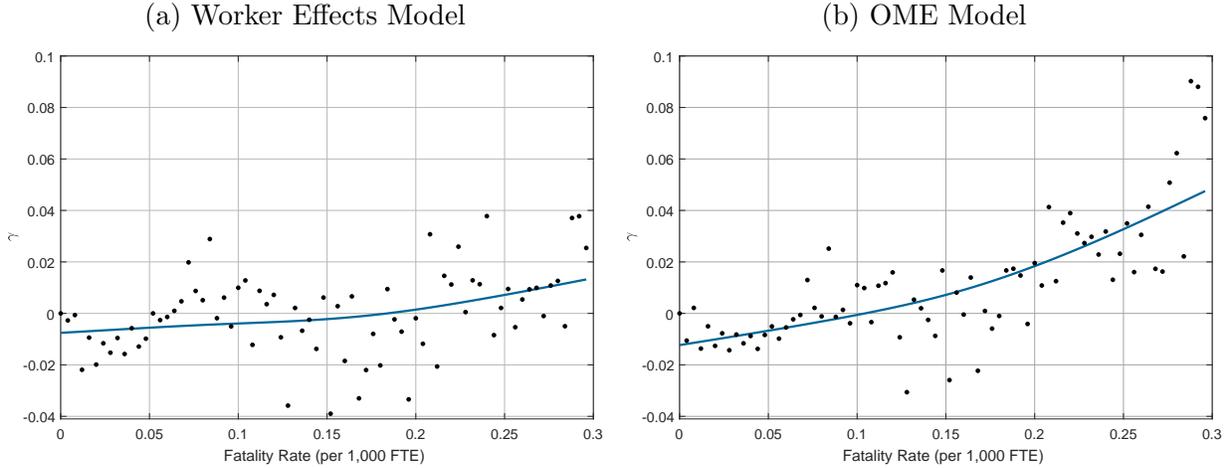
The IV estimate of the compensating wage differential in Column (4), $\hat{\gamma} = 0.210$, is slightly smaller than the effect estimated in the model controlling for establishment effects, and we reject the null hypothesis that the estimates are equal. The endogenous mobility bias that is corrected by the instrument relative to the establishment effects model appears to be modestly positive. However, when we use the IV sample to estimate the OME 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 OME model, the impact of any associated endogeneity bias is not economically meaningful.

6.6 Relaxing the Linearity Assumption: Semi-Parametric 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 necessary for our empirical approach or our conceptual framework. In Figure 7 we consider semi-parametric estimates, which suggest the linear model is not an unreasonable approximation at typical risk levels observed in our data.

To estimate the conditional relationship between risk and wage semi-parametrically, we discretize fatality risk into 75 equally-spaced bins for risk levels between 0 and 0.3 deaths per

Figure 7: Semi-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.

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 7 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).

6.7 Job Duration Models

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 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 10 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

Table 10: 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 parentheses. * Indicates significance at the 0.01 level.

from our model relate to separations initiated by workers, the dependent variable equals 1 if the worker 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.³²

Estimates in column (1) suggest that a one log-point increase 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.

³²The distinction between jobs that end due to a resignation as opposed to a termination is observed based on the cause of separation in the data.

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 10. 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 establishment effects is challenging because, as we show in Table 5, the majority of wage gains associated with job separations occur when workers move into higher paying establishments. Estimates in column (4) indicate 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.³³

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.

7 Conclusion

We 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 associated structural, theoretical, and reduced-form literatures. Specifically, we show that 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 illuminating the relationship between wages and job characteristics.

Relaxing key assumptions on the conditional exogeneity of the assignment of workers to jobs yields several important conclusions. We show that common panel-based approaches to eliminating ability bias can instead amplify bias caused by endogenous job mobility. The misspecification of these models can be largely corrected by incorporating firms into the model of equilibrium wage determination. We investigate one such model, and propose a theoretical analogue that combines elements of the hedonic search model proposed by Hwang et al. (1998) and the differentiated firms framework introduced by Card et al. (2018). This model

³³Although these estimates are from linear probability models, we find very similar patterns from logit specifications and in hazard models.

shows that the classic framework developed by Rosen (1974) can be extended to accommodate imperfect competition in labor markets while maintaining the equilibrium construct of a sorting-based hedonic pricing function. The model also highlights the importance of evaluating whether worker-firm match effects are economically meaningful in the wage-risk data generating process, as the presence of match effects in our proposed model changes estimates from having a preference-based interpretation to an equilibrium wage interpretation.

