THE UNIVERSITY OF CHICAGO

UNCERTAINTY OVER FATALITY RISKS AND THE VALUE OF STATISTICAL LIFE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
MAURICIO CHIKITANI

CHICAGO, ILLINOIS
JUNE 2021

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Acknowledgments

I would like to thank my dissertation committee, Dan Black, Stéphane Bonhomme and Peter Hull for their invaluable guidance. This work certainly would not be possible without their insight and support.

I would like to thank my family, including—and perhaps especially—the ones that were not even around when I first came to Chicago. The geographical distance between us taught me the most important lesson of my PhD. And sticking with Brazil, I would like to thank my friends back there (and those I met there but have left the country by now), for an often needed respite from my Economics-centered Chicago life.

Still, Chicago actually felt like home for the past six years, and I have many people to thank for that. Newport folks were absolutely key, especially in this pandemic period. My appreciation also goes for the Pub Night people, the Crazy Chat idiots and Department/University friends in general.

I avoided naming anyone other than my committee because I am bound to unfairly forget someone, and the contrast with explicitly named individuals would only add insult to the injury. Plus, if you are reading this, chances are you indeed helped me one way or another, so thanks!

CHAPTER 1

INTRODUCTION

The Value of Statistical Life (VSL; Schelling, 1968) is an important policy parameter when assessing the benefits of decreasing mortality risks.[1] Its estimation procedure often infers people's attitude towards fatality risks through their labor market choices. But do workers really know the fatality risk of their jobs? This has generally been assumed to be the case, because work is a routine part of life that people have a strong incentive to be informed about (Thaler & Rosen, 1976). Moreover, this assumption has the practical benefit of allowing analysts to use objective measures of risk as proxies for workers' perceptions, which are typically not observed but guide their decisions. Yet, if there is uncertainty over these risks, the persons we see employed at a job tend to have relatively optimistic beliefs about these hazards—which is precisely why they chose to work there. Therefore, by relying on these objective proxies, we would be undervaluing the VSL. In this paper, I show that this uncertainty does exist—and quantify the subsequent bias in valuation—but that the assumption of perfect information is partially warranted because workers eventually learn through work experience.

I start by developing a simple model that illustrates why uncertainty over fatality risks is concerning and the logic behind my empirical approach. Workers base their job choice on their perceptions of fatality risks, which they are unsure about. Job choice implies that, conditional on being employed, workers' perceptions are optimistic relative to an objective assessment of risk. This systematic difference leads to an undervaluation of fatality risks. Yet, workers do acquire more information as they gain experience, and update their beliefs in a Bayesian fashion. This on-the-job learning means that experienced workers tend to

[1] The VSL is the amount of money a large group of people is willing to pay to eliminate a marginal risk of fatality that is expected to kill one of them in the following year. For example, how much money 1,000 people would pay to decrease the fatality risk each one of them face by 1 in 1,000 per year. See Viscusi (2018) for a comprehensive review of the historical development of the VSL and its applications. Currently, the VSL is a key input for the cost-benefit analysis of COVID-19 related policies (Alvarez, Argente, & Lippi, 2020; Greenstone & Nigam, 2020).

have accurate beliefs. But along his tenure, the worker may receive a signal—a work-related fatality, in my case—that worsen his perception to the point that he decides to quit the job. Importantly, the Bayesian learning process implies that this updating should be stronger if his beliefs are based on less information, i.e., he is less experienced.

I test this prediction using employer-employee linked administrative data from Brazil. I use deaths in a work-related accident as a dire signal about the fatality risk a job poses to the surviving coworkers. This interpretation is possible because I focus my attention on coworkers in the same establishment and occupation as the deceased person, so they are likely to face similar job hazards. In order to isolate the informational effect of a work-related fatality, I compare my treated coworkers to work colleagues of someone that died for reasons unrelated to the job. In both cases, there is grief and disruption to the production process, but only the coworkers of someone who died in a work-related accident should be revising upwards their perceptions of occupational hazards.

In comparison to coworkers of a person that died for reasons unrelated to work, I see a clear spike in the number of quits in the year of death among those that received the dire signal, despite their behavior being similar in the years prior to the death. In establishments in which the death is work-related, the quitting probability of surviving coworkers increases by 1.29 p.p. in the year of the death, a 16% increase relative to the baseline level. Moreover, this effect is concentrated among relatively new coworkers; once I start restricting my sample to more experienced persons, the spike in quitting probability fades away: among coworkers with at least three years of tenure, there is no quit reaction, and this holds for even more experienced people.

These patterns are highly suggestive of the uncertainty and learning over job hazards predicted by theory, but they are not definitive. Importantly, work-related deaths could objectively increase the risks of a job for the surviving coworkers, relative to deaths unrelated to work. In this scenario, surviving coworkers should be more likely to quit precisely if they have accurate beliefs. Yet, I show there is no evidence in the data of an objective change in fatality risk driving my results, whether through non-fatal injuries, physical capital destruction or firm growth. Therefore, I conclude that uncertainty over fatality risks in Brazil is pervasive.

Moreover, the learning interpretation relies on preferences towards fatality risks being homogeneous. So I also show that the fatality rate wage premium in the data is not consistent with preference heterogeneity driving my results. Still, this evidence is more speculative, as the relationship between preferences and the wage premium is unclear (Hwang, Mortensen, & Reed, 1998; Hwang, Reed, & Hubbard, 1992). Therefore, while I am confident that uncertainty over occupational fatality risks is a pervasive characteristic of the Brazilian labor market, my learning interpretation is more tentative. Nevertheless, preference heterogeneity would imply that high tenure workers have lower willingness to pay (WTP) for reductions in fatality risk, which would be a force against the bias in valuation I document.

I also discuss some additional evidence to enhance the credibility of my results. First, the quitting patterns are the same in a smaller sample that only includes deaths of experienced workers, which indicates that my main results are not driven by tenure-specific fatality risks. Then, I show that these work-related deaths have no discernible impact on the wage of surviving coworkers, and that human capital interactions are unlikely to be driving my results. Furthermore, I also find some evidence that those quitting establishments after the work-related fatality struggle to find subsequent formal employment, which is consistent with these persons quitting because they were repelled by their former job—rather than being attracted to a new one.

Using my model, I can identify workers' WTP for reductions in fatality risks using quit rates, similar to Gronberg and Reed (1994). I then compare how different these estimates are if I focus on the group of workers that presumably have more accurate perceptions—and hence close to the objective fatality rate I can actually observe—versus an unrestricted sample. The results indicate that ignoring uncertainty means we underestimate the WTP for reductions in fatality risk by roughly 33%. This translates into a VSL of US$1.10 million in Brazil, versus a naive estimate that disregards the bias caused by uncertainty of US$0.73 million. As mentioned above, preference heterogeneity implies that more experienced workers have lower WTP for these risk reductions, so this bias is potentially larger in magnitude. Given that the VSL is instrumental in designing public policy, this bias is concerning (U.S. Office of Information and Regulatory Affairs, 2017; Viscusi, 2018). For example, the Clear Air Act alone may have had its cumulative benefit between 1990 and 2020 underestimated by almost US$1 trillion, if the bias in the U.S. is similar to the one I found in Brazil (Industrial Economics Inc., 2011).

1.1 Literature review

How to incorporate the potential for preventing premature deaths in a cost benefit analysis? Certainly, saving more lives is valuable, so it should be counted in the benefits of the intervention. But it is almost repugnant to put a monetary value on them, because this exercise may seem akin to "pricing" a human life. Nevertheless, these are often important decisions: in the United States, the Office of Information and Regulatory Affairs reports that among its major federal rules[2] "the largest benefits [...] are associated with regulations that reduce risks to life" (U.S. Office of Information and Regulatory Affairs, 2017). Therefore, the question is not whether we should price these benefits—a decision will have to be made—but instead to make sure that we are doing so properly.

The concept accepted as the proper measure of such benefits is the Value of Statistical Life (VSL), proposed by Thomas C. Schelling in 1968.[3] Before defining the VSL and discussing its properties, it is useful to entertain the alternative methods being proposed and used at the time. According to the theoretical discussion posed by Mishan (1971), there were essentially four other methods for pricing lives: (1) the present discounted sum of the

[2] A major federal rule is one that may result in expenditures of more than $100 million in at least one year, and therefore must be scrutinized with a cost-benefit analysis by the corresponding agency, according to the Unfunded Mandates Reform Act of 1995.

[3] See Schelling (1968).

deceased person's expected future earnings,[4] which was also the most common approach; (2) a slight modification of the previous method that removed the decedent's expected consumption from the earnings streams; (3) an implicit value based on democratic decisions over investments that could potentially increase or decrease the number of deaths; (4) an insurance based method, predicated on the premium a person is willing to pay and the probability of being killed as a result of engaging in some specific activity.

The first two methods would only make sense to the extent that the value of a life is limited to its contribution to national production, which is absurd. Moreover, they reveal very little about anyone's preferences and focus instead on their resource constraint. Given the goal of informing policy decisions, the third method is useless, for it is circular: any democratically chosen investment would pass a cost-benefit analysis by definition. The fourth method is better because the willingness to pay for insurance is more closely linked to a person's preferences. Yet, it falls short because an insurance policy in such cases would compensate others, rather than the decedent himself.

Following this logic, a better measure would be a person's willingness to pay to prevent their own premature death. There are two issues here: first, one may be concerned that people should be willing to pay as much as they could to avoid death, which would lead to the first two methods described above, depending on how you treat consumption. Yet, the very fact that people are willing to buy life insurance shows that they are also willing to forego some income to compensate others in case they die. This means people value the wealth of loved ones in a state of the world in which they are dead, bounding the efforts they are likely to make to prevent such state in the first place. Second, we cannot use insurance to back this valuation, because the policy cannot be paid to the victim, so the approach is not feasible.

Schelling's insight was to focus on the original goal of informing public policy: most such interventions do not target any specific individual. Instead, they tend to slightly change the

[4] Net of the deceased's nonhuman capital returns, for these would presumably remain despite his passing.

risk to life for a large number of people. Which leads to the definition of the VSL: how much money a large group of people would be willing to pay for an intervention that prevents the death of one of them. Because the identity of the life being saved is not known, each individual in the group would only be preventing their own death with a small probability. And trading-off small risks to life is something people do on a daily basis, e.g., by driving a car or wearing their seat-belts, so the amount they are willing to pay for small reductions in risks to life is limited, and the strategy is empirically feasible.[5] Finally, because most people would survive, the monetary cost of the intervention falls on the individual himself.

While conceptually appealing, the VSL would not enjoy its current widespread use[6] without some developments in its estimation process. The first major development is due to the seminal paper by Rosen (1974), in which he establishes a theoretical connection between the price of a good, its underlying characteristics and its consumer's preferences. In a competitive market, the variation in prices of close substitutes that differ across some of their characteristics (like build quality, features, etc.) is in equilibrium a reflection of the preference of the marginal consumer towards said characteristics. Intuitively, if consumers appreciate cars with air conditioning, then a car with this feature should command a premium that can be interpreted as the marginal consumer's willingness to pay for air conditioning. This method could be applied to the VSL estimation problem in a market in which people purchase similar goods with varying levels of risk to life.

Thaler and Rosen (1976) is the first such application. Instead of a regular consumer good, the authors focused on jobs; instead of a price tag, a job pays a monetary compensation in the form of wages, and different jobs pose different fatality risks for workers, who "shop" for their preferred one. In order to estimate the VSL, they used a regression of wages on fatality rates, i.e., an aggregate measure of deaths per hour worked for a given occupation, which implied VSL estimates of around $200,000 at the time. As I discuss below, the usage

[5] Curiously, Schelling himself did not seem to think this was the case at the time, noting that "The main problem is that people have difficulty knowing what is worth to themselves, cannot easily answer questions about it, and may object to being asked. Market evidence is unlikely to reveal much."

[6] See U.S. Office of Information and Regulatory Affairs (2017), U.S. EPA (2016), U.S. DOT (2016).

of labor market data would eventually become widespread—though not the only approach—when estimating the VSL, for such data is often of good quality and reflects a high stakes environment in which people make choices between monetary compensation and risks to life.

Yet, there are some caveats with the hedonic wage approach, as pointed out in the original article itself (Thaler & Rosen, 1976). First, it does not allow for any heterogeneity across workers, with the exception of different tolerance towards risk to life. Therefore, "[the] analysis of worker job choice is confined to people with identical personal characteristics. [...] It is as if there are separate risk markets for workers with each bundle of personal characteristics, and the present analysis of worker choice is confined to only one of those markets." But it is virtually impossible to argue that any empirical work is truly comparing workers with identical personal characteristics, as many of them are unobserved. For example, Hwang et al. (1992) re-iterate this point by arguing that if workers have different levels of innate productivity and safety is a normal good, then high productivity workers will select themselves into high-paying, low risk jobs. If this productivity is not accounted for, then the VSL estimates obtained with labor market data should be biased.

Furthermore, the competitive market framework used to justify the hedonic method may be too far removed from reality to be useful. Using a search model, Hwang et al. (1998) showed that the Rosen result about the relationship between preferences and the contract curve (i.e., the equilibrium mapping between job risk and wage) no longer holds. The intuition is simple: in a competitive market, equilibrium wages are set so that the marginal worker is indifferent between the jobs. Yet, in a search environment the workers may not seamlessly move between jobs, because search frictions prevent them from doing so. Moreover, if there is heterogeneity among firms, more productive ones offer both better wages and smaller risks.[7] In this equilibrium, search frictions make the fatality risk premium turn negative, even if workers would prefer not to risk their lives.

A recent paper by Lavetti and Schmutte (2018) argues that both of these difficulties can

[7] The opportunity cost of a vacancy is higher for more productive firms, so they offer a better compensation package to avoid that situation.

be overcome with richer labor market data structure. In particular, the authors argue that a regression specification that has both employer and employee fixed effects can ameliorate the issues above by removing the influence of unobserved heterogeneity that is fixed over time—both at employer and employee level. I build upon their work when estimating the VSL myself, but instead of a regression of log wages on fatality rates, I use a regression of quitting decisions on fatality rates and log wages, as in Gronberg and Reed (1994). The intuition is that the decision to quit a job is a function of the utility a worker derives from it, so it can be used to reveal one's preferences through the marginal rate of substitution between risk and monetary compensation. Another advantage of using quits is that in many settings, including the one I study, wages are constrained by institutional features such as minimum wages or mandatory risk premiums. This makes interpretation of hedonic wages harder, but it is not an issue for quitting decisions.

In any case, one must recognize that hedonic estimates of the VSL are used in practice. For example, both the U.S. Environmental Protection Agency (EPA) and the Department of Transportation (DOT)—probably the regulators that most heavily make use of the VSL in their decisions—rely on several hedonic method studies in their meta-analyses that inform the estimates they use.[8] Still, of the thirteen hedonic wage studies the agencies consider in their meta-analyses, only Kniesner et al. (2010) and Kniesner et al. (2012) include worker fixed effects in their specifications and none of them has employer fixed effects.

The goal of my dissertation is to raise a different concern, one that applies to the estimation of the VSL with labor market data, not just the hedonic method. A common feature of all these approaches is that they rely on revealed preferences: by observing people's choices in the labor market, we can infer how much they value different aspects of a job, and fatality risk in particular. Underlying this logic is the assumption that the worker is relatively

[8]For the EPA, the following hedonic method studies are used: Viscusi (2003); Viscusi (2004); Kniesner and Viscusi (2005); Viscusi and Aldy (2007); Aldy and Viscusi (2008); Viscusi and Hersch (2008); Hersch and Viscusi (2010); Scotton and Taylor (2011); Scotton (2013). In the case of the DOT, Viscusi (2004); Kniesner and Viscusi (2005); Evans and Kerry Smith (2008); Viscusi and Hersch (2008); Evans and Schaur (2010); Hersch and Viscusi (2010); Kniesner, Viscusi, and Ziliak (2010); Scotton and Taylor (2011); Kniesner, Viscusi, Woock, and Ziliak (2012).

well-informed, i.e., he understands the trade-off he is making. If the analyst and the worker differ in their understanding of the trade-off the worker made, the former will incorrectly infer the latter's preferences. And while monetary compensation is a relatively straightforward information about many jobs, e.g., $15 per hour, the same cannot be said about their fatality risks. As I show later, there is indeed empirical evidence that workers do not fully comprehend the fatality risks of their jobs.

This difficulty has not gone unnoticed in the literature. As mentioned above, Thaler and Rosen (1976) measure the fatality risk of a job, i.e., the risk the worker perceived when making his choice, by the occupation's fatality rate, as measured by number of deaths per hour worked in that occupation. Because it is not obvious that the two should be similar, the authors argued that "parties who voluntarily face such [occupational fatality] risks daily and as a major part of their lives, or production processes, have a special interest in obtaining reliable and objective information about the nature of the risks involved." Therefore, it is in the worker's best interest to ensure that his perception of the fatality risk of his job is accurate, and using the fatality rate as a proxy for worker's perception makes sense.

This potential difference between risk perceptions and the fatality rates we can measure was prominently a caveat in the early literature. It would often be pointed out that, conceptually, the analyst should be concerned with the worker's perception of the fatality risk. Mishan (1971) states that "what is strictly relevant to the analysis is not the change in the statistical risk per se but the person's response, if any, to such change. For the change in risk may go unperceived, and, if perceived, it may be improperly evaluated." In a literature review, Viscusi (1993) notes that empirically "the ideal risk measure would reflect subjective assessment of the fatality risk of the job by both the worker and the firm."

Along these lines, the paper that is most closely related to my research is Viscusi and O'Connor (1984), in which the authors try to assess whether workers' perceptions of risk were sensible and if they updated their beliefs in a Bayesian fashion. The authors conducted a survey among chemical industry workers and their reaction to new information regarding chemical components. The survey design allowed the authors to directly measure perceptions of risk, by asking workers whether they believed their job to be dangerous and a continuous variable that asked the workers to rank their jobs' risk relative to the U.S. average. The authors argue that this direct elicitation of perceptions is beneficial because "from the standpoint of the theoretical foundations of the compensating differentials theory, the wage-risk relationship should be driven by such subjective risk perceptions." In addition, they included a small experiment in their survey, in which the workers were presented with a scenario in which a new chemical compound—which had varying degrees of risk associated with it, and varied across experimental groups—would be introduced to their jobs. Then, they asked the workers for an updated risk assessment of the job, how much of a raise would be necessary for them to accept the job with the new chemical—if any, as one of the hypothetical chemical was harmless—or if they would rather quit. Their main conclusions were that employees were reasonably well-informed about the hazards of their jobs, because their wages and stated willingness to quit were consistent with a risk premium. Furthermore, they show that the introduction of new chemicals were duly incorporated into their beliefs about risks, as reflected by their updated perceptions of risk, wage premium and willingness to stay at the job.

Yet, a fatality risk premium will exist as long as the marginal worker requires such a premium, regardless of his (potentially inaccurate) perceptions of risk. Therefore I revisit this assumption in my work. Furthermore, the survey in Viscusi and O'Connor (1984) is limited to a single industry, which may not be representative of the broader labor market, and can only elicit willingness to quit. Neither is a limitation in my case, as I have a broad portion of the labor market in Brazil and I do observe actual quit behavior. Furthermore, I can emulate their learning experiment using a colleague's work-related death as the new information. Using actual quit behavior, I find evidence that workers seem to incorporate new information to their beliefs in a Bayesian fashion, which corroborates the findings of Viscusi and O'Connor (1984) in this regard.