The empirical and conceptual issues we highlight are likely not unique to Brazil. Our analysis is motivated by the contrast between cross-sectional estimates of the compensating differential for fatal injury and the much smaller estimates from U.S. panel data (Brown 1980; Kniesner et al. 2012). Our analysis is also motivated by the fact that in Brazil a large part of variation in wages can be explained by movements between firms. This pattern, which is also consistent with hedonic search, is a stylized feature of matched employer-employee data in many countries (Abowd et al. 1999; Card et al. 2013; 2016; Abowd et al. 2012). In particular, 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. It therefore seems likely that our empirical approach should be broadly applicable beyond the setting of this paper.

The analysis of hedonic wage models is fraught with challenges for applied work, and no study can resolve them all. Like many studies, we do not directly observe every job amenity. However, by controlling for unobserved establishment and occupation heterogeneity we substantially relax assumptions about unobserved amenities that tend to be employer- or occupation-specific, such as health insurance. If, however, fatal and non-fatal risk tend to be bundled together in the same way across jobs within establishments, then the quantitative estimates from our application of the model correspond to the change in wages associated with changes in this composite bundle. In this case, our estimates recover the treatment effect on wages of changes in fatality risk and other correlated amenities, but does not separately identify the hedonic pricing functions of each amenity.

Nearly forty-five years after Rosen proposed the model of hedonic prices in implicit markets, the consensus remains that labor economists have much to learn about this topic. In recent years, creative strands of this literature have developed, each of which approaches elements of Rosen's question from different perspectives. Mas and Pallais (2017) recover worker preferences from a field experiment, while Wiswall and Zafar (2018) elicit willingness-to-pay in a survey. Since the identification of preference-based parameters remains one of the most challenging applications of Rosen's model, these approaches can provide insights that are not only informative about policy-relevant preference parameters, but could also be used to learn why preferences are so difficult to infer from equilibrium wages. Sorkin (2018) and

Taber and Vejlin (2016) also contribute creative approaches to understanding the problem of equalizing differences in labor markets. These papers step back from Rosen’s framework and consider a different object of interest—the share of variation in equilibrium compensation that can be explained by unobserved amenities based on revealed preferences for jobs. We anticipate that future progress on understanding equalizing differences in wages is likely to come from combining lessons from these innovative perspectives with the enduring insights of the hedonic wage literature.

References

- Abowd, J. M., Kramarz, F., Lengerman, P., McKinney, K. and Roux, S. (2012). Persistent inter-industry wage differences: Rent-sharing and opportunity costs, *IZA Journal of Labor Economics* **1**(1): 1–7.
- 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. (2017). Modeling endogenous mobility in earnings determination, *Journal of Business & Economic Statistics* **to appear**.
- 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.
- Alvarez, J., Benguria, F., Engbom, N. and Moser, C. (2018). Firms and the decline in earnings inequality in brazil, *American Economic Journal: Macroeconomics* **10**(1): 149–89.
URL: <http://www.aeaweb.org/articles?id=10.1257/mac.20150355>
- Barth, E., Bryson, A., Davis, J. C. and Freeman, R. (2016). It’s 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. (2018). Firms and labor market inequality: Evidence and some theory, *Technical Report S1*.
- Card, D., Cardoso, A. R. and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women, *The Quarterly Journal of Economics* **131**(2): 633.
- Card, D., Heining, J. and Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality, *The Quarterly Journal of Economics* **128**(3): 967–1015.
- Cornwell, C., Rivera, J. and Schmutte, I. M. (2016). 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.
- 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.
- Gelbach, J. B. (2014). When do covariates matter? and which ones, and how much?, *Journal of Labor Economics, Forthcoming*.
- 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.
- Hammit, 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.
- Hotz, V. J., Johansson, P. and Karimi, A. (2017). Parenthood, family friendly firms, and the gender gaps in early work careers, *Working Paper 24173*, National Bureau of Economic Research.
- 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.
- Krueger, A. B. and Summers, L. H. (1988). Efficiency Wages and the Inter-industry Wage Structure, *Econometrica* **56**(2): 259–93.
- Lalive, R. (2003). Did we overestimate the value of health?, *Journal of Risk and Uncertainty* **27**(2): 171–193.
- Lamadon, T., Mogstad, M. and Setzler, B. (2017). Imperfect competition and rent sharing in the us labor market, *Working Paper*.
- 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. (2018). The estimation of compensating wage differentials: Lessons from the *Deadliest Catch*, *Journal of Business & Economic Statistics*.
- Lavetti, K. and Schmutte, I. M. (2017). Gender differences in earnings and sorting on occupational safety, *Working paper*, Working Paper, University of Georgia.
- Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements, *American Economic Review* **107**(12): 3722–3759.
- 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.
- Rosen, S. (1988). The value of changes in life expectancy, *Journal of Risk and Uncertainty* (1): 285–304.
- Schmutte, I. M. (2015). Job referral networks and the determination of earnings in local labor markets, *Journal of Labor Economics* **33**(1): 1–32.
- Sorkin, I. (2018). Ranking firms using revealed preference, *The Quarterly Journal of Economics* **133**(3): 1331–1393.
- 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 Vejlín, 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.
- Wiswall, M. and Zafar, B. (2018). Preference for the workplace, human capital, and gender, *Quarterly Journal of Economics* **133**(1): 457–507.
- Woodcock, S. (2008). Wage differentials in the presence of unobserved worker, firm, and match heterogeneity, *Labour Economics* **15**: 771–793.

Online Appendix to Accompany Kurt Lavetti and Ian M. Schmutte
“Estimating Compensating Wage Differentials with Endogenous Job Mobility,”
January 17, 2025

The full manuscript is available at http://www.kurtlavetti.com/CDEM_vc.pdf. Please contact the authors at lavetti.1@osu.edu or schmutte@uga.edu with comments or for additional information.

A Additional Tables and Figures

Table A.1: 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.2: 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.3: 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	transfer with cost to firm
31	trans s/onus	transfer 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.4: 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 Maintenance 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.5: Estimated Compensating Wage Differentials for Full-Time Prime-Age Men,
Excluding Industry and Occupation Effects

	(1)	(2)	(3)	(4)
	Pooled	Worker Effects	Match Effects	OME
Fatality Rate	0.343* (0.001)	0.068* (0.001)	-0.006* (0.001)	0.156* (0.001)
Zero Fatality Rate	0.211* (0.000)	0.022* (0.000)	-0.006* (0.000)	0.018* (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.6: Heterogeneity by Region in Brazil

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

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.7: Estimates with Clustered Standard Errors

	(1)	(2)	(3)	(4)
	Pooled	Worker Effects	Match Effects	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 <i>pardo</i> or <i>preto</i>	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:	$\ln(Wage)$		$\ln(Wage) - X\hat{\beta}$	
	(1) Pooled	(2) Worker Effects	(3) Match Effects	(4) Orth. Match Effects
Fatality Rate (3-Yr MA)	0.365* (0.003)	0.011 (0.014)	-0.006 (2.079)	0.203* (0.009)
Zero Fatality Rate	0.079* (0.001)	0.014* (0.003)	0.001 (0.188)	0.021* (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 Inference Based on Industry-Occupation Aggregates

We estimate a version of equation 4 that consolidates variation in wages and fatality risk to the same level of aggregation. To do this, we adjust the second-stage equation of the OME model, Equation 4, to include a completely unrestricted industry-occupation-year component, $\mu_{c(i,t),t}$.

$$w_{it} - x_{it}\hat{\beta} = \mu_{c(i,t),t} + \theta_i + \Psi_{J(i,t)} + \epsilon_{it}. \quad (\text{B.18})$$

The primary object of interest in this step is $\mu_{c(i,t),t}$, which is the smallest component of the aggregate variation in wages that fully contains any compensating wage differential associated with fatality risk, as well as any other industry-occupation level amenities.

We then collapse the estimated $\widehat{\mu_{c(i,t),t}}$ from the job-level model, and estimate the industry-occupation-year-level model:

$$\widehat{\mu}_{ct} = \check{\gamma}R_{ct} + \tau_t + Z_{ct}\delta + \xi_{ct}. \quad (\text{B.19})$$

The index c denotes an industry-occupation cell, measured at the same level of aggregation as fatality risk. The dependent variable, $\widehat{\mu}_{c,t}$ is the cell-year effect estimated from fitting equation (B.18) to the full analysis sample used in Table 2. The term Z_{ct} allows for occupation controls, industry-occupation interactions, and cell-average worker and firm effects.