The concern over whether decision makers are well-informed is not exclusive to labor markets, though. Ashenfelter and Greenstone (2004) leveraged a change in legislation that increased speed limits on highways to estimate the VSL. The authors provide anecdotal evidence suggesting that the legislators—the decision makers in their case—were fully aware that doing so would result in more traffic fatalities, as well as cutting travel time, i.e., they were aware of the trade-off they were making. Therefore, they argue, the decision to pass the legislation increasing speed limits reflected their preferences.[9] Pope (2008) showed that reminding people of publicly available information on airport noise had a large impact on house prices, suggesting that people may be poorly informed about the trade-offs they are facing even in highly consequential situations.

Yet, in practice most research using labor market data since then has largely assumed, like Thaler and Rosen (1976), that objective measures of risk such as fatality rates are a sensible proxy for workers' perceptions of risk.[10] And I believe this assumption is not warranted based on the early evidence on the topic. While it is certainly true that workers have the incentive to "obtain reliable and objective information" about their jobs, the evidence in Viscusi and O'Connor (1984) suggests that workers go through a period of experimentation, in which on-the-job experiences inform workers on the nature of their occupation. If anything, the fact that such experimentation is needed to begin with, despite the strong incentives workers already have, indicates that making these risk assessments may be a very difficult task. And as I show below, VSL estimates are not robust to uninformed decisions by workers.

Importantly for my purposes, the results in Viscusi and O'Connor (1984) suggest workers use job experience to incorporate new information on job hazards into their beliefs. This is consistent with the results in Rundmo (1995), which studies workers in offshore petroleum installations. Using a survey with nearly 1,000 employees on oil installations in Norway, Rundmo finds that those who experienced more work accidents tend to perceive their jobs as

[9] In fact, the authors argue that the legislators' decisions should, ideally, reflect the median voter's preferences.

[10] See the literature cited in footnote 8.

more dangerous and less satisfying, even if they had not been the accident victims themselves. Similarly, my empirical analysis relies on workers updating their beliefs on job hazards based on fatal work accidents happening to their coworkers.

A valid concern about my analysis is that workers may not necessarily be learning about these hazards, but simply be subject to availability bias (Kahneman & Tversky, 1972). Under this theory, rather than analyzing data in a rational way, people use an availability heuristic, in which they over-extrapolate the experience of events that are more readily available in their memory. For example, they could overestimate the risk of flying following some recent, widely covered plane crash. In my case, this could lead colleagues of someone that dies in a work-related accident overestimating the probability of falling victim to such an event as well. I believe this is not a problem in my study because, unless the availability heuristic makes people completely disregard their previous knowledge, the evidence I find would still be consistent with workers being unsure about the fatality risks of their jobs. In other words, the critical point is that a person with perfect knowledge does not need to incorporate new information to their beliefs, i.e., there should be no updating in beliefs under perfect information. But while I do find evidence of updating, the manner in which this updating occurs is less important for me.

Another difficulty when dealing with fatality risks is that deaths are likely to trigger an emotional response that could prevent people from evaluating probabilities properly, as suggested by Sunstein (2002). According to the author, an outcome that triggers a strong emotional response, like a death, may overwhelm the decision maker to the point that the actual probability of said outcome is disregarded when making decisions. This could lead to an overreaction, in which the decision maker behaves as if the emotional outcome is far more likely than it actually is. On the other hand, Sunstein (2002) also argues that, when a particularly bad outcome has a small probability, the bias may go the other way: in order to avoid cognitive dissonance, people choose to ignore the risk entirely and think they are "safe". In the context of fatal occupational accidents I am analyzing, both effects could be at play.

More broadly, Manski (2004) argues that the rational expectations framework is implausible, as it requires from decision makers a level of statistical sophistication often not found even among expert economists. And without rational expectations to discipline decision makers' behavior, it becomes virtually impossible to identify preferences based on their choices made under partial information alone. The reason is that, for any set of preferences, one can typically find a subjective probability distribution that could rationalize the observed choice. Therefore, Manski (2004) argues for the complementary use of elicited beliefs, i.e., directly asking decision makers about their beliefs and preferences.

Indeed, this is a common approach in the VSL literature,[11] commonly known as the stated preference approach, in contrast with the revealed preference approach that typically relies on labor market data. One benefit of the stated preference approach is that for many contexts in which the VSL is applied, the risk being evaluated is protracted, such as pollution exposure leading to respiratory diseases. On the other hand, the fatality risks we can study in labor markets are typically sudden risks to life, as the fatalities used to calculate risks in these studies are traumatic injuries, which typically involve acute exposure to energy, impact or absence of essentials such as oxygen. If people value mitigation of these risks differently, then perhaps labor market data cannot provide a meaningful VSL for many applications. Meanwhile, survey questions can present many flexible scenarios, which more accurately reflect the fatality risks potentially being mitigated in a particular application. Furthermore, the survey sample may be tailored to be representative of specific populations, which may be useful if the intervention being evaluated also has a limited scope. Finally, the risk trade-off being presented can be much clearer in an hypothetical survey question than in the labor market, so my concern about the fatality risks being properly understood by the decision maker is less relevant.

For all these benefits, I do believe supplementing my analysis with subjective probabil-

[11] Viscusi and O'Connor (1984) is one such example. For a more recent overview, see Cropper, Hammitt, and Robinson (2011); Lindhjem, Navrud, Braathen, and Biausque (2011); Robinson and Hammitt (2016).

ities information would be helpful. Yet, the data set I use does not have this information, and augmenting it is impracticable. I use an administrative data set that is collected by the Ministry of Labor in Brazil with the purpose of handling pensions and unemployment benefits. Each year, every establishment must inform the Ministry of all job spells they had in the previous calendar year, along with some features of the contract between employer and employee. Such information should be readily available on the contract itself or some internal registry of the firm, so in practice a single (or a few) employee is responsible for filing the data to the Ministry of Labor. It is not practical to include a survey on risk perceptions in this process, especially because studying workplace safety is not the goal of the Ministry when collecting these data to begin with. And a separate survey, specifically designed to elicit people's perceptions of risk, would present the downside of not allowing me to observe real, high stakes behavior that I do observe in the administrative data.

Finally, work-related fatalities are rare events, so the sheer magnitude of the data is instrumental for my empirical design, which is something I would not be able to do with a survey. Indeed, it has been growing increasingly common in the literature to leverage deaths as a source of exogenous variation. Azoulay, Graff Zivin, and Wang (2010) showed that the unexpected death of a scientist is detrimental to the productivity of their co-authors. Jäger and Heining (2019) used similar variation to show that these interactions between coworkers are not exclusive to highly skilled people. In fact, their findings presented a complication for me, because they document a direct effect of a colleague's death in the coworkers' decision to quit. I circumvent this difficulty by employing an empirical strategy similar to the one used by Sarsons (2019). Rather than defining the treatment group based on whether an event occurs—a death in this case—we define treatment as the specific circumstance in which an event occurs. In Sarsons (2019)'s work, she compared surgeons who lost a patient during surgery, but differ in their gender, showing that women had their professional reputation more adversely affected by such deaths than men. In my case, I compare coworkers who lose a work colleague, but only the treatment group does so in a way that is informative about job hazards, allowing me to emulate a learning experiment.

CHAPTER 2

EMPIRICAL EVIDENCE

This chapter is structured in the following way. In Section 2.1, I develop a simple conceptual framework to illustrate why we should be concerned about uncertainty over fatality risks, how learning can help dealing with it and how to test for it. In Section 2.2, I explain the data and my empirical strategy, and I show the main results in Section 2.3. In Section 2.4, I provide some auxiliary evidence and discuss potential alternative explanations. I assess the bias in the valuation of fatality risks in Section 2.5, and Section 2.6 concludes, with a brief discussion on policy implications.

2.1 Conceptual framework

My conceptual framework helps accomplish three goals. First, I illustrate how uncertainty over fatality risks can bias their valuation. Second, I discuss how on-the-job learning can help circumvent this bias. Third, I derive testable predictions that allow me to empirically assess whether there is uncertainty over fatality risks and, if workers learn about them on-the-job, how long it takes for them to do so.

A worker and a firm play a game to decide whether to continue their employment relationship or not. The worker moves last: after observing if the firm has experienced a work-related fatality, he decides whether to quit based on his perception of the job's fatality risk, the wage offer at hand and the outside opportunity available.[1] The firm has to decide on a wage offer before the idiosyncratic shocks are realized; she only knows the distribution of the value of the outside option, and that a work-related fatality may happen. Because the firm has a comprehensive understanding of the production technology, I assume that, unlike the worker, she perfectly knows the probability of such deaths, i.e., the objective fatality risk.

[1] I do not model the worker's utility if he dies. Instead, I directly assume that, although the worker is afraid to die, he may not die himself.

In particular, the firm knows that these deaths are extremely rare, as I detail in Section 2.2. Finally, the firm also knows the belief her worker has on these fatality risks.

To summarize, the timing of the model is:

1. Firm makes wage offer to maximize expected profit

2. Fatality signal and outside option value are realized

3. Worker chooses whether to quit the job

I assume that workers have preferences $u(w, \tilde{r}) = w - \gamma\tilde{r}$ that are additively separable between log wages w and perceived fatality risk $\tilde{r}$. The separability between pecuniary and non-pecuniary benefits is a common assumption in the literature[2], but here I explicitly acknowledge that perceptions of risk are what matters for the workers. Importantly, these perceptions may not necessarily be similar to an objective assessment we can make as outsider analysts. The wage is perfectly observed. Therefore, the expected utility of accepting a job is given by $\mathbb{E}[u(w, \tilde{r})|\mathcal{I}] = w - \gamma\mathbb{E}[\tilde{r}|\mathcal{I}]$, in which $\mathcal{I}$ is the information set of the worker.

The worker's decision depends on whether the expected utility of the job is higher or lower than the outside option, v. I assume that v is independently drawn from the distribution G, which is known by the firm. The worker quits if $\mathbb{E}[u(w, \tilde{r})|\mathcal{I}, x] < v$, in which x is the fatality signal, i.e., x is a binary variable denoting whether there is a death, so that $Prob(x = 1) = r$. This signal enters the worker's decision because it shapes his beliefs on the fatality risk. Then, the probability that a worker quits after observing the signal x and the wage offer w^* is

$$Q(\mathcal{I}, x) = 1 - G\left(\mathbb{E}[u(w^*, \tilde{r})|\mathcal{I}, x]\right) \tag{2.1}$$

I focus on the quitting probability because this is the one decision any worker always has at his disposal. Wages, for example, could change in reaction to the death of a coworker (Jäger & Heining, 2019), but perhaps the worker does not have the bargaining power to get any

[2]For example, Bonhomme and Jolivet (2009), Lavetti and Schmutte (2018), Card, Cardoso, Heining, and Kline (2018) and Lamadon, Mogstad, and Setzler (2019).

wage concessions from the firm.[3] When it comes to the decision to quit, though, the worker has full autonomy over this choice.

How does this all affect the firm's decision? I assume that the value of having an employed worker is constant, A. The firm's payoff depends on the worker's choice: if he quits, her payoff is zero; if he accepts the wage offer, her payoff is $A - w$. The firm knows what the worker initially knows $\mathcal{I}$, and his preferences. Therefore, the firm's problem is to choose w to maximize the expected profit:

$$\mathbb{E}[\pi] = (A - w)\left\{rG\left(\mathbb{E}[u(w, \tilde{r})|\mathcal{I}, x = 1]\right) + (1 - r)G\left(\mathbb{E}[u(w, \tilde{r})|\mathcal{I}, x = 0]\right)\right\} \tag{2.2}$$

Rather than solving the problem above, I use an approximation that relies on the fact that work-related fatalities are really rare. For example, there were 3.5 work-related fatalities per 100,000 full-time equivalent (FTE) years in the U.S. in 2018, which is fairly similar to what I observe in my Brazilian data, as shown in Section 2.2.[4] To better understand this number, note that in my main sample I restrict attention to establishments with at most 30 employees in a given occupation; at these fatality rates, it would take approximately 1,000 years for someone to expect a work-related fatality to occur among one of his coworkers. Therefore, I assume that $r \approx 0$, and because the firm knows what this value is, the expected profit (2.2) can be approximated by

$$\mathbb{E}[\pi] \approx (A - w)G\left(\mathbb{E}[u(w, \tilde{r})|\mathcal{I}, x = 0]\right)$$

This is helpful because, given that the firm knows $\mathcal{I}$, the optimal wage is no longer a function of r. That is, the optimal wage here is a function of the information set the worker already had, $w^*(\mathcal{I})$. Therefore, the worker cannot use the wage offer he receives to try to infer

[3] Indeed, I show in Section 2.4 that a work-related death does not seem to have any impact on surviving coworkers' wages, but the comparison to Jäger and Heining (2019) is not immediate, as I explain later.

[4] A full-time equivalent year consists of 40 hours of work per week, 50 weeks per year, for a total of 2,000 hours per year.

anything about the true risk r beyond what he already knew. In other words, w^* itself does not directly enter the information set the worker uses to form his beliefs when deciding to quit, which greatly simplifies the problem. If this was not the case, the worker would try to learn something about r through the wage offer he received w^*. But then the firm would anticipate that and incorporate those incentives in her decision to set w^* to begin with, and so on, which would make the problem intractable.

Therefore, x is the only new information regarding the fatality risk the worker receives before making his decision. In order to illustrate its impact on the learning process, I now put more structure on the worker's beliefs. Let $\tilde{r} \sim Beta(a, b)$, so its density f is such that $f(\tilde{r}; a, b) \propto \tilde{r}^{a-1}(1-\tilde{r})^{b-1}$, for $\tilde{r} \in [0, 1]$ and $a, b > 0$. Under this specification, the information set $\mathcal{I}$ can be characterized by the prior's hyperparameters (a, b), and assume further that people have unbiased priors, so that the prior mean is r, i.e.,

$$\mathbb{E}[\tilde{r}; a, b] = \frac{a}{a+b} = r$$

Under this restriction, we have that $Var\left(\tilde{r}|a(a+b)^{-1} = r\right) = \frac{r(1-r)^2}{b+1-r}$, which is decreasing in b. This means I can characterize the uncertainty over the beliefs with the hyperparameter b: the larger b is, the more precise is the prior.[5]

With this functional form, the posterior distribution is $\tilde{r} \sim Beta(a + x, b + 1 - x)$, which means that the posterior mean is given by $\mathbb{E}[\tilde{r}; a, b, x] = \frac{a+x}{a+b+1}$ and I can re-write (2.1) as

$$Q(a, b, x) = 1 - G\left(w^* - \gamma \frac{a+x}{a+b+1}\right) \tag{2.3}$$

So the probability of quitting is increasing in the posterior mean $\frac{a+x}{a+b+1}$, which is increasing in x, hence we should see a larger share of persons quitting the job following a colleague's work-related death; that much is obvious: the more dangerous the job is perceived to be, the

[5] The Beta distribution is a conjugate prior of the Binomial distribution, so it is commonly assumed in setups in which the signal is binary like mine (e.g., Israel, 2005). Intuitively, it describes the beliefs a Bayesian agent would form over the probability of a binary outcome if they had observed a successful draws and b failures.

more likely it is that an outside option is able to lure the worker out of his current job. The more relevant point is to describe how experience affects the reaction to new information.

First, note that workers' posterior beliefs about risk $\mathbb{E}[\tilde{r}|\mathcal{I}, x]$ have a tendency to be lower than r on average, despite their unbiased priors and signals. This is because people are more likely to quit if they observe a work-related fatality. More formally, the unconditional average posterior mean is given by $(1-r)\mathbb{E}[\tilde{r}|\mathcal{I}, x=0] + r\mathbb{E}[\tilde{r}|\mathcal{I}, x=1] = \frac{a+r}{a+b+1}$. Substituting the restriction of unbiased priors $\frac{a}{a+b} = r$, we have that $\frac{a+r}{a+b+1} = r$, so the unconditional average posterior mean is also unbiased. Yet, we already know that workers who receive a work-related death signal should be more likely to quit. Therefore, the average perception among the workers that do not quit—and hence would be observed in data—is lower than the objective risk r.[6]

This systematic difference between what we as analysts can measure—r—and what workers actually base their decisions upon—$\mathbb{E}[\tilde{r}|\mathcal{I}, x]$—creates a bias in the valuation of the fatality risks. For example, we could observe in the data a worker willing to accept \$$X$ for an objectively measured risk r, when in fact his real decision was to accept \$$X$ for a perceived hazard $\mathbb{E}[\tilde{r}|\mathcal{I}, x] < r$. In this scenario, it is not clear that the worker would have accepted the objective trade-off, but that is the inference we draw from the data if we use these objective proxies for the worker's perceptions. Therefore, valuation methods that use r to proxy for $\mathbb{E}[\tilde{r}|\mathcal{I}, x]$ tend to undervalue fatality risks.

Point #1: Uncertainty and job choice lead to bias in the valuation of fatality risks. Once we allow for perceptions of fatality risk to vary across people, job choice ensures that workers pick jobs for which their beliefs are particularly kind. Therefore, valuation methods that rely on objective assessments of risk tend to undervalue them, because the actual trade-off workers made was perceived by workers to be a better deal than it seems from an objective

[6]The direction of the bias relies on my assumption that priors are unbiased. Without it, the VSL would still be biased, but the direction would not be clear. The model predictions on the learning process and the way people should react to dire signals would remain the same though. As I find evidence of undervaluation of fatality risks, I believe the assumption of unbiased priors is a nice way to discipline the model and get clearer predictions from it.

standpoint.

So in general it is not advisable to use objective measures of risk as proxies for the perception of workers, because they are systematically different, leading to a bias in valuation. Yet, there is an exception if we expect people to be learning with experience. For a worker that is very experienced—a high value of b, conditional on the prior being unbiased (see footnote 5)—the posterior mean should be virtually the same as the prior mean, regardless of the signal they observe: $\frac{a+x}{a+b+1} \approx \frac{a}{a+b} = r$, for large b.[7] Therefore, for sufficiently experienced workers, the signal x should barely change their beliefs. Moreover, the difference between the objective fatality risk and their perceptions of it should be negligible.

Point #2: On-the-job learning means experienced workers' beliefs are close to the objective level of risk. This is simply a Bayesian maxim applied to my context: the more information you have, the more confident you should be in your beliefs. Therefore, experienced workers' perceptions should not be as easily swayed by new information.

So even though I do not observe perceptions of risk, I know that I can approximate them relatively well with objective proxies so long as workers have enough experience. Then, the question is how much experience is necessary for this approximation to be sensible. In other words, at which point is the posterior belief relatively unchanged by new information? I can use the reaction to a work-related fatality in terms of quits as a test for whether the updating in beliefs is strong or not, as we can see from the quitting probability (2.3). Inexperienced workers, whose beliefs are still pliable, should quit their jobs at a higher rate after observing these fatal accidents. In addition, if this dramatic signal does not elicit a higher quitting probability, then I can assume that workers' perceptions are relatively stable, and thus close to the objective level r.[8]

Point #3: The reaction to a work-related death can inform me on both uncertainty and learning. If workers are not certain about the fatality risks of their jobs, a work-related

[7] More formally, substitute the unbiased prior condition $\frac{a}{a+b} = r \Rightarrow a = \frac{br}{1-r}$ on $\frac{a+x}{a+b+1}$ and take the limit as b grows larger to see that it converges to r.

[8] I consider the case of a death happening $x = 1$, instead of the case with no death $x = 0$, because the former is when the contrast between new information and prior beliefs should be starker.

fatality is likely to shape their beliefs, which should induce some of them to quit their jobs. Moreover, if there is on-the-job learning, this reaction should be particularly noticeable among the less experienced workers, but should weaken as they grow more experienced.

This differential quit reaction is something I can test in my data, and it is informative of whether there is uncertainty around fatality risks and whose perceptions are likely to be stable. It follows from point #2 that the increase between prior and posterior means is smaller if workers are more experienced, so their reactions in terms of quitting probability are more likely to be muted. Conversely, first impressions matter, so the death signal increases the worker's perceptions of risk more if they have less experience, which should be reflected in their decision to quit.