Table B.10 presents estimates from Equation B.19. The model reported in Column (1) only controls for year effects, and yields a point estimate on fatality risk, $\widehat{\gamma} = 0.211 \pm 0.079$. In Column (2) we control for 1-digit industry-occupation effects, and add back the cell-average worker and plant effects estimated in Equation B.18. We include these controls because $\widehat{\mu}_{ct}$ has a fairly high correlation with θ_i and $\Psi_{J(i,t)}$, and if these components enter the error term ξ_{ct} they could potentially still cause some endogeneity bias. This more saturated model gives a point estimate of $\widehat{\gamma} = 0.237 \pm 0.041$.

These results suggest, first, that our finding of a statistically significant and positive compensating differential is robust to analyzing variation in wages and fatality risk at the same level of aggregation. Second, while the confidence intervals overlap, the results in this aggregate analysis are slightly larger than our benchmark point estimate. In the remainder of the paper we evaluate a range of alternative approaches aimed at quantifying the potential influence of any remaining unmodeled match-specific heterogeneity. Consistent with these aggregate patterns, we generally find that attempting to correct for match-specific heterogeneity has only modest positive effects on γ relative to estimates from the OME model.

Table B.10: Estimates Based on Industry-Occupation Aggregates

Dependent Variable:	$\hat{\mu}_{ct}$	
	(1)	(2)
Fatality Rate (3-Yr MA)	0.211* (0.079)	0.237* (0.041)
Zero Fatality Rate	0.050 (0.009)	0.004 (0.005)
Year Effects	X	X
1-Digit Industry-Occupation Effects		X
Avg. Worker and Estab. Effects		X
N Obs	6,264	6,264
N Clusters	1,179	1,179
R-Sq	0.053	0.617

Notes: Estimates from an aggregate model with one observation for each year in each 2-digit industry by 3-digit occupation cell. The dependent variable is a cell-by-year effect in log wages, estimated from equation (B.18). Fatality Rate is measured in deaths per 1,000 full-time full-year equivalent workers. Observations are weighted by the number of job-level observations. Standard errors are clustered by 3-digit occupation by 2-digit industry. * indicates significance at the 0.01 level.

C Institutional, Data, and Model Details

C.1 Formal Employment in Brazil

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.

C.2 Wage Regulations in Brazil

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 Regulamentadora de Segurança e Saúde 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.

C.3 Details of Fatality Rate Calculations

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

$$R_c = \frac{F_c}{(H_c/2,000)} \times (100,000). \quad (\text{C.20})$$

Throughout the paper, cells are defined as unique combinations of 2-digit industry code, 3-digit occupation code, and year. These cells are mutually exclusive and exhaustive. 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.³⁴ For each job, j , in the cell c , we count the number of hours worked as $H_i = (\text{MonthsWorked}/12) * 50 * (\text{Hours/Week})$. 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.

C.4 Relationship to Value of Statistical Life

We compare our estimates of the value of statistical life (VSL) to estimates not derived from Brazilian wage data. Doing so serves two purposes. First, this widely-used policy parameter provides a useful benchmark for interpreting the magnitude of our estimates. Second, the comparison validates our conclusion that the within-worker estimate is severely biased toward zero.

³⁴Changes 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.

To benchmark our estimates, we rescale VSL estimates from the US into Brazilian reais using the benefit transfer method recommended by Hammitt and Robinson (2011) and as implemented by Miller and Façanha (2016). This approach is necessary because there are no comprehensive studies of the VSL for Brazil. The benefit transfer method adjusts VSL estimates for differences in per-capita income, which has large effects on the scale of willingness to pay. 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:

$$VSL_{Brazil} = VSL_{US} \times \frac{\text{GNI per capita in Brazil}}{\text{GNI per capita in U.S.}}. \quad (\text{C.21})$$

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 (C.21) 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 implies the VSL in Brazil should be between between 2.14 and 3.10 million 2003 reais. The VSL implied by our pooled model, 2.20 million reais (Column 1 of Table 2,) is at the bottom end of this range. However, our benchmark OME model (Column 4 of Table 2) gives a VSL of 1.73 million reais. The VSL estimate derived from the within-worker models is much lower than the benefit transfer calculation suggests.

C.5 Bias Decomposition

We consider the nature of omitted variable bias following the approach suggested by Abowd et al. (1999) and in the spirit of Gelbach (2014). Our goal is to show how much bias is contributed by observables, and how much by unobserved worker and employer heterogeneity.