Ultimately, I am concerned with the impact uncertainty may have on the valuation of fatality risks, which is directly connected to the marginal disutility of expected fatality risk γ, as I show in Section 2.5. Differentiating (2.1), I can identify γ with:

$$\gamma = -\frac{\partial Q/\partial \mathbb{E}[\tilde{r}|\mathcal{I}, x]}{\partial Q/\partial w} \tag{2.4}$$

This is very similar to the method proposed by Gronberg and Reed (1994), with the difference being that the numerator uses the worker's mean belief, rather than the objective level of risk. As discussed above, this is unlikely to be an issue if workers have accurate perceptions of risk, but as I show Section 2.3, this does not seem to be the case in Brazil. Yet, I also show that they seem to be learning with experience, so I know I can use the objective level of risk to proxy for the perception of some workers, according to point #3. So in Section 2.5, I compare a valuation estimate that disregards the potential bias created by uncertainty over fatality risks, with one that focuses on a sub-sample that is likely to be well-informed—workers that have a minimum level of job tenure.

2.1.1 Summary

The conceptual framework above has three key points I want to reiterate:

Point #1: Uncertainty and job choice lead to bias in the valuation of fatality risks. Once we allow for perceptions of fatality risk to vary across people, job choice ensures that workers pick jobs for which their beliefs are particularly kind. Therefore, valuation methods that rely on objective assessments of risk tend to undervalue them, because the actual trade-off workers made was perceived by workers to be a better deal than it seems from an objective standpoint.

Point #2: On-the-job learning means experienced workers' beliefs are close to the objective level of risk. This is simply a Bayesian maxim applied to my context: the more information you have, the more confident you should be in your beliefs. Therefore, experienced workers' perceptions should not be as easily swayed by new information.

Point #3: The reaction to a work-related death can inform me on both uncertainty and learning. If workers are not certain about the fatality risks of their jobs, a work-related fatality is likely to shape their beliefs, which should induce some of them to quit their jobs. Moreover, if there is on-the-job learning, this reaction should be particularly noticeable among the less experienced workers, but should weaken as they grow more experienced.

The first point is the cause for concern; the valuation of fatality risks is an important input in many policy decisions, which may be misguided by a biased estimate. The second point is how I can address the previous one, allowing me to devise an improved valuation for the fatality risk; by focusing on more experienced workers, I can be more confident that my objective measures of risk actually resemble these workers' beliefs and, hence their actual decisions. And the third point provides me a way of testing my model and timing the learning process.

2.2 Data description and empirical strategy

I want to test the following prediction: a work-related fatality—due to its informational content about occupational hazards—should induce workers to quit their jobs, but this effect should grow weaker as they accumulate more tenure. Before doing so, I describe the data I use, as it provides context and guidance for my empirical analysis. Then, I discuss the empirical strategy itself, its limitations and how it can elucidate some of the questions I raised with my model.

2.2.1 The data

I use the Relação Anual de Informações Sociais (RAIS) between 2007 and 2017, a matched employer-employee Brazilian administrative data used by the Ministry of Labor to administer formal employment benefits. A caveat to this data set is that it does not cover the relatively large informal labor market sector in Brazil, so my results should be interpreted with caution (Ulyssea, 2018). From now on, the labor market and its agents are implicitly understood to be formal, unless stated otherwise.

Every year, virtually all employers in Brazil are required to submit information on their entire workforce to the Ministry of Labor; each establishment is responsible for informing the Federal Government on every active labor contract they have had at some point in the previous year. They must inform the date in which the contract started and potentially ended, the reason for termination, average monthly wage over the period, hours contracted, some demographic characteristics of the worker and occupation codes.[9] Both workers and establishments can be identified by a national identification number, with the latter also identifying the firm to which an establishment belongs.

The key variable here is the reason a job spell was terminated. It provides two important pieces of information: first, if the spell is terminated because the employee has died, I also observe some details on the circumstances of the death; second, terminations that are

[9] I report every monetary value in 2016 US$. See Appendix D for details.

voluntarily initiated by the worker are explicitly categorized, which I call quits. Both are intrinsically connected to my model: work-related deaths not only allow me to estimate the objective fatality risk of an occupation, but also provide the observable signal I predict should precipitate quits.

Given the importance of this information, it is worth discussing it in detail, with a special attention to concerns about measurement error. Start with quits. The reason why a job spell was terminated is categorized in several different ways: retirement, employee's death, transfer, contract expiration, among others. I define quits as "*terminations without cause initiated by the employee*"; similarly, layoffs are "*terminations without cause initiated by the employer*".

In my simple model, quits are the only type of separation that can exist, because firms are always willing to hire the worker for the wage they offer. Yet, a more prevalent assumption is to allow for wage bargaining, so that separations are efficient, i.e., employees are willing to separate if and only if their employers are also willing to do so (Jäger, Schoefer, & Zweimüller, 2019). If this was the case, the distinction between quits and layoffs in the data would be moot. Yet, empirically this is not the case. As shown in the Appendix Table A.1, persons that quit are more likely to find new jobs and, conditional on doing so, find them faster and with higher wages. These patterns indicate that separations labeled as quits in the data happen to workers that are in a relatively advantageous position in the job market, lending credence to the idea that they are indeed voluntary.

Another reason to believe quits are voluntary separations by the part of the employee is related to the institutional environment in Brazil. Many employment benefits administered by the Federal Government are not available to persons that quit their jobs.[10] Therefore, while in principle firms can misreport layoffs as quits, they risk being liable to fines and retroactive compensation if they are denounced by workers, who have the financial incentive to do so.

[10]For example, unemployment insurance, severance pay or access to a compulsory savings account administered by the Federal Government.

The same applies to the case of an employee's death. Such events can be labeled in four mutually exclusive categories: caused by work-related accident—which is the category I use to define my treatment group; caused by injury in the commute to work; caused by professional disease; and a residual category that is simply described as "death", presumably because it is not directly related to the job. A natural concern is that establishments would misreport work-related fatalities, because they may be liable to sanctions or enhanced regulatory costs, especially if the establishment was negligent in its safety procedures. But just like with quits, workers have a strong incentive to monitor their employer. Whenever a work-related accident occurs, the establishment is mandated to inform Social Security through a Work Accident Communication (Comunicação de Acidente de Trabalho), which details the accident and is a necessary document for requesting public pensions. Firms that do not inform Social Security on time are liable to fines. Importantly, the worker or a relative—who have a direct stake on the matter, as they are often eligible recipients of such pensions—are also allowed to inform Social Security directly, but that does not exempt the firms from fines, creating an incentive for employers to quickly report these cases.

Another way to evaluate the quality of my data is to compare the number of fatalities reported in Brazil and in the U.S., through the Census of Fatal Occupational Injuries (CFOI). In particular, I compare occupation-specific fatality rates, because this is a useful way to measure the objective level of fatality risk of a job, r.[11] I focus on occupations because they are the most precise description of job responsibilities I have in the data, so people with the same occupation are likely to face the same job hazards. Following the U.S. Bureau of Labor Statistics (BLS) methodology, the fatality rate of an occupation is the number of deaths caused by a work-related accident divided by the number of hours worked, for a given occupation.

[11] While occupation-specific fatality rates are useful to assess the quality of the RAIS data, for my empirical exercise described below, I would ideally calculate establishment-specific fatality rates. But as previously mentioned, work-related fatalities are rare, so establishment- or even firm-specific fatality rates would be zero for the vast majority of firms, and most of its variation would be driven by size, rather than number of deaths.

Here I present a brief summary, but a detailed analysis on the calculation of fatality rates and the comparison between RAIS and CFOI can be found in the Appendix C. The overall fatality rate in my data is 2.71 per 100,000 FTE years. This average hides substantial heterogeneity across occupations. The first, second and third quartiles of the fatality rate distribution are 0, 1.25 and 3.55, respectively; the occupations with the highest fatality rates are close to 100 per 100,000 FTE years. There were no work-related fatalities for 141 of the 555 occupations, while 133 of them had at least one work-related fatality per year, on average. More importantly, the fatality rates I calculate are remarkably similar to the CFOI ones, especially considering how rare these events are and the difficulties for accurate estimation that presents. Unsurprisingly, the larger differences tend to occur in occupations with fewer hours worked. To the extent that these rates reflect some technological constraints of the job that are similar between the U.S. and Brazil, this is evidence that work-related fatalities are reported at similar rates in both countries.

Hence, the RAIS does seem to provide reliable information on work-related fatal accidents. These are rare and noteworthy events, so we can expect them to leave a strong impression on those witnessing them. Moreover, it is hard to argue they could have been predicted, otherwise these accidents would have been avoided.

2.2.2 The empirical strategy

If workers do not fully comprehend the fatality risk of their jobs, my model predicts that they are more likely to quit following a work-related death. Furthermore, if they learn about these risks on-the-job, this effect would be particularly important among workers with relatively low tenure, because their limited experience implies their perceptions are more malleable—and also less likely to be close to an objective assessment of such risks.

The obvious caveat is that a dead worker cannot quit. But his death may still inform those around him. More specifically, I focus on the behavior of persons that work in the same establishment and have the same occupation as the decedent. The restriction on occupation is related to the information I am expecting such deaths to convey: as mentioned above, occupation is the best proxy I have on the nature of the work performed by people, including the fatality risks they are exposed to. Therefore, a fatal work-related accident is likely to reveal or make salient dangers that are common to the people with the same occupation. Moreover, in order for the signal to affect people's perceptions, it must be received by them. Hence, I also restrict attention to persons in the same establishment as the deceased worker, because they are likely to become aware of his death.

I also need a counterfactual for these coworkers, a way to assess how they would have behaved in the absence of the risk signal. This presents a difficulty because the impact of a work-related fatality goes well beyond the risk signal they provide. At the very least, the establishment also faces some disruption to its production process, through grief and a sudden decrease to its labor force, and that could affect the surviving workers' probability of quitting directly (Jäger & Heining, 2019).

I address this by using the death of employees under circumstances unrelated to work. As described above, when the death of the employee is used as the cause for a job spell termination in the RAIS, it can be classified in four mutually exclusive ways: the death was caused by a work-related accident, which I am using as the signal about risk; the death happened in the commute to work or the death was caused by a work-related illness, both of which are related to work, but do not satisfy the CFOI criteria for work-related fatality; and a residual category simply called "death". Even though it is not explicitly stated, this latter category is clearly meant to encompass deaths under circumstances unrelated to the job. Therefore, while I can expect grief and productive disruptions to ensue in these cases as well, there is no reason for the surviving coworkers to perceive their jobs as more dangerous. Hence, I use the coworkers of persons that died in a manner unrelated to work as my control group. From now on, I refer to deaths in a work-related accident simply as "work-related", while my control group is defined by deaths denoted "unrelated (to work)" or "other deaths".

Conceptually, my empirical strategy is the same as the one employed by Sarsons (2019).

Rather than defining treatment based on whether an event occurs, we define it as the specific circumstances under which the event occurs. In my case, the only systematic difference between the coworkers are the circumstances in which their colleague died: it was either due to a work-related accident or unrelated to work. A difference-in-difference specification arises naturally: if risk perceptions are malleable, coworkers in the treatment group should be more likely to quit after the fatality, presumably because, unlike a death unrelated to work, a work-related death conveys a grim message about the job hazards they face. Furthermore, if the control group is indeed a sensible counterfactual, we should also expect no systematic difference in their behavior prior to the death.

2.2.3 Sample restrictions

I focus on establishments that have a single death that is caused by either a work-related accident or unrelated to work between 2007 and 2017. Doing so allows me to clearly define the pre- and post-period in my analysis, as each establishment in my sample experiences a single death. That is not to say that repeated deaths would be uninformative; in fact, Israel (2005) used repeated signals to document evidence of learning in the auto-insurance market. The reason I focus on single deaths is that (luckily) work-related deaths are very rare events, and few establishments in my data experience more than one such case. The 12,853 work-related deaths in my sample occur in 10,530 distinct establishments, among which 9,495 of them experience a single case, so I do not have enough data to leverage this variation.

The few establishments in which multiple work-related deaths do occur also tend to be very large, and those can be hard to study. Part of the reason a work-related death is such an effective risk signal is because they are very rare, and therefore unexpected news for the coworkers. In very large establishments, the information about such accidents may not spread far enough to reach most coworkers and, even if it does, the death may not be unexpected, given the magnitude of its workforce.[12] In order to avoid this issue, I restrict my sample to establishments in which there are at most thirty coworkers in the same occupation as the decedent, in the year of the death.[13]

Finally, my treatment and control groups by construction are systematically different in an important dimension: the fatality risk of their jobs. If the quit reaction to a death—irrespective of its cause—is a function of how risky the job is, the difference-in-difference estimator would be biased. For example, the death of an employee may imply a higher workload for his surviving coworkers until a replacement is found, which may be particularly undesirable for more dangerous jobs; if that is the case, workers in more dangerous occupations would be more likely to quit following a death even if there was no learning about risk, and regardless of the circumstances of the death. Therefore, I use inverse probability weighting to re-weigh the deaths unrelated to work and replicate in them the occupation—and hence risk—distribution found among the work-related deaths. Because deaths unrelated to work are far more numerous, they cover all the occupations in which a work-related death occurs.

My sample therefore consists of workers in a given occupation-establishment combination, defined by the death of a worker in either a work-related accident or an unrelated one. This decedent must have had between one to thirty coworkers with the same occupation at the time of his death. I also restrict the sample to cohorts between three years before and two years after the death. Finally, I weigh occupations so that their weighted distribution among the workers that die in a way unrelated to work is the same as the occupation distribution among the workers who die in a work-related accident. Therefore fatality risks are balanced across treatment and control, because they are defined at the occupation level.[14]

[12]For example, the largest establishment-occupation with a single death in my data has more than 10,000 job spells in a single year. For a group this large, the unusual thing in a relatively risky occupation would be for no fatalities to occur. Therefore, even the interpretation of a work-related fatality as a dire signal may no longer be warranted in very large establishments.

[13]Using the main difference-in-difference specification, which I describe below, I show in the Appendix Figure B.1 that my results are robust to changes in the size threshold, as long as the establishment is not too large. At that point, there is no statistically significant reaction to a work-related death, which is consistent with the signal not being properly observed by all coworkers or even losing its edge.

[14]I have constructed my weights such that the weighted distribution of occupations across the decedents are the same. As the number of coworkers they had may vary, the fatality risk among the coworkers is not

Table 2.1: Characteristics of coworkers in year immediately before death

Type of colleague's death	Unrelated to work		Work-related accident	t-stat
	(1)	(2)	(3)	(3) - (2)
Hourly wage (2016 US$)	3.031	3.191	3.266	1.286
High School	0.479	0.463	0.445	-2.375*
College	0.052	0.029	0.028	-0.382
Male	0.741	0.891	0.922	7.473*
Age	35.141	36.613	35.803	-6.218*
Fatality rate (per 100,000 FTE)	4.671	12.051	12.216	0.674
Job experience	2.344	2.077	1.951	-2.763*
Quit	0.080	0.079	0.083	1.67
Layoff	0.228	0.240	0.251	2.125*
Weighted by occupation	No	Yes	NA	
# of establishments	144,185		5,970	150,155
Observations	919,620		43,079	962,699

*p<0.05. Sample averages over workers by establishment treatment status, in the year immediately before the death. Columns (1) and (2) refer to the control group, with the latter being re-weighted; column (3) refers to the treatment group. After the re-weighting, the control group is much more similar to treatment, even though this is not a necessary condition for the difference-in-difference estimator. Weights are set so the occupation distribution across establishments in the control group is the same as in the treatment one. Job experience is measured as number of years since spell started, rounded down. The t-stat is calculated with a regression of variable on an indicator of whether the death is work-related, with weights and standard error clustered at the establishment level.

After all these restrictions, I am left with 5,970 work-related deaths, and 144,185 deaths unrelated to work. In Table 2.1, I show the sample averages of the coworkers in the occupation-establishments I track, exactly one year before the death. Balance in observable characteristics across treatment status is not a necessary condition for the difference-in-difference estimator to be unbiased, but it would be reassuring if the groups were similar. Without the weighting (Column 1), there are some stark differences: coworkers in the treatment group (Column 3) are much less likely to be college educated (2.8% versus 5.2%), more likely to be male (92.2% versus 74.1%) and work in occupations with much higher fatality rates (12.22 versus 4.67 per 100,000 FTE). After weighting (Column 2), these differences become much smaller, but some of them are still statistically significant, which is not surprising given that there are nearly a million observations here. Still, the differences in sample averages are small in magnitude and are not statistically significant for my most important variables, which are "quit" and "fatality rate".

2.2.4 Regression specification

The prediction I am testing is that a work-related fatality should precipitate an increased likelihood of workers quitting, especially among relatively inexperienced ones. Therefore, I compare the quitting pattern of coworkers in an establishment in which the death is work-related versus those in which the death is unrelated. The regression model I use is the following:

$$q_{ijt} = \sum_{k(j,t)=-3}^{2} \gamma_{k(j,t)} D_{k(j,t)} + \underbrace{\sum_{k(j,t)=-3}^{2} \beta_{k(j,t)} D_{k(j,t)} W(j)}_{\text{Work-related relative to Other deaths}} + \theta' x_{ijt} + \xi_t + \xi_j + \varepsilon_{ijt} \quad (2.5)$$

in which q_{ijt} is an indicator of whether the coworker i, who works at establishment j in calendar year t, quits his job—although later I use this specification for different outcome

exactly the same.

variables; $D_{k(j,t)}$ is an indicator that the death in establishment j has occurred k years after the year t, i.e., $k(j,t) = -3$ indicates that the death at establishment j occurs at calendar year $t+3$; and $W(j)$ is an indicator for whether the death is work-related.

The regression model also includes worker age and its square, education, race, gender, and fixed effects for the establishment and treatment-specific calendar year. Moreover, I include treatment and relative-time interactions with worker age, in order to try to tease apart tenure and age effects. The standard errors are clustered at the establishment level and I use the occupation weights I described above to ensure my results would not be driven by a higher propensity to quit in riskier jobs. Given that I have eleven years of data, I kept the time window around the death relatively short, so it spans between three years before and two years after the death.

The parameters of interest are $\beta_{k(j,t)}$, which are associated with the interaction between $D_{k(j,t)}$ and $W(j)$. They indicate the difference in the probability of a coworker quitting k years away from a work-related death, relative to an unrelated one. I normalize β_{-1} to zero, so according to my model, I should expect $\beta_{k(j,t)} > 0$ if $k(j,t) \geq 0$, because the work-related death makes the coworkers perceive their jobs as more dangerous, making them more likely to quit. Moreover, if deaths unrelated to work are indeed a sensible counterfactual, I should expect $\beta_{k(j,t)} = 0$ if $k(j,t) < 0$, because there should be no systematic difference between treatment and control prior to the death.

Furthermore, I estimate the model using samples restricted to workers with more job tenure because I also expect any increase in the probability of quitting to fade away as workers accumulate more experience. In particular, if workers' perceptions of risk are stable because they are built upon a rich array of previous experience, then they are also likely to be close to the objective fatality risk.

2.3 Main results

In this section, I provide the main empirical evidence I use to conclude that workers are not certain about the fatality risks of their jobs, but learn about them through experience. I present my estimates of Equation (2.5), and show that they conform to my predictions: coworkers are more likely to quit when a work-related accident kills a colleague, but that only happens among relatively inexperienced ones. In the next section, I discuss some auxiliary evidence and potential alternative explanations for the results below.