To do so, we compare the estimated compensating differential from our preferred OME estimator with an estimate from a version of the pooled model that includes only year effects:

$$y_{it} = \gamma R_{c(i,t)} + \gamma^0 \mathbf{1}(R_{c(i,t)} = 1) + \tau_t + \varepsilon_{it}. \quad (\text{C.22})$$

We continue to include an indicator for jobs with zero risk, $\mathbf{1}(R_{c(i,t)} = 1)$, for comparability and because we have found that selection into jobs with no fatality risk is a bellwether for endogenous selection, as in Caetano (2015). Let z_{it} be the vector of additional covariates included in our preferred OME specification (experience interacted with race and 1-digit occupation controls). Also, we use $\hat{\gamma}$ and $\hat{\gamma}^0$ to denote the estimates of interest from equation (C.22). Finally, for exposition, we use $\hat{\gamma}^*$ and $\hat{\gamma}^{0*}$ to denote the corresponding estimates based on the preferred OME specification.

Following Abowd et al. (1999) and Gelbach (2014), we can express the bias in equation

Table C.11: Decomposition of Bias in the Estimated Compensating Differential: Pooled vs. OME

Component	Risk		Risk=0	
	(1) Raw	(2) Share	(3) Raw	(4) Share
Total Bias	-.351	100%	0.381	100%
Worker FE, κ_θ	-.212	60%	0.288	70%
Estab. FE, κ_ψ	-.272	77%	0.101	24%
Observed, $\kappa_z^T \hat{\beta}_z$	0.134	-37%	0.027	6%

Table entries decompose bias in the estimated coefficient on fatality risk (columns 1 and 2) and the estimated coefficient on a dummy for whether fatality risk is equal to zero (columns 3 and 4). The bias is measured as the contrast between a pooled model that only controls for year effects, and our preferred OME specification.

(C.22) relative to OME as

$$\delta^R = \hat{\gamma} - \hat{\gamma}^* = \kappa_\theta + \kappa_\psi + \kappa_z^T \hat{\beta}_z \quad (\text{C.23})$$

$$\delta^0 = \hat{\gamma}^0 - \hat{\gamma}^{0*} = \kappa_\psi^0 + \kappa_z^{0T} \hat{\beta}_z^0 \quad (\text{C.24})$$

where κ_θ is the correlation between R and θ , conditional on $R > 0$ and the year effects. That is, it is the coefficient on R in the auxiliary projection of worker effects, θ , onto the variables in the right-hand side of equation (C.22). The terms κ_ψ and κ_z are the analogous expressions for establishment effects and the vector of regressors in z . The term $\hat{\beta}_z$ is the estimated vector of coefficients on z from the OME specification. The terms in the second line are analogously defined for γ^0 .

We find $\hat{\gamma} = -0.181$ in the pooled model that does not include the variables in z . The total bias relative to OME is, therefore, $\delta_R = -0.181 - 0.170 = -0.351$. Columns (1) and (2) in Table C.11 decompose this bias into contributions from worker effects, establishment effects, and observables. We find, as expected, that omission of worker and establishment effects both induce strong negative bias of roughly equal shares. Working against these effects, the omission of observed worker characteristics contributes a positive bias. This may be because less experienced workers tend to be in safer jobs, reflecting structural changes in the labor market opportunities for workers from different entry cohorts. The majority of the downward bias is from omitted establishment effects.

In columns (3) and (4), we decompose the sources of misspecification that contribute to the estimated coefficient on the dummy for jobs with $R = 0$. We find $\hat{\gamma}_0 = 0.430$ in the pooled model that does not include the variables in z . By contrast, the comparable coefficient estimate from the OME model is 0.014. The total bias is, therefore, $\delta_0 = 0.430 - 0.014 = 0.416$. We find 70 percent of this bias is from high-wage workers sorting into the safest jobs. A further 24 percent is due to the safest jobs being concentrated in high-wage establishments. The balance, six percent, is due to the observed covariates.

Table C.12: Decomposition of Bias in the Estimated Compensating Wage Differential:
Within-worker vs. OME

Component	Risk		Risk=0	
	(1) Raw	(2) Share	(3) Raw	(4) Share
Total Bias	-.103	100%	0.011	100%
Estab. FE, κ_{ψ}^0	-.103	100%	0.003	38%
Observed, $\kappa_z^{0T} \hat{\beta}_z^0$	0.000	0%	0.008	62%

Table entries decompose bias in the estimated coefficient on fatality risk (columns 1 and 2) and the estimated coefficient on a dummy for whether fatality risk is equal to zero (columns 3 and 4). The bias is measured as the contrast between a model that controls only for year effects and for worker effects, and our preferred OME specification.