In Figure 2.1, I plot my estimates for β_k of model (2.5) without any tenure restriction. As expected, in the year of the death, there is a spike in quits in the work-related establishments, relative to the other ones. In that year, the probability of quitting increases by 1.29 p.p., which is a 16% increase off the baseline share of 8.06%. This reaction is indicative that workers do not know the fatality risk of their jobs for sure. The increase in quits is consistent with workers updating their assessment of risk for the worse, and voluntarily leaving their jobs. This spike in quits does fade away quickly over time, with the probability of a worker quitting no longer being statistically significant a single year after the death and virtually zero after two years. Reassuringly, there does not seem to be a systematic difference in the probability of quitting between treatment and control groups before the death, which is consistent with my control group being a sensible counterfactual for the treatment one.

If workers are indeed uncertain and use their on-the-job experiences to hone their perceptions of risk, my model predicts this quit reaction should be weaker for persons with longer job tenure. In order to investigate that, I re-estimate the model in Equation (2.5), but removing the workers with the least experience from the sample. My measure of tenure is the rounded down number of years since the job spell started, so I removed every coworker with zero years of experience. That is, I remove anyone whose job spell started in the same calendar year t in which I am observing their decision to quit $q_{i,j,t}$.

As I increase the lower bound on tenure, I expect workers' perception to grow increasingly stable: employees should eventually accumulate enough work experience that their percep-

Figure 2.1: Relative quits - Work-related versus Unrelated

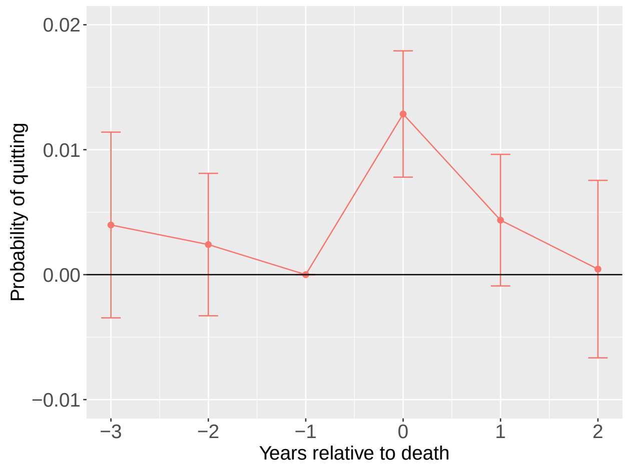

The graph plots the β_k coefficients estimated with the model described in Equation (2.5). I normalize β_{-1} to zero.

tions are barely budged by new fatality risk signals. This would be reflected in a flattening of the probability of quitting around the time of death.

In order to facilitate the comparison, in Figure 2.2, I plot these new results in red along with the previous estimate of β_k's, i.e., the ones in Figure 2.1, as a shaded gray line. There is still a significant spike, with an increase in the probability of quitting of 0.81 p.p. in the year of the death, but it is less pronounced than before, as we would expect if there was learning. The overall pattern remains similar, which indicates that workers' perceptions of risk are still pliable in this sub-sample. Moreover the pre-trends still indicate that my control group is a proper counterfactual for the treatment one.

So I estimated (2.5) using increasingly stricter lower bounds on tenure, expecting the increase in quits to disappear and combined them in Figure 2.3. As expected, we can see the probability of quitting flattening around the time of the death. With two years of tenure as the lower bound, the increase of 0.40 p.p. in the probability of quitting in the year of death is no longer statistically significant at a 5% confidence level, although it is so at 10%. Once the tenure restriction is at least three years, there is no longer a discernible increase in the probability of quitting in the year of death. If anything, the probability of quitting

Figure 2.2: Relative quits - Work-related versus Unrelated, at least 1 year of tenure

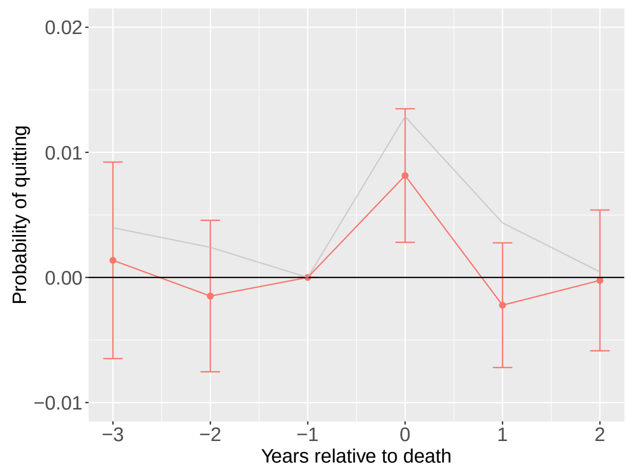

The graph plots the β_k coefficients estimated with the model described in Equation (2.5). I normalize β_{-1} to zero.

seems to decrease slightly in the following years, though those coefficients are typically not statistically different from zero. Among this group, the baseline probability of quitting still is 2%, so the similarity across circumstances of death is not a mechanical effect that could arise if experienced workers never quit to begin with.

Moreover, there is no qualitative difference between workers with at least three, four or five years of tenure. This is what we would expect if it took about three years of job experience for perceptions of fatality risk to become stable. After this point, even a dramatic signal like a work-related fatality no longer budges perceptions in a meaningful way, because priors are based upon a vast enough collection of previous knowledge. And again, it is reassuring to see that regardless of the tenure restriction I apply, there is no evidence of a differential trend in the probability of quitting between treatment and control groups before the death happened.

Based on these findings, I use three years of tenure as my threshold for the experience required to reach stable perceptions of fatality risk. Note that, under the assumption that on-the-job experience provides unbiased signals, my model predicts that perceptions can only be stable around the true fatality risk. In other words, the readily available measures

Figure 2.3: Relative quits - Work-related versus Unrelated, increasing tenure threshold

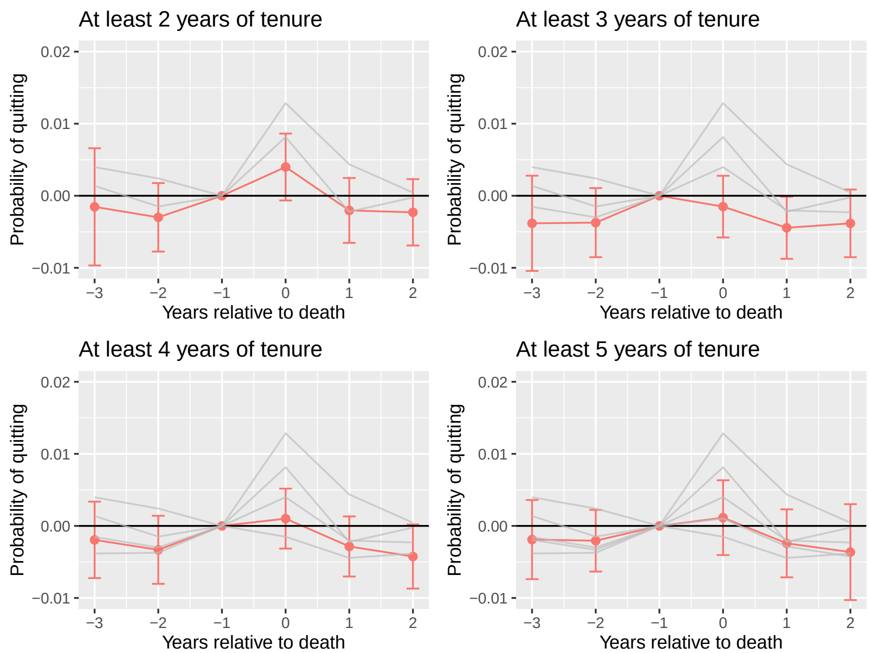

The graph plots the β_k coefficients estimated with the model described in Equation (2.5). I normalize β_{-1} to zero.

of objective risk should be proper proxies for workers' perceptions if they are stable, i.e., if they have at least three years of job experience.

Figure 2.4 summarizes the main finding here; it contains the estimated β_k's for two different sub-samples: one with workers with two or less years of tenure and another with those with more than three years. Reassuringly, before the death happens, the behavior across treatment and control is similar, with no indication that there is a pre-trend in their propensity to quit, regardless of their tenure.[15] Once the death happens though, the treatment group receives a dramatic negative signal about the fatality risk of their jobs. Yet experienced and inexperienced workers react very differently to it, in terms of quitting propensity: while there is no significant reaction among high tenure workers, the probability of quitting spikes by 1.79 p.p. among low tenure workers in the year of the death, which is a 18%

[15] In the Appendix Table A.2, I show that the balance across observable characteristics between treatment and control holds even conditional on tenure.

Figure 2.4: Relative quits - Work-related versus Unrelated, low and high tenure

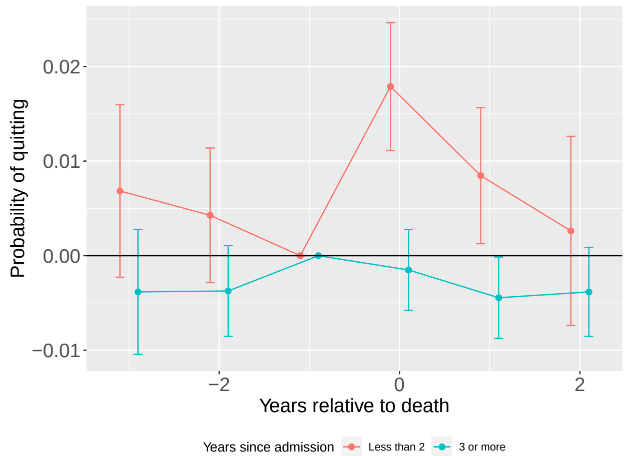

The graph plots the β_k coefficients estimated with the model described in Equation (2.5). I normalize β_{-1} to zero.

increase relative to the baseline probability in this group.[16]

My interpretation is that these signals influence job mobility through their ability to shape workers' perceptions. High tenure employees have accumulated a lot of experience, so they have an accurate understanding of the hazards of their jobs; this means their perceptions are not easily shaken, not even by the death of a colleague in a work-related accident. Accordingly, we do not see a meaningful difference in the probability of a high tenure worker quitting after such deaths. The same cannot be said for relatively inexperienced workers. That is not surprising; after all, first impressions matter. In fact, the important role of early experiences in shaping people's perceptions have been documented in settings as diverse as car insurance (Israel, 2005) and classical music (Kim, 2020). In my case, this means that workers who experience a very dire event early on in their job tenure are much more likely to quit these jobs.

[16]In Appendix F, I do a simple calibration exercise of my toy model to show that the quit magnitudes observed here are consistent with workers receiving 2.23 signals per week, i.e. the number of x draws, which is a fairly sensible number.

2.4 Additional evidence

The empirical evidence provided in Section 2.3 conforms with the idea that workers are not certain about the fatality risks of their jobs, but take roughly three years to learn about them with on-the-job experience. Yet, this interpretation is not entirely justified at this point, so now I provide some additional evidence to try and enhance the confidence in my interpretation. If the reader is already convinced that workers take three years of tenure to fully appreciate the fatality risks of their jobs, this section can be skipped without much of a loss. Alternatively, the reader can go directly to the subsection they deem most relevant, as they are self-contained.

First, I discuss the key assumption that the difference between deaths unrelated to work and work-related ones is mainly the informational content of the latter. In particular, I show that the quitting patterns observed above do not seem to be driven by an *objective* increase in fatality risk. Then, I discuss how preference and fatality risk heterogeneity could affect the learning interpretation of my results. I also discuss the effect—or absence of one—of work-related deaths on wages and the limited role human capital interactions could have in my results. Finally, I show that workers who quit following a work-related death struggle to find subsequent employment, which is consistent with an increased perception of risk.

2.4.1 Accidents objectively increase risk

I have compared unrelated to work-related deaths under the assumption that the difference between them is merely the signal the latter sends about the fatality risk of these occupations. This is not necessarily the case; in particular, the fatal accident may *objectively* increase the risk of performing the job, so surviving coworkers would choose to quit even if—and perhaps because—they were fully informed.

The most natural approach would be to check if the likelihood of a work-related fatality increases after one such occurrence. Yet, as discussed in section 2.2, this strategy is not

Figure 2.5: Nonfatal injuries - Work-related versus Unrelated

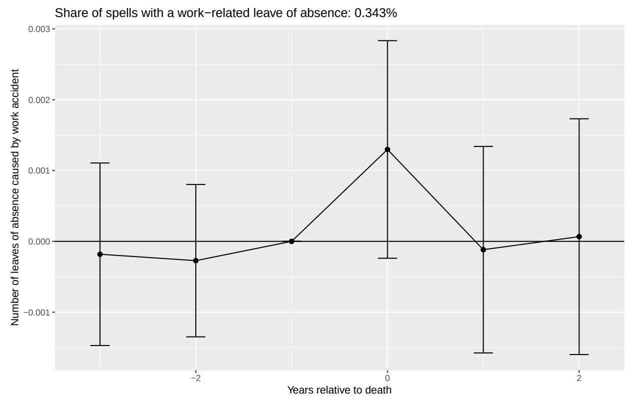

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but using an indicator of whether the leave of absence was caused by a work-related accident as dependent variable, instead of quits. I normalize β_{-1} to zero.

feasible because work-related deaths are rare, so we do not see many repeated instances within an establishment. Instead, I use information on leaves of absence. When a worker takes a leave of absence, the information is recorded in the RAIS, along with the reason why he took the leave. One such reason is that the worker was injured in a work-related accident, which is also indicative of occupational hazards. As discussed in Appendix C, I refrained from using this information in my fatality risk measure because it is hard to compare these nonfatal injuries with deaths, but that is not a problem if I analyze them separately. So I estimate model (2.5), but this time my dependent variable is an indicator for work-related nonfatal injuries that cause a leave of absence instead of quits.

Figure 2.5 presents the results. There is a clear, though not statistically significant at the 5% level, spike of 0.13 p.p. in the probability of a leave of absence caused by a work-accident in the year of the death, which is a big change relative to the baseline of 0.34%. In the other years, there is no difference across establishments in which the death is work-related or not. While it is reassuring to see that there is no difference in the years after the death, the contemporaneous spike is harder to interpret. On the one hand, this spike is worrisome, because it may be indicative that there were more accidents, so the job may be objectively riskier; on the other hand, this is absolutely expected because these nonfatal injuries may be directly related to the fatal accident itself. That is, this spike happens because the work accident may have killed an employee while also injuring some others, though not fatally.

A natural subsequent concern is whether the quits I had found earlier were not driven precisely by those who got a nonfatal injury. While in principle one could interpret such quits as a result of learning, an equally (if not more) plausible interpretation is that these workers may have been left incapacitated to perform the job and decided to quit. If that was the case, quits would not be indicative of uncertainty, but would have been directly caused by the work-related fatality. In order to assuage this concern, I replicated Figure 2.4, but removed from the sample any worker that took a leave of absence caused by a work-related accident. Because there are very few such cases, the results are virtually unchanged, so I have omitted it from the main text and left it in the Appendix Table B.2; just as in Figure 2.4, there is a spike in the probability of quitting only among low tenure workers in the year of the death, but now we know that these were not driven by workers that got non-fatally injured in a work-accident—at least to the extent that it would have warranted a leave of absence.

Another way through which a work-related death could objectively affect the risk of a job is through physical capital. That is, the fatal accident may damage equipment in a way that makes future accidents more likely.[17] In order to address this issue, I start by analyzing the quitting patterns of workers in service occupations. These jobs are relatively less intensive in physical capital, so it is less likely that the work-related fatality caused physical capital damage that could lead to an increased risk of accident afterwards. For example, some common occupations found in my data are cashiers and doormen, whose work-related deaths are likely due to robberies, rather than any physical capital destruction.[18]

[17]On the other hand, work fatalities could lead to a review of production practices and less complacency with safety protocols by the part of employers and employees, diminishing risks. Therefore, I cannot rule out this possibility without data.

[18]In the Appendix Figures B.3 and B.4, I show that the spike in non-fatal injuries found in Figure 2.5 seems to be driven entirely by workers that are not in Service occupations. This is consistent with the idea that a work-related fatality is more likely to impose some collateral damage on other workers if the work

Figure 2.6: Relative quits (service occupations) - Work-related versus Unrelated

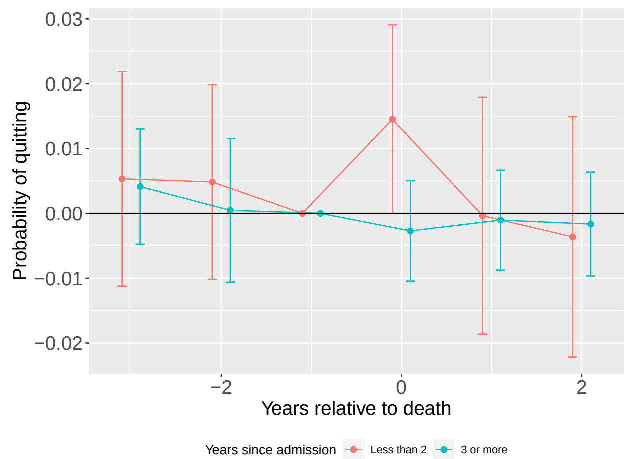

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but focusing on service occupations. I normalize β_{-1} to zero.

In Figure 2.6, I plot the coefficients of model (2.5) in an attempt to replicate Figure 2.4, but now restricted to service occupations. The pattern is familiar: high tenure workers are unaffected, but among workers with less than two years of tenure, there is a spike in the number of quitting workers in the year of the death, although its 95% confidence interval does barely contain zero. The fact that there is a quit reaction even in service occupations is indicative that physical capital damage is unlikely to be the sole explanation for my main results.

A more indirect approach to test whether the job becomes objectively more dangerous after a work-related fatality is to compare the growth rate of the establishments, depending on the type of death they faced. A more dangerous job is also a less attractive one, making it harder for establishments to hire and retain workers, hampering their growth.[19] Therefore, I collapse the individual worker sample at the establishment level to estimate the following

requires large machinery.

[19]I focus on growth here, but the firm could also pay higher wages as a result of the increased risk. As shown in Figure 2.10 that does not seem to be the case as well.

Figure 2.7: Establishment growth - Work-related versus Unrelated

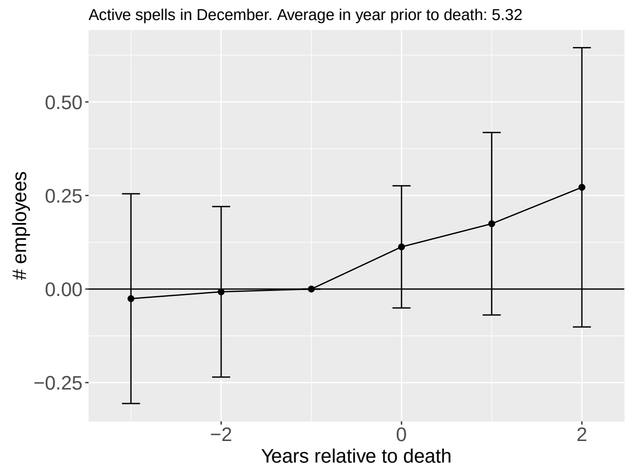

The graph plots the β_k coefficients estimated with the model described in Equation (2.6). I normalize β_{-1} to zero.

model:

$$n_{j,t} = \sum_{k(j,t)=-3}^{2} \gamma_{k(j,t)} D_{k(j,t)} + \sum_{k(j,t)=-3}^{2} \beta_{k(j,t)} D_{k(j,t)} W(j) + \xi_t + \xi_j + \varepsilon_{j,t} \tag{2.6}$$

in which $n_{j,t}$ is the number of active job spells in December of establishment j in the calendar year t, and all the other variables are defined in the same way as I did in (2.5). Just like before, the parameter of interest here are the β_k's: they capture the difference in size between establishments around the death, depending on whether it was work-related or not.