In Table C.11, we repeat the same exercise to compare bias in the within-worker model to the preferred OME specification. In the within-worker model that does not include the covariates in z , we find $\hat{\gamma} = 0.067$. Hence, the total bias relative to OME is $\delta_R = 0.067 - 0.170 = -0.103$. Strikingly, we find the bias is entirely explained by omitted establishment effects. The omission of observed covariates does not contribute any net bias. This finding is consistent with our findings, reported in Appendix Tables A.1 and A.5, that estimates of the OME model are not very sensitive to changes in the included controls. The estimated bias on the $R = 0$ dummy is negligible.

C.6 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 eliminate 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 m th 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,

$$\Pr[Z_{im} = 1|\omega_{im} = 1] = \sum_{\omega_i|\omega_{im}=1} \Pr[Z_{im} = 1|\omega_i, \omega_{im} = 1] \Pr[\omega_i|\omega_{im} = 1]. \quad (\text{C.25})$$

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

$$\Pr[Z_{im} = 1|\omega_i, \omega_{im} = 1] = \frac{\exp(\bar{u}_{s(i)m})}{\exp(\bar{u}_{s(i)0}) + \sum_{m'} \omega_{im'} \exp(\bar{u}_{s(i)m'})}. \quad (\text{C.26})$$

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

$$\Pr[\omega_{im} = 1] = \frac{L_{sm}}{N_s} + \lambda \left(1 - \frac{L_{sm}}{N_s}\right). \quad (\text{C.27})$$

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

$$\Pr[\omega_i|\omega_{im} = 1] = \sum_{m'} \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \Pr[M_{im'} = 1|\omega_{im} = 1]. \quad (\text{C.28})$$

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 $m \neq m'$

$$\begin{aligned} & \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \\ &= \frac{\prod_{\ell} \Pr[\omega_{i\ell} = 1]}{\Pr[\omega_{im} = 1] \Pr[\omega_{im'} = 1]} \\ &= \frac{\Pr[\omega_i]}{\Pr[\omega_{im} = 1] \Pr[\omega_{im'} = 1]}. \end{aligned} \quad (\text{C.29})$$

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

$$\Pr[\omega_i|\omega_{im} = 1, M_{im} = 1] = \frac{\prod_{\ell} \Pr[\omega_{i\ell} = 1]}{\Pr[\omega_{im} = 1]}. \quad (\text{C.30})$$

So we have

$$\Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] = \frac{\Pr[\omega_i|\omega_{im} = 1, M_{im} = 1]}{\Pr[\omega_{im'} = 1]} \quad (\text{C.31})$$

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

$$\Pr[M_{im'} = 1|\omega_{im} = 1] = \frac{\Pr[\omega_{im} = 1|M_{im'} = 1] \Pr[M_{im'} = 1]}{\Pr[\omega_{im} = 1]}. \quad (\text{C.32})$$

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

$$\frac{\Pr[\omega_{im} = 1|M_{im'} = 1] \Pr[M_{im'} = 1]}{\Pr[\omega_{im} = 1]} = \frac{\lambda \frac{L_{sm'}}{N_s}}{\Pr[\omega_{im} = 1]} = \frac{\lambda L_{sm'}}{L_{sm}(1 - \lambda) + \lambda N_s} \quad (\text{C.33})$$

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

$$\frac{\Pr[\omega_{im} = 1|M_{im} = 1] \Pr[M_{im} = 1]}{\Pr[\omega_{im} = 1]} = \frac{\frac{L_{sm}}{N_s}}{\Pr[\omega_{im} = 1]} = \frac{L_{sm}}{L_{sm}(1 - \lambda) + \lambda N_s} \quad (\text{C.34})$$

Finally, we can rewrite the RHS of (C.28)

$$\begin{aligned} & \sum_{\{m':\omega_{im'}=1\}} \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \Pr[M_{im'} = 1|\omega_{im} = 1] \\ &= \Pr(\omega_i|\omega_{im} = 1) \left[\Pr[M_{im} = 1|\omega_{im} = 1] + \sum_{\{m':m' \neq m; \omega_{im'}=1\}} \frac{\Pr[M_{im'}=1|\omega_{im}=1]}{\Pr[\omega_{im'}=1]} \right] \end{aligned} \quad (\text{C.35})$$