The results are shown in Figure 2.7. Before the death, there is no meaningful difference in the size of establishments depending on the circumstances of the death. Yet, growth seems to be, if anything, larger after a work-related fatality, which is exactly the opposite of what we would expect if they had an objective impact on the job's risk. Granted, this effect is not statistically significant, and even its economic magnitude is not large, with an increase of roughly 0.25 workers at its peak, two years after the death. Moreover, this result is not driven by outliers nor is it concentrated in specific industries, occupations or years.

Overall, I conclude that it is unlikely that an objective change in fatality risk that is concomitant to the work-related death is driving my results. Objective changes in risk failed to manifest themselves clearly in any of my tests. First, there is no effect in the number of nonfatal work-related accidents *after* the death; second, removing people that could claim to have objective reasons to quit due to a work-related nonfatal accident has no impact in the main quitting patterns; third, quits happen even in occupations in which physical capital damage is unlikely; fourth, there is no evidence that firms with work-related fatalities encountered any difficulty in replacing the quitting workers, and if anything they actually grew relatively larger.

2.4.2 Preference heterogeneity

An important assumption I have made is that the preference parameter that determines people's WTP for decreases in fatality risk, γ, is homogeneous. Although it is not uncommon in the literature, especially when using a revealed preference approach (Lavetti & Schmutte, 2018), it is certainly a strong assumption. Preference homogeneity allows me to interpret the absence of a differential quit reaction among high tenure workers as evidence that they do not update their perceptions significantly. Implicitly, the fact that low tenure workers do quit shows that, under preference homogeneity, everyone cares about fatality risks. Therefore, the muted reaction of high tenure workers could only be explained by an unchanged perception of risk.

Yet, if γ is not constant this is not necessarily the case. An alternative rationale for the tenure gradient on quits is simply that high tenure workers do not really care about fatality risks, i.e., their γ is close to zero.[20] In that case, perceptions would not need to be stable in order for the quitting probabilities to be the same among high tenure workers, when comparing work-related versus other deaths; in principle, perceptions could be changing dramatically, but if workers do not care about these risks, such changes would not be reflected in their quit behavior. This cannot rule out learning by itself, because it fails to explain why

[20]Remember that I am assuming that $\gamma \geq 0$, so nobody strictly prefers to endanger their lives.

low tenure workers do quit more in the first place. But it does mean that I cannot time the learning process, which would present a difficulty when trying to value fatality risks, as I do in Section 2.5.

One dimension in which preferences could differ is within individuals, across time. This age heterogeneity could be driving the differential response between high and low tenure workers, because job experience and age are correlated. Yet, previous research has not reached a consensus on the relationship between age and WTP for diminished fatality risks, neither theoretically nor empirically (Evans & Smith, 2006; Aldy & Viscusi, 2008). Moreover, the concern typically found in this literature relates to a potential decline in the WTP late in life, for people past their 60s, but my sample is far younger than that, even conditional on being a high tenure worker (see Appendix Table A.2).

A more consequential concern is that γ may differ across individuals, because then endogenous selection is an irrefutable reason why we should expect high tenure workers to have lower γ's—and potentially close to zero. A worker with smaller γ would need a worse perception of risk in order to quit; therefore, high tenure workers should have lower γ on average. The question is how much lower.

I use a simple hedonic regression in order to investigate this. As proposed by Rosen (1974), the wage premium of a certain job characteristic can be interpreted as the worker's willingness to pay for them under some assumptions. And while this strong relationship has been shown not to be robust to alternative assumptions (Hwang et al., 1992; Hwang et al., 1998), it remains intuitive that the two concepts should at least bear a resemblance. Moreover, in Brazil there are restrictions to equilibrium prices, due to mandated wage premiums related to safety concerns that are similar to minimum wage regulations.[21] Therefore, I cannot interpret the risk wage premium as the preference parameter γ, and the following exercise is merely meant to be suggestive.

[21] See Appendix E for details.

In order to estimate the fatality risk wage premium, I use

$$w_{i,j,t} = \kappa r_{o(i,j)} + \Xi x'_{i,j,t} + \varepsilon_{i,j,t} \tag{2.7}$$

in which $w_{i,j,t}$ is the log hourly wage of person i who works at establishment j in calendar year t. The fatality rate $r_{o(i,j)}$ is calculated at the occupation level, as described in Section 2.2; its coefficient κ is the parameter of interest here, capturing the (log) wage premium associated with a higher fatality risk. Regression controls $x_{i,j,t}$ include a quartic polynomial in establishment size, and fixed effects for age, education, race, gender, industry, calendar year and State.

I estimate model (2.7) separately by tenure and for two different samples. The full sample has all individuals aged eighteen to sixty-five years old in full-time contracts in the private sector; the restricted sample includes only job spells in the full sample that I observe, at some point in my data, to have had exactly five years of tenure. This latter sample is meant to be "selected", in the sense that it only includes worker-establishment combinations that I know are relatively successful, because they last a long time. Therefore, comparing the wage premiums the restricted sample has received across their tenure with the full one may be informative of the potential impact selection on γ has.

The results are shown in Figure 2.8. As expected if there is selection on preferences, the fatality wage premium decreases over tenure. On the other hand, it is reassuring to see that the wage premium is not very different between workers with zero to five years of tenure, although the full sample wage premium for zero years is much higher and the restricted sample premium for five years is much lower. The fact that the premium profile is relatively flat around those years suggests that the underlying γ for these populations is not likely to differ too much as well. The reason why I focus in this particular tenure range is that it comprises most of the data, with roughly 90% of the job spells in the full sample data, as shown in the shaded bars.[22] Furthermore, the wage premium profile for both samples is

[22]In the sample of establishments with unique deaths that is used in Section 2.3, the share of job spells

Figure 2.8: Fatality risk premium, by tenure

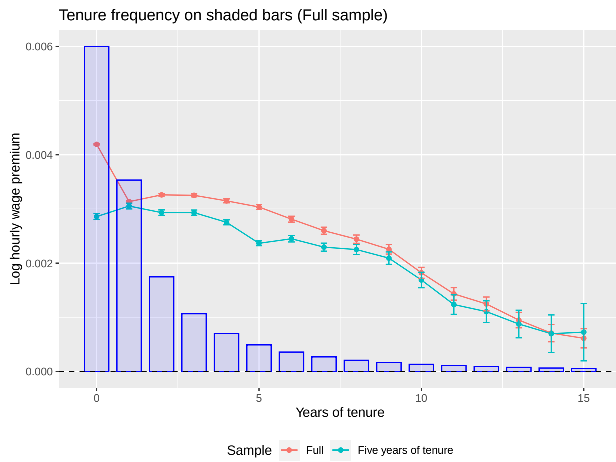

The graph plots the κ coefficient of estimated with model (2.7), for different tenure levels. The full sample consists of individuals aged eighteen to sixty-five in full-time contracts in the private sector; the restricted sample includes only jobs spells in the full sample that I observe, at some point in my data, to have exactly 5 years of tenure. The shaded bars are the relative tenure frequency in the full sample. The y-axis refers to the risk premium, not the frequencies.

fairly similar, even though the restricted sample only includes long-lasting job spells.

The similarity of the wage premium for the tenure levels that define my low and (most of) my high tenure workers suggests it is hard for selection on γ to be driving my main results. Nevertheless, if the wage premium is being determined by the marginal worker, then the infra-marginal distribution of preferences may be changing dramatically even with a constant wage premium, a possibility I cannot really rule out (Rosen, 1986). Therefore, I judge this evidence to be merely suggestive of the limited role heterogeneity in preferences may have. Importantly, preference heterogeneity would not rule out learning, because it cannot explain the spike in quits among low tenure workers, but it would cast doubts in my conclusion that learning about fatality risk is essentially over after three years of job experience.

However, the bias in valuation of fatality risks caused by uncertainty goes in the opposite direction than what selection on γ predicts. Because low tenure workers tend to underestimate the risk of their jobs, uncertainty over fatality risks predicts the WTP of high tenure

with at most five years of tenure is 90.91%.

Figure 2.9: Relative quits (experienced decedents) - Work-related versus Unrelated

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but focusing on deaths of workers who had at least three years of tenure. I normalize β_{-1} to zero.

workers should be higher. But that is the opposite of what we would expect to find under preference heterogeneity: workers with lower γ should stay longer in their jobs. As I show in the next section, uncertainty seems to dominate, so even though I cannot rule out preference heterogeneity here, it does not qualitatively affect my results on valuation bias.

2.4.3 Risk heterogeneity

I assume that fatality risks are homogeneous within a given occupation. This is done out of necessity, as it would be impossible to reliably estimate these risks otherwise. But it is clearly a strong assumption that I should be wary of: for example, a security guard that works night shifts may face higher risks than a security guard that works during the day (Black & Kniesner, 2003).

A more concerning scenario here would be if these fatality risks differ significantly across job tenure. In particular, 80% of the work-related deaths in my sample happen to workers with two or less years of tenure. If an important component of fatality risks is tenure-specific, then the absence of a reaction among high tenure workers may not be evidence that learning is over, but that they face different hazards instead.

In order to address this concern, I re-estimated Equation (2.5) using only deaths of experienced workers, with results shown in Figure 2.9. If fatality risk heterogeneity is driving my results, I should not observe the same patterns I see in Figure 2.4. As mentioned above, work-related deaths of experienced workers are a minority, so I lose a lot of statistical power. Nevertheless, the qualitative pattern observed in Figure 2.9 is familiar: in the year of the death, there is a spike in the quitting probabilities among low tenure colleagues, but nothing of the sort among the experienced ones. Therefore, I conclude that it is unlikely that my main results are driven by risk heterogeneity.

2.4.4 Wages and human capital interactions

I find no evidence that work-related deaths have any discernible impact on wages. This is related—but with a key difference, as I discuss below—to Jäger and Heining (2019), who showed that the death of a worker may raise the firm's demand for the best substitute input available: the labor of his coworkers in the same occupation. In their study, there is substitution between human capital among coworkers, so a colleague's death caused wage increases and, importantly, a higher retention rate—which should imply a lower quitting rate—among survivors. They find this to be important because the surviving coworkers should provide firm-specific human capital that is not easily replaced with outside hires. In order to assess this hypothesis, I estimate model (2.5), but with hourly wages as my dependent variable, instead of the indicator for quits.

By differentiating between high and low tenure workers, I allow for the effect of the work related death to have a differential impact, depending on whether the coworkers have more or less human capital, respectively. The results in Figure 2.10 do not indicate any discernible pattern, across either group, which is not what we would expect if human capital interactions were important. Still, given the turnover consequences I have already documented, these results should compound productivity and composition changes, so this is not conclusive evidence.

Figure 2.10: Hourly wage - Work-related versus Unrelated

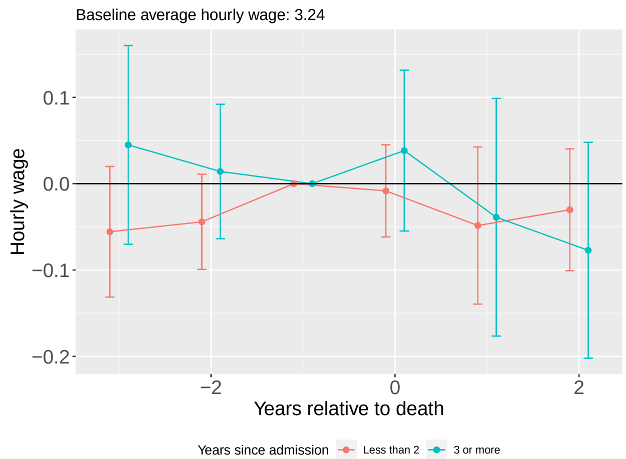

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but using hourly wage as the dependent variable, instead of quits. I normalize β_{-1} to zero.

A key difference between my analysis and Jäger and Heining (2019)'s is that I am comparing two groups that do lose a work colleague, but under different circumstances. In their study, the comparison is between firms in which a worker dies versus observably equivalent firms that did not face this problem. Therefore, if I want to understand the role interactions between human capital could have in my setup, I need to establish whether the comparison between work-related versus unrelated deaths should be interpreted as a relative loss or gain of human capital.[23]

Table 2.2 provides some insight, with the summary statistics of the dead workers. Persons who died in a work-related fatality tend to be younger and less experienced, which suggests they have less human capital. Yet, they also earn more than people whose death was unrelated to work, even after the weights equalize their fatality risks. Therefore, the summary statistics do not paint a clear picture regarding the relative loss in human capital, so I need to consider both possibilities.

[23]While a death is always a loss of human capital, persons that die in a work-related accident may have more or less human capital, on average, than those that died for reasons unrelated to work. Therefore, the relative effect here could in principle be interpreted either way.

Table 2.2: Summary statistics of dead workers

Type of death	Unrelated to work	Work-related accident
Hourly wage (2016 US$)	3.402	3.701
High School	0.435	0.449
College	0.030	0.028
Male	0.914	0.955
Age	42.692	38.076
Fatality rate (per 100,000 FTE)	12.171	12.171
Job experience	3.071	2.127
Observations	144,185	5,970

Sample averages over dead workers, in the year of death. Weights are set so the occupation distribution across establishments in the control group is the same as in the treatment one. Job experience is measured as number of years since spell started, rounded down.

Because I am focusing on coworkers in the same occupation as the deceased one, I assume their human capital are substitutes. Therefore, in order for human capital interactions to explain the spike in quits in establishments with work-related deaths, the relative loss in human capital must be greater when the death is unrelated to work. In that scenario, the incentives for retention would be greater in establishments with deaths unrelated to work, and that could explain why we see a relatively higher quit rate when the death is work-related. The problem with this argument is that it cannot explain what we observe among high tenure coworkers. Given their higher human capital, the retention pull being described here should be even stronger and they would have been less likely to quit in the treated establishments. Yet, the effect among these coworkers is null in the data. Therefore, both wages and quit rates indicate that human capital interactions likely have a small role in my setup.

2.4.5 Following the career of quitters

A person that quits his job because he suddenly perceived it to be too dangerous should have worse subsequent work opportunities, relative to someone that quits under more prosaic circumstances. While the latter is likely being pulled to a new job, the former is essentially being “pushed out”. Therefore, in order to further scrutinize my interpretation, I analyze the career progression of persons that quit these establishments with unique deaths. Presumably, the death unrelated to work does not push away the surviving coworkers, unlike work-related fatalities. Overall, there is evidence that those being scared away struggle to find subsequent employment, but this seems limited to the extensive margin.

I used the regression specification (2.5) again with different outcome variables, but this time I do not break it between sub-samples of workers with up to two years of job tenure and those with three or more years. I do so due to a lack of statistical power. First, only about 8% of the job spells in establishments with unique deaths were terminated with a quit, for a total of 406,484 quitting spells. Second, 21.80% of these people cannot be found in the RAIS data subsequently, so I only know the following employment information for 317,879 job spells. The other ones may have found employment in the informal sector, or left the labor force entirely. Conditional on having quit an establishment in which the death is work-related, I have only 20,186 quitters and am able to find 16,633 of them in their next job.

The first outcome I analyze is an indicator for not being found in the RAIS after quitting. Given that quitters are relatively young—about thirty-one years old, on average—it is reasonable to imagine they are not planning on leaving the labor force yet. Thus, being subsequently found in the RAIS can be meaningful in and of itself: at the very least, these workers were able to find new (formal) jobs. The results are shown in Figure 2.11. While the pre-trend is somewhat worrying, I do observe that workers who quit in the year of the death are 2.19 p.p. more likely to not be found subsequently in the RAIS if the death was work-related, which is roughly a 10% increase relative to the baseline probability. This is consistent with work-related deaths pushing some workers away, so they are less likely to face favorable conditions when seeking their next job. This conclusion must be tempered with the understanding that persons who cannot be found on RAIS may be employed in the informal sector. And even though that is usually associated with lower productivity, lower

Figure 2.11: Probability of not being found in RAIS after quitting

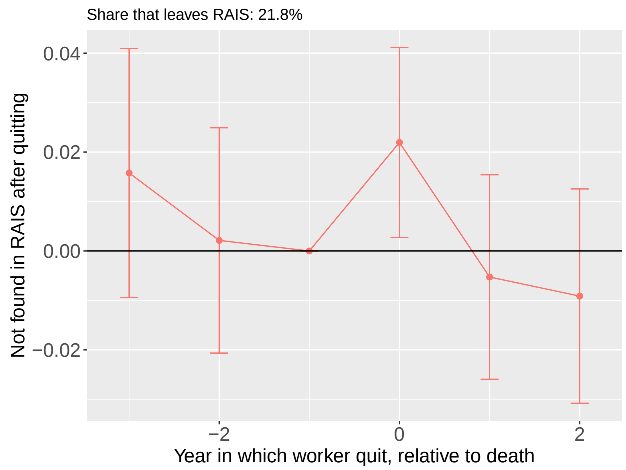

The graph plots the β_k's coefficients estimated with the model described in Equation (2.5), in which the outcome variable is an indicator of whether the worker can be found in the RAIS after quitting one of the establishments with unique deaths in my sample. I normalize β_{-1} to zero.

wages and worse working conditions, some recent research in Brazil has painted a far more nuanced picture (Meghir, Narita, & Robin, 2015; Ulyssea, 2018).

For the workers that are found later in the RAIS, I compare them across three different outcomes. Yet, none show evidence of a systematic difference between people that quit establishments with work-related deaths versus unrelated ones, before or after the event. The outcomes are: an indicator for whether the following job has the same occupation code as the previous one, the number of months between the two jobs and the percentage change in wages.

Figure 2.12 shows the results for the percentage change in wage between jobs. Intuitively, a worker may quit because he has a good outside offer that entails a wage increase. If you quit your job because you deemed it too dangerous, you are less likely to have this better outside offer in hand. Therefore, we should expect wage increases among those that quit the treated establishments to be lower. Yet, there is no discernible difference between people that quit establishments with work-related deaths versus unrelated ones. This null result is also found when I look at both the time it takes to find a new job and occupation change.[24] Therefore,

[24]These results can be found in the Appendix Tables B.5 and B.6, respectively.

Figure 2.12: Percentage change in wage between jobs

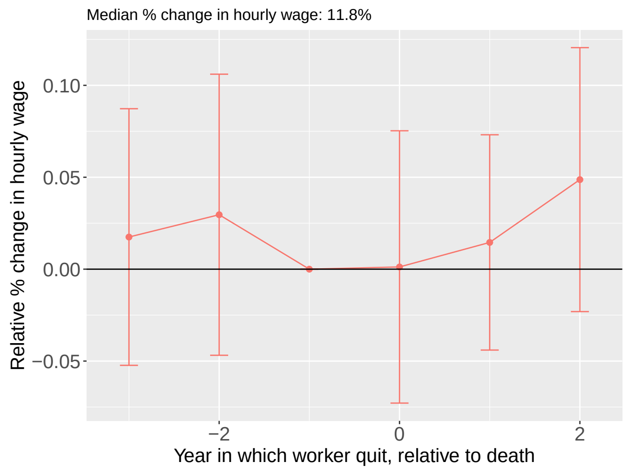

The graph plots the β_k's coefficients estimated with the model described in Equation (2.5), with the percentage change in wage between jobs as the outcome variable. I normalize β_{-1} to zero.

conditional on finding a formal job, it does not seem that people who quit following a work-related death are at a particular disadvantageous position. Nevertheless, they are less likely to find a job to begin with.

2.5 Valuation of fatality risks

My conceptual framework has three main points. First, workers whose perceptions of fatality risk are malleable should be more likely to quit their jobs following a work-related fatality, which is what I found in the data. Second, if workers are learning through on-the-job experience, their beliefs grow stabler over job tenure. According to the data, the quit reaction is exclusive to workers with two or less years of experience, which indicates that their beliefs about fatality risks are fairly stable after three years of tenure. At that point, objective proxies of fatality risks must approximate these workers' perceptions reasonably well.

Now I can address the third point, the one with clear normative implications: what is the extent of the bias in valuation of fatality risks that is caused by uncertainty? My model predicts that workers' perceptions tend to be too optimistic relative to the objective assessment of fatality risk r we can make as analysts, if workers are uncertain about such risks. This systematic difference means that we likely undervalue fatality risks if we do not account for uncertainty.

In order to assess this bias, I revisit my model and estimate γ in the following way:

$$\gamma = -\frac{\partial Q/\partial \mathbb{E}[\tilde{r}|\mathcal{I}, x]}{\partial Q/\partial w} \tag{2.4 revisited}$$

As previously mentioned, this is very similar to the approach in Gronberg and Reed (1994), with the difference being that I explicitly acknowledge that workers' decisions are based on their beliefs—$\mathbb{E}[\tilde{r}|\mathcal{I}, x]$—and not on what we can measure, r. Importantly, $\mathbb{E}[\tilde{r}|\mathcal{I}, x]$ is not observed, but my empirical evidence suggests it can be approximated by r relatively well if the workers have at least three years of tenure. I compare this estimate with a naive approach that disregards the potential discrepancy between $\mathbb{E}[\tilde{r}|\mathcal{I}, x]$ and r early on in people's careers, i.e., the same procedure, but in an unrestricted sample.

Once I have an estimate for γ, I can convert it to a monetary value simply as $\gamma \times \omega$, in which ω is the hourly wage—not its logarithm. Then, I can evaluate it at the average wage, and scale it so that the change in risk implies an expected loss of one life. This calculation yields the Value of Statistical Life (VSL), a key parameter in many policy decisions, so I report it alongside my estimates of γ (U.S. Office of Information and Regulatory Affairs, 2017; Viscusi, 2018).[25]

Regarding my sample, at this point I am just trying to estimate the conditional expectation $Q(w, \mathbb{E}[\tilde{r}|\mathcal{I}, x])$, so there is no need to focus just on persons in establishments with single deaths. Therefore, I use all private sector workers in full-time contracts aged eighteen to sixty-five years old. Table 2.3 shows summary statistics for the unrestricted sample. I also exclude workers in occupations with less than 10,000 FTE years, because their fatality rates are likely too noisy to be relied upon, as discussed in Appendix C.

[25] To be more precise, I report the VSL in millions of US$. In order to make the calculation simpler, I also re-scaled my fatality rate for the regression table: from deaths per 100,000 FTE, to deaths per 1,000 FTE. This means the VSL in millions is simply $2 \times \hat{\gamma} \times \bar{\omega}$.

Table 2.3: Summary statistics

	Mean	SD
Quit	0.057	0.233
Hourly wage	3.938	5.606
Fatality rate	3.132	5.158
College	0.106	0.308
Less than High School	0.363	0.481
White	0.564	0.496
Male	0.625	0.484
Tenure	3.747	4.533
Age	34.872	10.235
Observations	262,911,425	

Summary statistics of full sample. Includes only private sector workers in full-time open-ended contracts aged between eighteen to sixty-five years old. To ease exposition, fatality rates are reported in deaths per 100,000 FTE years here, but I use deaths per 1,000 FTE years in the regression.

My regression model here is the following:[26]

$$q_{ijt} = \alpha_w w_{ijt} + \alpha_r r_{o(i,t)} + \psi' x_{ijt} + \xi_j + \epsilon_{ijt} \tag{2.8}$$

in which q_{ijt} is an indicator of whether the person i, who works at establishment j in calendar year t, quits his job; w_{ijt} is his log hourly wage; and $r_{o(i,t)}$ is the fatality rate of his occupation. I also include fixed effects for age, years of tenure, establishment, education, race, gender and calendar year. Again, controlling for age and tenure is important here because both are expected to be key determinants in one's decision to quit. As before, I also estimate this model in a restricted sample that only includes workers with at least three years of tenure, because I expect them to have accurate beliefs about fatality risks. The comparison between these two samples allows me to gauge the impact uncertainty over

[26]I present the linear probability model here to remain consistent with the approach in Section 2.3, but in Appendix Table A.3, I present the results for a logistic regression model of quits. Although qualitatively similar, the magnitude of the bias is smaller, close to 17%.

Table 2.4: Estimating γ

Quit	Full sample	High tenure
Log wage (w)	-0.02734^{***}	-0.00228^{***}
	(0.00026)	(0.00014)
Fatality rate (r)	0.00254^{***}	0.00032^{***}
	(0.00008)	(0.00005)
Establishment FE	Y	Y
Mean wage	3.938	
$\hat{\gamma}$	0.093	0.140
	(0.00145)	(0.00053)
VSL (million US$)	0.733	1.103
Observations	262,911,425	112,530,665
R^2	0.1073	0.0723

*p<0.1; **p<0.05; ***p<0.01. Linear probability model of quits. I changed the scale of r to make it easier to convert to a VSL: Fatality rate is measured as number of work-related deaths per 1,000 FTE years, at the 4-digit occupation code level. Regressors omitted from the table are: fixed effects for educational attainment, race, gender, age, years of tenure and calendar year. Standard errors, in parentheses, are clustered at the establishment level.

fatality risks has on its valuation.

The results are shown in Table 2.4. As expected, higher wages are associated with a lower probability of quitting: among high tenure workers, a 1% increase in hourly wage is, ceteris paribus, associated with a 0.23 p.p. decrease in the probability of a worker quitting. When I consider the whole sample, this effect is much stronger, as low tenure workers are much more likely to quit: the same increase in hourly wage is associated with a 2.73 p.p. decrease in the propensity to quit. This is also the case when looking at fatality rates: its effect is eight times larger for the whole sample.

The most important result in this table is the estimate $\hat{\gamma} = \frac{-\hat{\alpha}_r}{\hat{\alpha}_w}$. Using the unrestricted sample, I have $\hat{\gamma}^U = 0.093$, which implies a VSL of US$0.73 million. Once I restrict my attention to workers whose perceptions are likely to be accurate—and therefore close to the fatality rate I am using to proxy for them—the estimate increases to $\hat{\gamma}^R = 0.140$, which implies a VSL of US$1.10 million.[27] This increase in the estimate of γ when accounting for uncertainty was expected. As discussed above, uncertainty means there is dispersion in perceptions around the fatality rate, and job choice means those employed tend to have particularly optimistic views. Therefore, by using the fatality rate to proxy for perceptions, we can rationalize the choices we see in the data with a lower γ, when in fact the actual driver of these decisions was a lower perception of risk. Hence, disregarding the role of uncertainty over fatality risks leads to a 33% undervaluation of the VSL.[28]

Note that this conclusion depends on three years of tenure being enough experience for me to reasonably approximate perceptions of risk with its objective rates. Yet, as discussed in Section 2.4.2, this conclusion relies on the assumption that γ is homogeneous across people. This is a strong assumption, and one I cannot definitively back with my data, so it ultimately presents a caveat to my results. Nevertheless, if γ was indeed heterogeneous across persons, we should expect workers with higher tenure to have lower γ's. If that was indeed the case, this 33% undervaluation would be, if anything, a lower bound on how undervalued fatality risks can be if we disregard uncertainty, because my restriction to high tenure workers would also be selecting low γ persons.

2.6 Conclusion

I showed that workers do not fully grasp the fatality risks of their jobs at first, but seem to eventually learn about them through work experience. Using a simple theoretical framework, I showed that this uncertainty leads to an undervaluation of fatality risks, that workers' perceptions grow accurate with experience, and that the quit reaction to a work-related

[27] Lavetti and Schmutte (2018) estimate a VSL of US$1.73 million (in 2010 US$) for Brazil using the hedonic method with RAIS data between 2003 and 2010. As mentioned before, the hedonic method has many known flaws and there are institutional constraints to wage premiums in Brazil that complicate interpretation, but for completeness I have also estimated the VSL this way with my data, yielding US$1.63 million. These results can be found in Appendix Table A.4.

[28] One may be concerned about how to differentiate between tenure and age effects here. As a robustness check, I have run tenure-specific regression in which I allow for wages and risk to interact with age. I have done so for workers with up to five years of tenure, and the qualitative result in which the VSL increases with tenure is the same, as shown in Appendix Table A.5.

fatality could provide empirical backing for the theory.

I tested these predictions on quit rates using Brazilian data. In the year of a work-related fatality, coworkers in the same occupation-establishment as the decedent are 16% more likely to quit their jobs, relative to coworkers whose colleague died due to reasons unrelated to work. Moreover, this effect is entirely concentrated among the coworkers that had at most two years of tenure at the time of death. Workers with at least three years of job experience do not exhibit the same quit reaction, which is consistent with the idea that people learn with job experience, so their perceptions about fatality risks are nearly crystallized after three years. Supplementary analyses on the quitters' subsequent job opportunities, and on a few alternative explanations further lend credence to the conclusion that workers do not fully comprehend the fatality risks of their jobs.

Finally, I showed that this uncertainty leads to a sizable bias in the estimated willingness to pay for reductions in fatality risks in Brazil. If I do not account for the fact that workers' perceptions of fatality risks are not always accurate—and therefore may differ significantly from an objective assessment of such risks—my estimate of the VSL is 33% lower.

A bias of similar magnitude in the U.S. could have important policy implications. Granted, it is not clear that the uncertainty over fatality risks is comparable across Brazil and the U.S., but its impact could still be significant. According to the U.S. Report to Congress on the Benefits and Costs of Federal Regulations and Agency Compliance with the Unfunded Mandates Reform Act (U.S. Office of Information and Regulatory Affairs, 2017),

> "*Across the Federal government, the rules with the highest estimated benefits as well as the highest estimated costs, by far, come from the Environmental Protection Agency and in particular its Office of Air and Radiation. Specifically, EPA rules account for 71 to 80 percent of the monetized benefits and 55 to 64 percent of the monetized costs.* [...] *the largest benefits* [of EPA regulations] *are associated with regulations that reduce risks to life.*"

As an illustration of the importance of these benefits, according to a report by the Environmental Protection Agency (EPA), the Clear Air Act alone may have had an accumulated impact worth nearly US$2 trillion between 1990 and 2020, most of which is attributed to improvements in human health through the VSL (Industrial Economics Inc., 2011). A simple extrapolation of my results would imply that these benefits are undervalued by nearly US$1 trillion.

Moreover, there are other valuation methods that use surveys and rely on stated preferences, rather than revealed ones. Notwithstanding the concerns over survey participants' incentive to truly elicit their preferences, this methodology does have the advantage of clearly stating the trade-off being made, so concerns over the uncertainty I have discussed here are diminished. Therefore, my results indicate that the relative weight given to Stated Preferences studies should perhaps increase in a meta-analysis of the VSL. Contrary to my findings though, this literature seems to reach smaller estimates of the VSL (Robinson & Hammitt, 2016). Yet, these surveys do not face the selection issues that arise in real labor markets, which is the root of the undervaluation of fatality risks I have found.

REFERENCES

Aldy, J. E., & Viscusi, W. K. (2008). Adjusting the value of a statistical life for age and cohort effects. *Review of Economics and Statistics*. doi: 10.3766/jaaa.20.6.1

Alvarez, F., Argente, D., & Lippi, F. (2020). *A Simple Planning Problem for COVID-19 Lockdown.*

Ashenfelter, O., & Greenstone, M. (2004). Using Mandated Speed Limits to Measure the Value of a Statistical Life. *Journal of Political Economy*, *112*(1). doi: 10.2139/ssrn.331463

Azoulay, P., Graff Zivin, J. S., & Wang, J. (2010). Superstar extinction. *Quarterly Journal of Economics*(May). doi: 10.1162/qjec.2010.125.2.549

Biddle, J., & Roberts, K. (2003). Claiming behavior in workers' compensation. *Journal of Risk and Uncertainty*, *70*(4), 759–780. doi: 10.1046/j.0022-4367.2003.00074.x

Black, D. A., & Kniesner, T. J. (2003). On the Measurement of Job Risk in Hedonic Wage Models. *Journal of Risk and Uncertainty*, *27*(3), 205–220. doi: 10.1023/A:1025889125822

Bonhomme, S., & Jolivet, G. (2009). The pervasive absence of compensating differentials. *Journal of Applied Econometrics*, *24*, 763–795. doi: 10.1002/jae

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics*, *36*(9850). doi: 10.3386/w22850

Cropper, M., Hammitt, J. K., & Robinson, L. A. (2011). Valuing Mortality Risk Reductions: Progress and Challenges. *Annual Review of Resource Economics*, *3*(1), 313–336. doi: 10.1146/annurev.resource.012809.103949

Evans, M. F., & Kerry Smith, V. (2008). Complementarity and the measurement of individual risk tradeoffs: Accounting for quantity and quality of life effects. *Environmental and Resource Economics*, *41*(3), 381–400. doi: 10.1007/s10640-008-9197-9

Evans, M. F., & Schaur, G. (2010). A quantile estimation approach to identify income and age variation in the value of a statistical life. *Journal of Environmental Economics and Management*, *59*(3), 260–270. Retrieved from `http://dx.doi.org/10.1016/j.jeem.2009.11.004` doi: 10.1016/j.jeem.2009.11.004

Evans, M. F., & Smith, V. K. (2006). Do we really understand the age-VSL relationship? *Resource and Energy Economics*, *28*(3), 242–261. doi: 10.1016/j.reseneeco.2006.02.004

Greenstone, M., & Nigam, V. (2020). *Does Social Distancing Matter?*

Gronberg, T. J., & Reed, W. R. (1994). Estimating Workers' Marginal Willingness to Pay for Job Attributes Using Duration Data. *Journal of Human Resources*, *29*(3), 911. doi: 10.2307/146258

Hersch, J., & Viscusi, W. K. (2010). Immigrant status and the value of statistical life. *Journal of Human Resources*, *45*(3), 749–771. doi: 10.3368/jhr.45.3.749

Hwang, H.-s., Mortensen, D. T., & Reed, W. R. (1998). Hedonic Wages and Labor Market Search. *Journal of Labor Economics*, *16*(4), 815–847. doi: 10.1086/209907

Hwang, H.-s., Reed, W. R., & Hubbard, C. (1992). Compensating Wage Differentials and Unobserved Productivity. *Journal of Political Economy*, *100*(4), 835–858. doi: 10.1086/261842

Industrial Economics Inc. (2011). *Health and Welfare Benefits Analysis to Support the Second Section 812 Benefit-Cost Analysis of the Clean Air Act.*

Israel, M. (2005). Services as Experience Goods: An Empirical Consumer Learning in Automobile Insurance. *American Economic Review*, *95*(5), 1444–1463. doi: 10.1257/000282805775014335

Jäger, S., & Heining, J. (2019). *How Substitutable Are Workers? Evidence from Worker Deaths.*

Jäger, S., Schoefer, B., & Zweimüller, J. (2019). *Marginal jobs and job surplus: a test of the efficiency of separations.*

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, *3*(3), 430–454. doi: 10.1016/0010-0285(72)90016-3

Kim, Y. (2020). *Customer retention under imperfect information.* doi: 10.2139/ssrn.3709043

Kniesner, T. J., & Viscusi, W. K. (2005). Value of a Statistical Life: relative position vs. relative age. *American Economic Review, Papers and Proceedings*, *95*(2), 141–146.

Kniesner, T. J., Viscusi, W. K., Woock, C., & Ziliak, J. P. (2012). The value of a statistical life: Evidence from panel data. *Review of Economics and Statistics*, *94*(1), 74–87. doi: 10.1162/REST$_{a0}$0229

Kniesner, T. J., Viscusi, W. K., & Ziliak, J. P. (2010). Policy relevant heterogeneity in the value of statistical life: New evidence from panel data quantile regressions. *Journal of Risk and Uncertainty*, *40*(1), 15–31. doi: 10.1007/s11166-009-9084-y

Lamadon, T., Mogstad, M., & Setzler, B. (2019). Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market. *SSRN Electronic Journal*. doi: 10.2139/ssrn.3400197

Lavetti, K. (2018). The Estimation of Compensating Wage Differentials: Lessons From the Deadliest Catch. *Journal of Business and Economic Statistics*, 1–18. doi: 10.1080/07350015.2018.1470000

Lavetti, K., & Schmutte, I. M. (2018). *Estimating Compensating Wage Differentials with Endogenous Job Mobility.*

Lindhjem, H., Navrud, S., Braathen, N. A., & Biausque, V. (2011). Valuing Mortality Risk Reductions from Environmental, Transport, and Health Policies: A Global

Meta-Analysis of Stated Preference Studies. *Risk Analysis*, *31*(9), 1381–1407. doi: 10.1111/j.1539-6924.2011.01694.x

Manski, C. F. (2004). Measuring expectations. *Econometrica*, *72*(5), 1329–1376. doi: 10.1111/j.1468-0262.2004.00537.x

Meghir, C., Narita, R., & Robin, J.-M. (2015). Wages and Informality in Developing Countries. *American Economic Review*, *105*(4), 1509–1546. doi: 10.1257/aer.20121110

Mishan, E. J. (1971). Evaluation of Life and Limb: A Theoretical Approach. *Journal of Political Economy*, *79*(4), 687–705. Retrieved from `https://www.jstor.org/stable/1830825`

Pope, J. C. (2008). Buyer information and the hedonic: The impact of a seller disclosure on the implicit price for airport noise. *Journal of Urban Economics*, *63*(2), 498–516. doi: 10.1016/j.jue.2007.03.003

Robinson, L. A., & Hammitt, J. K. (2016). Valuing reductions in fatal illness risks: implications of recent research. *Health Economics*. doi: 10.1002/hec

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, *82*(1), 34–55. Retrieved from `https://www.jstor.org/stable/1830899`

Rosen, S. (1986). The theory of equalizing differences. In O. Ashenfelter & R. Layard (Eds.), *Handbook of labor economics* (Vol. I, pp. 641–692). Elsevier Science. doi: 10.1016/s1573-4463(86)01015-5

Rundmo, T. (1995). Perceived risk, safety status, and job stress among injured and noninjured employees on offshore petroleum installations. *Journal of Safety Research*, *26*(2), 87–97. doi: 10.1016/0022-4375(95)00008-E

Sarsons, H. (2019). *Interpreting Signals in the Labor Market: Evidence from Medical Referrals.*

Schelling, T. C. (1968). The life you save may be your own. In S. B. Chase (Ed.), *Problems in public expenditure analysis* (pp. 127–162).

Scotton, C. R. (2013). New risk rates, inter-industry differentials and the magnitude of VSL estimates. *Journal of Benefit-Cost Analysis*, *4*(1), 39–80. doi: 10.1515/jbca-2012-0015

Scotton, C. R., & Taylor, L. O. (2011). Valuing risk reductions: Incorporating risk heterogeneity into a revealed preference framework. *Resource and Energy Economics*, *33*(2), 381–397. Retrieved from `http://dx.doi.org/10.1016/j.reseneeco.2010.06.001` doi: 10.1016/j.reseneeco.2010.06.001

Sunstein, C. R. (2002). Probability Neglect: Emotions, Worst Cases, and Law. *The Yale Law Journal*, *112*(1), 61–107.

Thaler, R., & Rosen, S. (1976). The Value of Saving a Life: Evidence from the Labor Market. In *Household production and consumption* (pp. 265–302). NBER.

Ulyssea, G. (2018). Firms, Informality and Development: Theory and evidence from Brazil.