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

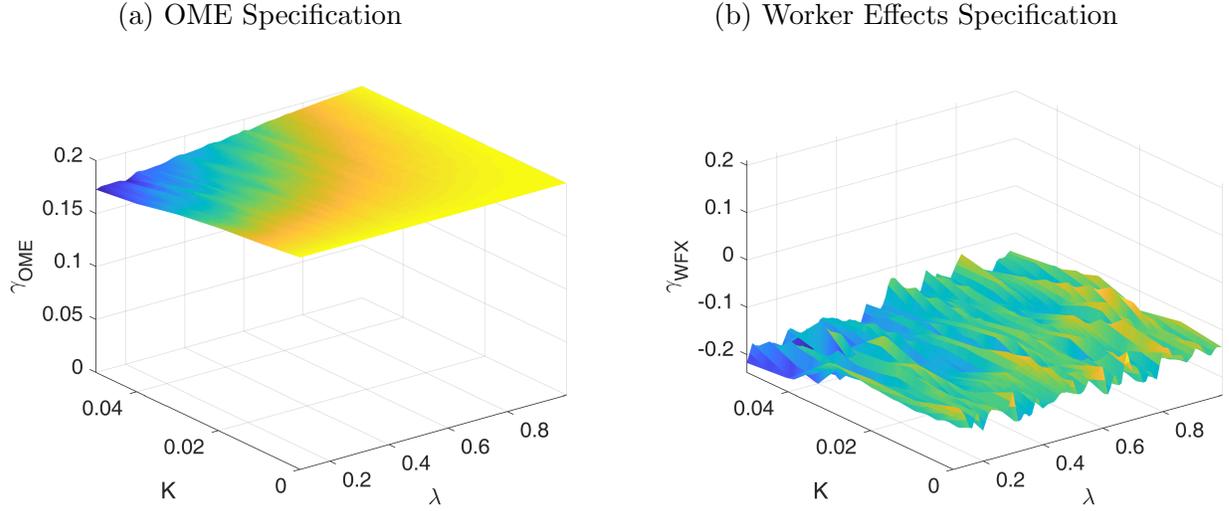
$$\begin{aligned} & \sum_{\{m':\omega_{im'}=1\}} \Pr[\omega_i|\omega_{im} = 1, M_{im'} = 1] \Pr[M_{im'} = 1|\omega_{im} = 1] \\ &= \Pr(\omega_i|\omega_{im} = 1) \frac{1}{L_{sm}(1-\lambda)+\lambda N_s} \left\{ L_{sm} + \sum_{m' \neq m} \frac{\lambda L_{sm'}}{\Pr[\omega_{im'}=1]} \right\} \end{aligned} \quad (\text{C.36})$$

C.7 Monte Carlo Simulation Details

The monte carlo simulation model is designed to evaluate the performance of the OME and worker effects model specifications over a range of assumptions about labor market conditions. The simulation differs from the theoretical model in several ways. First, we abstract from heterogeneous firm technology across industries and occupations. Instead we consider a constant marginal cost of providing safety, and constant marginal willingness to accept safety, that satisfy the first order conditions of the model. Second, we instead allow both λ and K to vary instead of fixing one parameter relative to the other and relative to the number of workers and firms. A high λ implies, all else equal, that more total offers are made. If λ increases holding K fixed, this implies a reduction in the expected number of firms in the model. That is, if each firm makes offers to a larger share of workers, and yet the probability of acceptance for a firm offering the average \bar{u} remains fixed, there must be fewer firms making offers.

The distribution of worker skill types in the model is parameterized such that $\ln \theta_s$ matches the mean and standard deviation of the empirically estimated distribution of worker

Figure C.1: Monte Carlo Estimates of $\hat{\gamma}$ when True $\gamma = 0.2$
Assuming 20 External Offers per Period