American Economic Review, *108*(632), 2015–2047. doi: 10.1257/aer.20141745

U.S. DOT. (2016). *Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses* (Tech. Rep.). U.S. Department of Transportation.

U.S. EPA. (2016). *Valuing mortality risk reductions for policy: a meta-analytic approach* (Tech. Rep.). U.S. Environmental Protection Agency.

U.S. Office of Information and Regulatory Affairs. (2017). *2017 Report to Congress on the Benefits and Costs of Federal Regulations and Agency Compliance with the Unfunded Mandates Reform Act* (Tech. Rep.). U.S. Office of Management and Budget.

Viscusi, W. K. (1993). The value of risks to life and health. *Journal of Economic Literature*, *1*(4), 385–452. doi: 10.1016/B978-0-444-53685-3.00007-6

Viscusi, W. K. (2003). Racial differences in labor market values of a statistical life. *Journal of Risk and Uncertainty*, *27*(3), 239–256.

Viscusi, W. K. (2004). The value of life: Estimates with risks by occupation and industry. *Economic Inquiry*, *42*(1), 29–48. doi: 10.1093/ei/cbh042

Viscusi, W. K. (2018). *Pricing lives: Guideposts for a safer society.* Princeton University Press. doi: 10.2307/j.ctvc772d8

Viscusi, W. K., & Aldy, J. E. (2007). Labor market estimates of the senior discount for the value of statistical life. *Journal of Environmental Economics and Management*, *53*(3), 377–392. doi: 10.1016/j.jeem.2006.10.005

Viscusi, W. K., & Hersch, J. (2008). The mortality cost to smokers. *Journal of Health Economics*, *27*(4), 943–958. doi: 10.1016/j.jhealeco.2008.01.001

Viscusi, W. K., & O'Connor, C. J. (1984). Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers? *American Economic Review*, *74*(5), 942–956. Retrieved from `https://www.jstor.org/stable/554`

APPENDICES

A Additional tables

Table A.1: Transitions within calendar year by cause for separation

Cause	End of contract	Fired	Laid off	Quit
No job within the year	0.45	0.57	0.54	0.44
Months between jobs	3.19	3.63	4.42	2.43
Wage change	-0.15	0.17	-0.31	0.64
Share of contracts	5.71%	0.58%	19.52%	7.18%

Sample averages for transitions within a calendar year. I restrict to separations that occur in the first 3 months of the year. The first row is the share of separations in which the person finds a new job in the RAIS within the calendar year. For persons that do find another job, I report the average number of months it took and hourly wage change. Share of contracts is not conditional on occurring until March, and is relative to the overall number of links, of which roughly two thirds are not terminated.

Table A.2: Sample average in year immediately before death, by tenure group

	Low tenure ≤ 2		High tenure ≥ 3	
Work-related death?	No	Yes	No	Yes
Hourly wage (2016 US$)	2.908	3.010	4.139	4.170
High School	0.489	0.472	0.378	0.353
College	0.024	0.023	0.044	0.045
Male	0.895	0.925	0.878	0.915
Age	34.826	34.214	42.432	41.241
Fatality rate (per 100,000 FTE)	12.267	12.480	11.345	11.314
Job experience	0.589	0.594	6.925	6.591
Quit	0.097	0.101	0.018	0.022
Layoff	0.269	0.279	0.145	0.155
# of establishments	144,185	5,970	144,185	5,970
Observations across all years	3,553,838	181,308	1,287,001	55,608

Sample averages over workers by establishment treatment status and job experience, in the year immediately before the death. Weights are set so the occupation distribution across establishments in the control group is the same as in the treatment group. Job experience is measured as number of years since spell started, rounded down.

Table A.3: VSL estimation through logit model

Quit	Full sample	High tenure
Log wage (w)	-0.61095^{***}	-0.09804^{***}
	(0.00084)	(0.00178)
Fatality rate (r)	0.07553^{***}	0.01446^{**}
	(0.00105)	(0.00732)
Establishment FE	Y	Y
Mean wage	3.938	
$\hat{\gamma}$	0.124	0.148
	(0.00001)	(0.00557)
VSL (million US$)	0.974	1.162
Observations	262,911,425	112,530,665

*p<0.1; **p<0.05; ***p<0.01. Logistic regression model of quits. I changed the scale of r to make it easier to convert to a VSL: Fatality rate is measured as number of work-related deaths per 1,000 FTE years, at the 4-digit occupation code level. Regressors omitted from the table are: fixed effects for educational attainment, race, gender, age, years of tenure and calendar year. Standard errors, in parentheses, are clustered at the establishment level.

Table A.4: VSL estimation through the hedonic method

	(1) Log hourly wage	(2) Log hourly wage
Fatality rate (r)	0.317***	0.207***
	(0.00482)	(0.00722)
Fatality rate (r) X Less than 3 years		0.133***
		(0.00711)
Zero fatality	0.0205***	0.0227***
	(0.00181)	(0.00219)
Zero fatality X Less than 3 years		-0.00393
		(0.00234)
Less than 3 years		-0.0789***
		(0.00119)
Observations	256,073,557	256,073,557
Mean hourly wage	3.946	
VSL (million US$)	2.499	1.634
Worker and Establishment FE	X	X
R-squared	0.888	0.887

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors in parentheses.
Workers with less than 3 years are unlikely to be aware of the fatality risks of their occupation. I have omitted quadratics of age and establishment size from the table. All models include 1-digit occupation codes, worker, establishment, calendar year and State fixed effects. Fatality rates are calculated using the whole sample period at the 4-digit occupation level, which are only included for occupations with at least 10,000 FTE workers. Only males between 23 and 65 years old working at least 30 hours per week in the private sector are included here. Standard errors are clustered at the firm level. All monetary values are in 2016 US$.

Table A.5: Tenure specific quit regressions

	Dependent variable:					
	Quit					
Years of tenure	(0)	(1)	(2)	(3)	(4)	(5)
Log wage (w)	-0.11781^{***}	-0.03267^{***}	-0.01551^{***}	-0.00824^{***}	-0.00884^{***}	-0.00619^{***}
	(0.00091)	(0.00073)	(0.00067)	(0.00065)	(0.00066)	(0.00050)
Log wage (w) $\times$ Age	0.00159^{***}	0.00037^{***}	0.00021^{***}	0.00015^{***}	0.00020^{***}	0.00015^{***}
	(0.00002)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Fatality rate (r)	0.01305^{***}	0.01128^{***}	0.00898^{***}	0.00746^{***}	0.00609^{***}	0.00311^{***}
	(0.00037)	(0.00032)	(0.00033)	(0.00035)	(0.00041)	(0.00037)
Fatality rate (r) $\times$ Age	-0.00027^{***}	-0.00030^{***}	-0.00024^{***}	-0.00020^{***}	-0.00016^{***}	-0.00008^{***}
	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)
γ	0.066	0.063	0.080	0.098	0.120	0.156
Mean age	31.92	33.30	34.67	35.89	36.96	37.94
Mean wage (US$ per hour)				3.491		
VSL (million US$)	0.461	0.440	0.559	0.684	0.838	1.089
Observations	78,197,995	45,114,245	27,068,520	22,573,943	18,896,879	13,307,210

*p<0.1; **p<0.05; ***p<0.01. Linear probability model of quits. Each column limits the sample to a given level of years of tenure. I changed the scale of r to make it easier to convert to a VSL: Fatality rate is measured as number of work-related deaths per 1,000 FTE years, at the 4-digit occupation code level. Regressors omitted from the table are: fixed effects for educational attainment, race, gender, age, years of tenure and calendar year. Standard errors, in parentheses, are clustered at the establishment level.

B Additional figures

Figure B.1: Quit reaction by establishment size

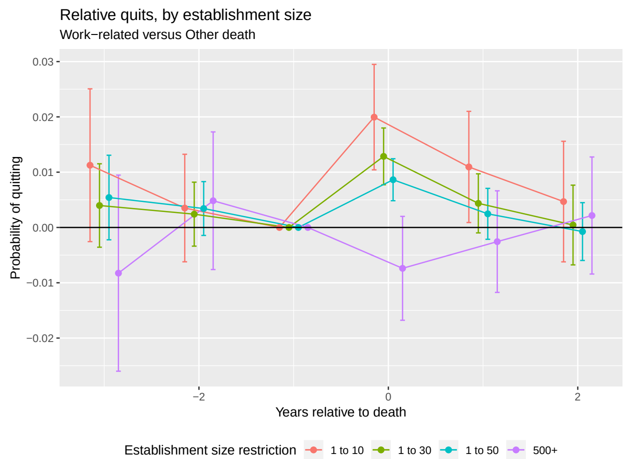

Figure B.2: Relative quits - Work-related versus Unrelated

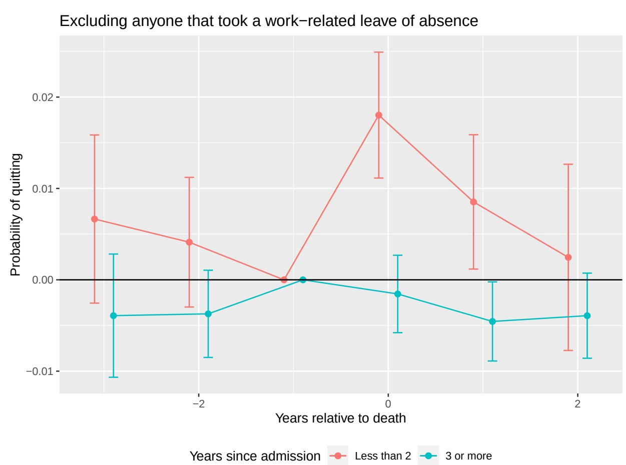

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but removing from the sample any worker that took a leave of absence caused by a work-related accident. I normalize β_{-1} to zero.

Figure B.3: Non-fatal injuries in non-service occupations

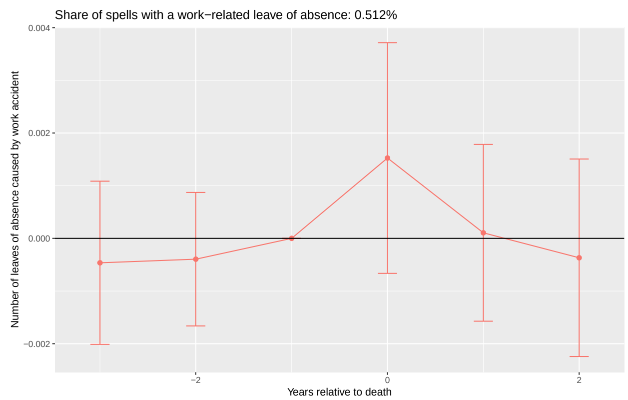

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but using an indicator of whether the leave of absence was caused by a work-related accident as dependent variable, instead of quits and a sample of occupations that are not in the Service industry. I normalize β_{-1} to zero.

Figure B.4: Non-fatal injuries in service occupations

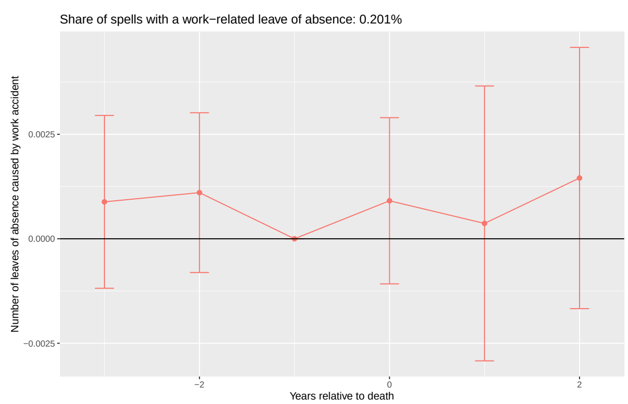

The graph plots the β_k coefficients estimated with the model described in Equation (2.5), but using an indicator of whether the leave of absence was caused by a work-related accident as dependent variable, instead of quits and a sample of occupations in the Service industry. I normalize β_{-1} to zero.

Figure B.5: Length of time until new job

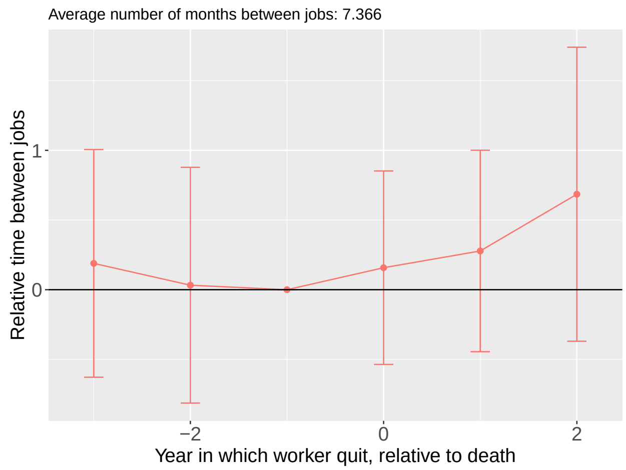

The graph plots the β_k's coefficients estimated with the model described in Equation (2.5), with the time gap between jobs, measured in months, as the outcome variable. I normalize β_{-1} to zero.

Figure B.6: Probability of moving to the same occupation

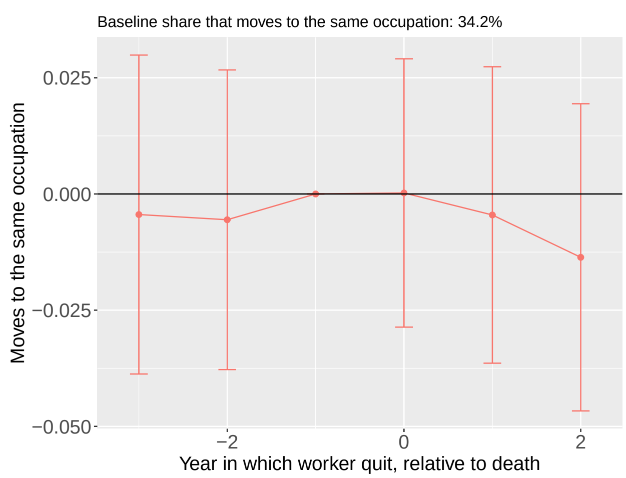

The graph plots the β_k's coefficients estimated with the model described in Equation (2.5), in which the outcome variable is an indicator of whether the quitter has moved to a job with the same occupation. I normalize β_{-1} to zero.

C Calculating fatality rates

In order to assuage concerns about the quality of the reports of work-related fatalities in the RAIS, I compare its fatality rates with the ones provided by the BLS through the CFOI, which is regarded as the official count of work-related fatalities in the U.S., and whose methodology I follow. Moreover, I also use fatality rates as the objective measure of fatality risk for a job. This latter usage drives my choice of group at which the fatality rates are calculated, as I describe below.

The (scaled) fatality rate of a group is simply the number of relevant occurrences per hour worked. This raises two main questions: which fatalities are relevant and how should we define groups. The CFOI methodology uses three criteria to determine whether a workplace fatality is deemed relevant and enters their count: (1) the incident that led to the death must have occurred in the United States, its territories, or its territorial waters or airspace; (2) it must be related to work; and (3) it must have resulted from traumatic injury.

The geographical criterion rules out fatalities that occur outside of Brazilian territory, but the RAIS does not provide information on the whereabouts of its work-related accidents. As these are very specific circumstances that are unlikely to make up a significant share of the reports, I chose to ignore this criterion. Hence, if a worker registered in a Brazilian establishment dies during a business trip outside of Brazilian territory, his death does enter my calculation, in disagreement with the CFOI guidelines.[1]

The CFOI uses data from a broad set of sources, including reports from news media, police, and death certificates, so the second criterion requires the fatal injury to be related to work. Even though the RAIS comes exclusively from reports by the work establishment itself, this criterion still matters. As mentioned earlier, one of the death categories in the RAIS is simply described as “death”, a residual case that presumably cannot be connected to the job like the other three. Even among these three, cases in which the fatal injury

[1] I know for a fact that at least some of the work fatalities in the data did occur outside of Brazilian territory, as the event itself was widely reported in Brazilian media. Yet, the laborious but feasible verification of every work-related death would raise grave privacy concerns, so I never really considered it.

happens during the normal commute to work are also not related to the nature of the job, so the CFOI and I do not count them.

The final criterion is for the death to result from a traumatic injury, defined by the BLS as "*any wound or damage to the body resulting from acute exposure to energy, such as heat or electricity; impact from a crash or fall; or from the absence of such essentials as heat or oxygen, caused by a specific event or incident within a single workday or shift.*" This definition excludes illnesses, because cases that are actually related to work may be hard to pinpoint: "*The long latency period for many fatal occupational illnesses makes it very difficult to compile a complete roster of these cases. It is also difficult to definitively link some cases to a workplace exposure — for example, a coal miner who worked 30 years in a coal mine and smoked two packs of cigarettes a day who dies of lung cancer*", again according to the BLS. Applying the same principle to the RAIS, I exclude deaths due to work-related diseases.

Therefore, I define as relevant occurrences all the deaths that are due to work-related accidents, excluding deaths in the commute to work, due to work-related illnesses or deaths that are not explicitly linked to the job.[2] The RAIS also has information on non-fatal work accidents, because leaves of absence are also reported. Like spell terminations, the reason for the leave of absence must be stated, and work-related injuries is one of the possibilities. I do not incorporate this information in my fatality risk measure for two main reasons. First, injuries may vary widely in their severity, even conditional on a leave of absence being necessary. While this may partially be addressed with the usage of leave duration as a proxy for severity, sequelae may differ dramatically between occupations. Therefore it is hard to compare such injuries across different jobs or with a fatality. Second, under-reporting of work-related injuries is more likely for less severe cases (Biddle & Roberts, 2003).

Admittedly, a drawback of the RAIS relative to the CFOI is the accuracy of the death reports. The CFOI cross-checks every fatality using several different sources of information,

[2]That is, I only include the cases in which the cause for termination is code 62. The excluded categories are codes 63, 64 and 60, respectively.

as its main purpose is precisely to provide an accurate census of these occurrences, which is not the case for the RAIS. And even though work-safety auditors in Brazil do undertake a more detailed analysis of each work-related fatality, there is no systematic compilation of these data I can leverage, to the best of my knowledge.

Using the panel structure of the RAIS, I can at least screen some suspicious cases: a dead person cannot work. If a worker dies, for any reason, he should not be found working subsequently. I have found 514 cases of a worker "dying" in a work-related accident that were either already supposed to be dead (for any possible reason) or are eventually found working later, which are distributed across 212 unique establishments. I excluded these entire establishments from my analysis and not just the suspiciously dead workers, because the establishments are the ones responsible for providing all the information, which becomes hardly credible at this point. Ultimately, I am left with 12,853 work-related fatalities.

In addition to the number of fatalities that occur due to work-related accidents, I need a measure of exposure to calculate the risk. Following the CFOI, I use hours worked as the denominator of the fatality rates. This is an improvement over the employment-based measure often used in the early literature, as some forms of employment may be seasonal or have a large fraction of part-time workers (Viscusi, 2018). Specifically, I define objective risk by the fatality rate of work-related deaths per 200,000,000 hours worked, i.e., the expected number of fatalities in a group of 100,000 FTE individuals for a year. An advantage of the RAIS is that it also has information on hours worked, so I can calculate them directly, while the CFOI relies on the Current Population Survey (CPS) to provide the denominator of their fatality rate.