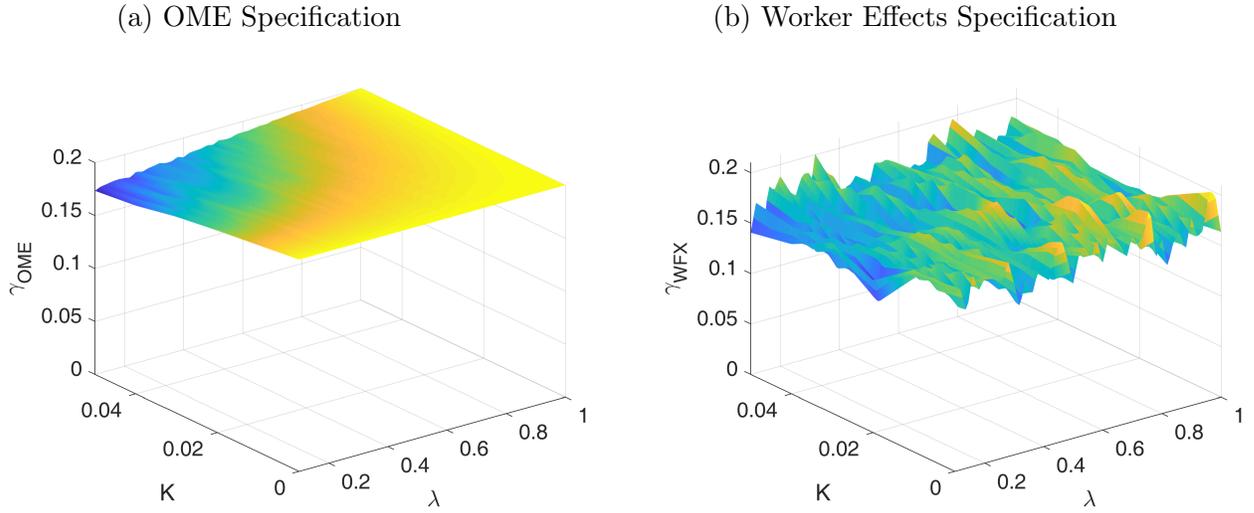
Notes: Estimates are based on 25000 simulated workers over 30 periods for each (λ, K) pair.

effects in the OME model. Specifically, $\ln \theta_s$ is normally distributed with mean zero and standard deviation 0.456. Similarly, the distribution of firm productivity is parameterized such that $\ln T_j$ and $y(R)$ have the same means and covariance matrix as the empirical joint distribution of establishment wage effects and R in the OME model. Specifically, $(\ln T_j, y(R))$ is jointly normally distributed with mean vector $(0, 0.083)$ and covariance matrix $\begin{bmatrix} 0.089 & -0.003 \\ -0.003 & 0.008 \end{bmatrix}$. Preference shocks ϵ are assumed to be distributed EV1 with mean zero and standard deviation 0.2.

The model is simulated for 25,000 individuals over 30 periods. In the first period of the model workers are assigned a random draw from the distributions of $\ln \theta_s$, $\ln T_j$, and $y(R)$. Beginning in the second period, workers receive a fixed number of outside offers per period, in addition to the option to remain in the same job, in which case all of the outcomes from period $t - 1$ carry over to period t . The exogenously assigned outcomes in the first period are not included in the simulated data used to evaluate model performance; only data from periods 2 to 30 are used. In the baseline model, the number of external offers is fixed at 3. Simulation results in Figure C.1 show that the outcomes are very similar when the number of offers per period is increased from 3 to 20.

The simulated data are slightly different than our empirical observational data from Brazil. First, the true worker and firm wage effects are observed. Second, the model abstracts from any variation in wages within jobs. In this sense, the wage outcome is comparable to the dependent variable in the second stage of our OME specification, after removing the effects of experience and time. Given this, the OME specification is equivalent to regressing the log wage on the fatality rate, simulated worker effect, and simulated firm effect. The full

Figure C.2: Monte Carlo Estimates of $\hat{\gamma}$ when True $\gamma = 0.2$
 Assuming $cov(\ln T_j, R) = 0$



Notes: Estimates are based on 25000 simulated workers over 30 periods for each (λ, K) pair.

design matrix of binary worker and establishment identifiers is not required for estimation.

Figure C.2 depicts simulated estimates when the covariance between $\ln T_j$ and R is set to zero. The figure shows that even when $\ln T_j$ and R have zero correlation in the offer function, the worker effects model has a larger bias than the OME model. This occurs because workers' decisions to switch jobs when offered an increase in utility creates a relationship between observed changes in R and changes in T_j . The maximum bias in the worker effects model is -40.7%, while the maximum bias in the OME model is -7.2%. The mean biases when $1 \geq \lambda \geq 0.5$ and $0.05 \geq K \geq 0$ are 3.96% and 0.46% in the worker effects and OME specifications, respectively.