Using the whole period between 2007 and 2017, Brazil's overall fatality rate is 2.71 per 100,000 FTE years, while the U.S. had a rate of 3.50 in 2018. This may seem surprising and perhaps indicative of under-reporting, because the U.S. is about four times richer than Brazil on a per capita basis and safety is generally regarded as a normal good. There are two main reasons why I believe this is reasonable though. First, unlike the CFOI, a work-related fatality count is a by-product of the RAIS, not its main goal. The sources used in the CFOI are far broader in scope, encompassing medical and police reports, for example, which probably contributes to a higher count of deaths. Second, the RAIS only covers the Brazilian formal employment sector. If the informal one is more dangerous—and the very apparatus of labor safety regulation suggests that is the case—then my fatality rate is likely a lower bound on the one encountered in the informal sector. And while this may be true for the U.S. as well, Brazil's informality rate is higher. More importantly, this is a reminder that my analysis may not be valid outside of the Brazilian formal labor market.

Finally, I can calculate group-specific fatality rates. My goal is to use these group-specific fatality rates as estimates of the true fatality risk a job poses. Hence, it is natural to use industry and occupation codes, as they are designed precisely to group workers based on the nature of their jobs, although my assumption that risks are homogeneous within these groups may not be warranted (Black & Kniesner, 2003). I chose to ignore industry codes and define a group by its 4-digit occupation code, aggregating all years, and discarding occupations with less than 10,000 FTE years in my data. The literature has often combined industry and occupation codes, although the latter are often at a more aggregated level, and this is widely recognized as an important choice when calculating fatality rates (Scotton, 2013). Therefore, I now discuss my choices in detail and why I believe them to be better than their alternatives.

The level of aggregation at which I calculate fatality rates poses the trade-off between a finer definition that may reflect the relevant job risks more accurately, but lack the data to be precisely estimated. The smaller the group, the more likely it is that no deaths occur at all, implying a fatality rate of zero. Yet, if the number of hours worked is small, even a single death may yield a very large fatality rate. Hence, I set a lower bound of 10,000 FTE years on group size, following Lavetti and Schmutte (2018).

I aggregate the data across all the eleven years, ignoring time variation. The job risks I want to capture are an undesirable by-product of a production function that is unlikely to be fluctuating significantly between years; and while technological change or capital accumulation may imply a time trend in how dangerous a job is, yearly fluctuations in fatality rates are more likely due to sampling variation. Furthermore, doing so allows me to be more stringent in a dimension I believe to be more important when defining groups.

Namely, I use the 4-digit level of the occupation codes to define my fatality rates. While early work often used groups as aggregated as industries due to data limitation (Viscusi, 1993), more recent work typically combines both occupation and industry codes. For example, Lavetti and Schmutte (2018), working with the RAIS between 2003 and 2010, use a combination of 2-digit industry and 3-digit occupation codes to define their fatality rates. But according to the methodology of the Brazilian Occupation Classification system used in the RAIS, while the 3-digit level "*gathers occupations that present close familiarity both in terms of the nature of the work and required qualifications*", the 4-digit level code is "*the unit of the classification system. For all practical purposes, it is defined as the set of jobs substantially similar in terms of the nature of the work and required qualifications*". This is precisely what I want: whatever hazard may have killed someone while they were working is therefore relevant for all the workers in that same 4-digit occupation code.

Because ignoring industry codes is a deviation from the literature, I now provide a comparison to the 2-digit industry and 3-digit occupation codes combination, aggregating all years, to show why I believe my choice is better.[3] While there are 619 different 4-digit occupation codes in my data, there are 15,581 combinations of 2-digit industry and 3-digit occupation codes. This is reflected in the size of each group: sixty-four, or 10.34% of the 4-digit occupation cells have less than 10,000 FTE years. When combining 2-digit industry and 3-digit occupation, 12,857 groups fail this size requirement, or 82.52% of the total. And these groups that are too small to use do have important information: 564 work-related deaths and about 16 million FTE years are discarded using the industry and occupation

[3] In their study, Lavetti and Schmutte (2018) calculate year-specific fatality rates and smooth sampling errors by using 3-years moving averages, with data between 2003 and 2010. I aggregate 11 years of data here, so the size criterion should be even more stringent for them.

combination. If I use the 4-digit occupation, though, these numbers fall to two work-related deaths and about 275,000 FTE years discarded due to group size. In fact, even if I use 1,000 FTE years as the lower bound for the industry and occupation combination, 8,988 of the groups would still fail this requirement.[4] Therefore, combining occupation—even at a more aggregated level—with industry makes the task of accurately estimating fatality rates much harder due to sample size limitations.

Still, the most important aspect is which group better reflects the nature of the job, and I believe strict occupation codes are better than laxer ones combined with industry codes. This is because industry codes are set at the establishment level and refer to their main economic activity, while the occupation code is set at the individual worker level. So grouping at the 4-digit occupation level not only loses less information, which is more accurately estimated, but also emphasizes a differentiation more likely to reflect the working conditions for individual employees. Therefore, I do not use industry codes to define my fatality rate groups. From now on, I refer to the 4-digit code occupation simply as occupation, unless stated otherwise.

Thus, my measure of fatality risk is fatality rate at the occupation level during the 2007 to 2017 period, excluding the sixty-four occupations with less than 10,000 FTE years. Precisely because there are so few people in these occupations, the overall fatality rate remains virtually unchanged at 2.71 after excluding them. This average hides substantial heterogeneity across occupations. The first, second and third quartiles of the fatality rate distribution are 0, 1.25 and 3.55, respectively. The same statistics for the FTE years distribution are 49,849.87, 173,576.79 and 542,869.31. There were no work-related fatalities for 141 of the 555 occupations I keep, while 133 of them had at least one work-related fatality per year, on average. In Table C.1, I present the highest fatality rates. The jobs listed here are not surprising, but an important point to emphasize is that even the highest fatality rates are close to 100 per 100,000 FTE years. This indicates just how rare these work-related fatalities

[4]Note that in a group as small as 1,000 FTE years, a single work accident implies a fatality rate of 100 per 100,000 FTE years, which is roughly the same as the highest fatality rates I calculate using the 4-digit occupation code.

Table C.1: Highest fatality rates - 2007 to 2017

Description	Code	Fatalities	FTE years	Fatality Rate
Pilots (small planes) and flight mechanics	3,411	23	23,277.18	98.81
Manufacturing of ammunition and chemical explosives	8,121	26	28,771.91	90.37
Roofers	7,162	4	10,412.45	38.42
Pilots (large planes, prototypes and special flights)	2,153	29	88,848.46	32.64
Truck driver	7,825	2,803	9,688,040.00	28.93
Delivery biker or motor-biker	5,191	303	1,389,984.00	21.80
Forestial extraction of medicinal, aromatic and toxic species	6,325	3	15,269.58	19.65
High sea fisherman	6,312	4	20,411.10	19.60
Ship technician and fisherman	3,412	8	43,759.73	18.28
Ship machinery technician	3,413	4	26,749.73	14.95

Fatalities are work-related deaths that occur during normal work activities. A full-time equivalent (FTE) year is defined as 2,000 hours. The Fatality Rate is defined as number of fatalities per 100,000 FTE years. I have excluded from this table occupations with less than 10,000 FTE years.

are, which is why I approximate them to zero in my model.

In order to validate my fatality rate estimates, I have compared them with the CFOI ones. I did not use a complete list of fatality rates by occupation, because the CFOI only publishes the rates for occupations that meet a lower bound on hours worked or number of fatalities. In any case, occupations with a large number of hours worked are also the ones that are more precisely estimated. I have aggregated the CFOI data for the years between 2011 and 2017, and manually matched them based on their description, as shown in Table C.2.

In Figure C.1, I plot both fatality rates and an identity line to help visualize their discrepancies, with circle size based on the number of hours worked in the RAIS. In order to aid visualization, this is a subset of occupations for which the fatality rates are not above 50 per 100,000 FTE years in neither source. The fatality rates are remarkably similar, especially considering how rare these events are and the difficulties for accurate estimation that presents. Unsurprisingly, the larger differences tend to occur in occupations with fewer hours worked. The comparison of fatality rates that also includes occupations with a value larger than 50 per 100,000 FTE years is presented in the Figure C.2. Among the more extreme fatality rates, the differences are larger, but some of them were expected. For example, fishing in the Bering Sea is known to be an incredibly dangerous job (Lavetti, 2018), but such extreme circumstances are not encountered in Brazil.

Table C.2: List of occupations matched between CFOI and RAIS

English description	Portuguese description
Aircraft pilots and flight engineers	Pilotos de aviação comercial, mecânicos de vôo e afins
Automotive service technicians and mechanics	Mecânicos de manutenção de veículos automotores
Carpenters	Supervisores em indústria de madeira, mobiliário e da carpintaria veicular
Cashiers	Caixas e bilheteiros (exceto caixa de banco)
Construction laborers	Ajudantes de obras civis
Construction managers	Supervisores da construção civil
Driver/sales workers and truck drivers	Motoristas de veículos de cargas em geral
Electrical power or telecommunications line installers and repairers	Instaladores e reparadores de linhas e cabos elétricos, telefônicos e de comunicação de dados
Electricians	Eletricistas de manutenção eletroeletrônica
Firefighters	Bombeiros e salva-vidas
Fishers and related fishing workers	Pescadores diversos (codes 6310, 6311, 6312)
Heating, air conditioning, and refrigeration mechanics and installers	Mecânicos de manutenção e instalação de aparelhos de climatização e refrigeração
Industrial machinery installation, repair, and maintenance workers	Mecânicos de manutenção de máquinas industriais
Janitors and building cleaners	Trabalhadores nos serviços de manutenção de edificações
Laborers and freight, stock, and material movers, hand	Trabalhadores de cargas e descargas de mercadorias
Logging workers	Extrativistas e reflorestadores de espécies produtoras de madeira
Mining machine operators	Trabalhadores de extração de minerais sólidos (operadores de máquinas)
Painters, construction and maintenance	Pintores de obras e revestidores de interiores (revestimentos flexíveis)
Pipelayers, plumbers, pipefitters, and steamfitters	Encanadores e instaladores de tubulações
Refuse and recyclable material collectors	Trabalhadores da coleta e seleção de material reciclável
Registered nurses	Enfermeiros e afins
Retail salespersons	Operadores do comércio em lojas e mercados
Roofers	Telhadores (revestimentos rígidos)
Security guards and gaming surveillance officers	Vigilantes e guardas de segurança
Structural iron and steel workers	Trabalhadores de traçagem e montagem de estruturas metálicas e de compósitos
Supervisors of food preparation and serving workers	Trabalhadores no atendimento em estabelecimentos de serviços de alimentação, bebidas e hotelaria
Taxi drivers and chauffeurs	Motoristas de veículos de pequeno e médio porte
Welding, soldering, and brazing workers	Trabalhadores de soldagem e corte de ligas metálicas

Fatality rates are aggregated between 2011 and 2017 for the CFOI and 2007 and 2017 for the RAIS. Selected occupations based on CFOI availability and a matching description on the RAIS. I have combined two distinct CFOI occupations into one: "Electrical power-line installers and repairers" and "Telecommunications line installers and repairers" were combined because they are reported as one in the RAIS.

Figure C.1: Fatality rate comparison

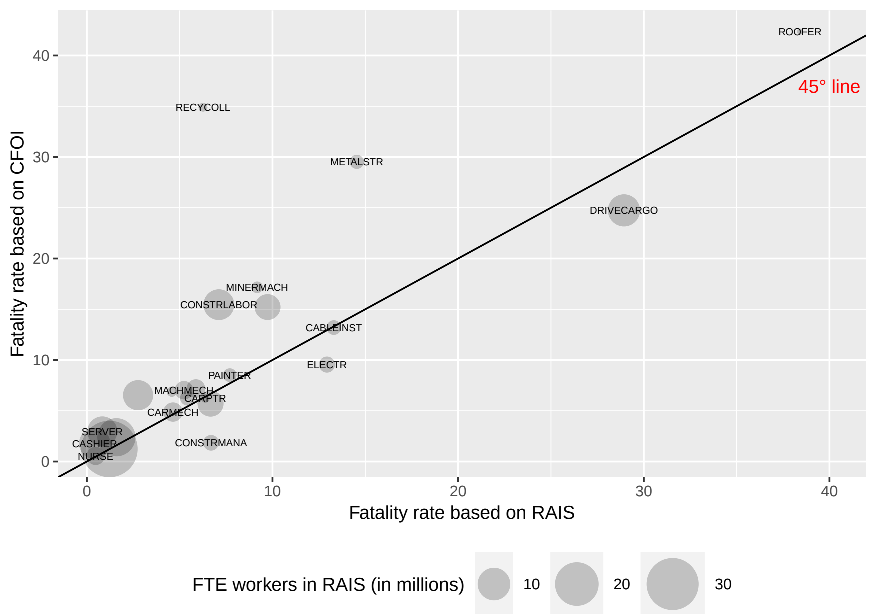

Selected occupations. CFOI between 2011 and 2017, RAIS between 2007 and 2017.

Figure C.2: Fatality rate comparison - expanded

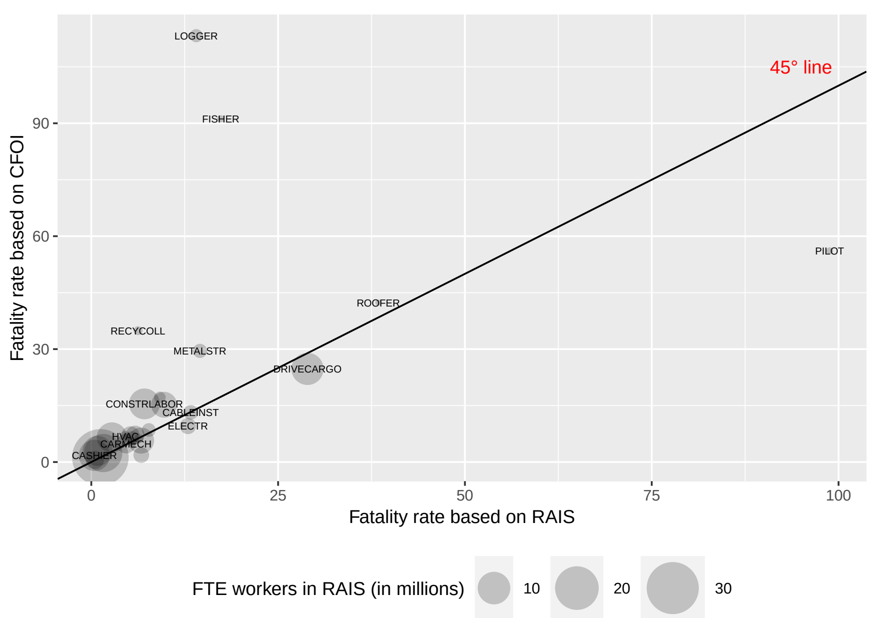

Selected occupations. CFOI between 2011 and 2017, RAIS between 2007 and 2017.

D Wage details

The norm in Brazil is for wages to be paid monthly. The RAIS has three ways in which wages have been consistently reported throughout my data: the base monthly pay, as stated in the work contract, the wage for the month of December and the average monthly wage for the period. The latter two may include bonus and other gratifications, but no severance payment. To make comparisons with the U.S. simpler, I have transformed the average monthly wage for the period into hourly wage using the number of hours contracted per week. I assume a year has 50 weeks of work, so that hourly wage is

$$HourlyWage = 12 \times AvgMonthlyWage/(50 \times HoursPerWeek)$$

I also converted nominal values to 2016 R$ using the IPCA (Índice Nacional de Preços ao Consumidor Amplo), Brazil's main price index. I then convert these values to 2016 US$ using the exchange rate at the time, 3.26R$/US$.

E Mandated risk premium

While the regulatory environment differs by contract type, the formal workers in the private sector considered here typically fall under the Consolidação das Leis de Trabalho (CLT), or Labor Law Consolidation. Concerning work safety, it establishes that firms are responsible for informing their workers of the potential risks of the job and for providing protective equipment that allows activities to be conducted within tolerable limits of risk, which are defined by the Ministry of Labor and assessed by audits that may be requested by the firms themselves, workers or their unions. If those limits are proved to be exceeded following an audit, workers may be granted two, not exclusive, types of wage premiums: an insalubrious one for exposure to hazardous chemical products, that varies from 10% to 40% of the minimum wage depending on the graveness of the situation; and a risk one for exposure to inflammable or explosive substances at excessive levels, which is worth an additional 30% of the worker's wage. Both premiums are only due if and while the firm fail to rectify the safety situation.

A natural concern is whether these rules distort the wage contract curve, in which case the hedonic wage regression would be confounding the effects of the regulation as a wage premium. My risk measure is not immediately connected to insalubrious conditions, because I am not including professional diseases in it. Nevertheless, it is still possible that insalubrious working conditions are correlated with fatality rates, which would bias the wage premium. The second mandated premium is an even clearer concern. If work conditions are dangerous for persons dealing with explosives and inflammable substances—and some occupations in the data do suggest they are—the risk premium is likely to be distorted by the regulation.

Even if the data do not explicitly list the nature of the fatal accident, the occupation description is often a good indicator. For example, "*workers in the manufacturing of ammunition and chemical explosives*", or "*workers in the extraction of solid minerals*" are likely to be subject to such explosion risks and are among the most dangerous occupations in my measure, with a fatality rate of 90.37 and 12.90 per 100,000 FTE workers, respectively. Therefore, these safety mandates present a further challenge in interpreting risk wage premium as workers' WTP for reductions of fatality risks.

F Is the magnitude of the quit reaction sensible?

In my main results, I find a 1.79 p.p. increase in the probability that a low tenure worker quits his job in the year in which a colleague dies in a work-related death, relative to low tenure coworkers of people who died for reasons unrelated to work. Is this magnitude sensible?

In order to investigate this, I did a simple calibration of the quitting probabilities to see if they can be rationalized with reasonable beliefs. Assuming that $\tilde{r} \sim Beta(a,b)$, and that priors are unbiased, i.e., $\frac{a}{a+b} = r$, I want to find values for (a,b) that are consistent with the quitting patterns I find. I do so separately for both low and high tenure workers. The quitting probability after a work-related death is

$$Q(\mathcal{I}, x = 1) = 1 - G\left(w^* - \gamma \frac{a+1}{a+b+1}\right)$$

From the data, I have the a baseline quitting probability $Q^0(\mathcal{I})$ and after death $Q^1(\mathcal{I}, x)$. I use my own estimate $\gamma = 0.140$, and the empirical distribution of log wages as G.[5] So with $\frac{a}{a+b} = r$, I can use $Q^0(\mathcal{I})$ to solve for w^*. Then, I use $Q^1(\mathcal{I}, x = 1)$ and $\frac{a}{a+b} = r$ to solve for (a, b). For low tenure workers, I have $Q^0(\mathcal{I}) = 10\%$, $Q^1(\mathcal{I}, x = 1) = 11.79\%$, which imply $a = 0.00185$ and $b = 15.08$. Note that with these parameters, the posterior following a death signal is $\frac{a+1}{a+b+1} = 6.23\%$, which is far higher than the unbiased level of 0.0123%.

For high tenure workers, I have $Q^0(\mathcal{I}) = 2\%$, but there is no increase in the quitting probability after a work-related death. In practice, I allow for $w^* - \gamma\frac{a+1}{a+b+1}$ to increase up to the closest point in the distribution G, which here would be the equivalent of a imperceptible change in the quitting probability. The implied parameters are $a = 0.03956$ and $b = 350.21$. The change from prior to posterior is far less dramatic, from 0.0113% to 0.296%.

The parameters of the Beta distribution can be roughly interpreted as the number of draws with which the beliefs are based upon, a of which were successes while b were failures. Therefore, these results imply that high tenure workers have roughly 335 additional signals on average. Given the extra three years of tenure and assuming a year has 50 weeks of work, this is equivalent to 2.23 signals per week, which I find to be a sensible number.

[5] Actually, G is the distribution of outside option value, which differs from log wages by $\gamma\mathbb{E}[\tilde{r}]$, which I ignore here as if workers take this outside option as risk-free.