

Workplace Injury Risk and the Gender Wage Gap

Francesco Del Prato, Salvatore Lattanzio

April 2026

Men experience workplace injuries at twice the rate of women. We study whether compensating differentials for injury risk contribute to gender differences in firm pay policies. We develop a search model microfounding an AKM wage equation that decomposes firm effects into productivity and injury-risk components. Using Italian administrative data with individual injury records, we estimate gender-specific firm wage effects and firm effects for injury risk. Injury channels explain 8 percent of the gender gap in firm effects, rising to 17 percent in manufacturing. Conditioning on occupation eliminates within-firm disparities, indicating the channel operates through between-firm sorting and within-firm occupational allocation.

KEYWORDS: Gender wage gap, workplace injuries, compensating differentials, AKM, rent sharing.

JEL CLASSIFICATION: J16, J28, J31, J64, J71.

Francesco Del Prato: *Department of Economics and Business Economics, Aarhus University*, francesco.delprato@econ.au.dk. Salvatore Lattanzio: *Bank of Italy, RFBerlin, and Dondena*, salvatore.lattanzio@bancaditalia.it. We thank Francesca Carta, Federico Cingano, Marta De Philippis, Nina Roussille and Rune Vejlin for helpful comments. The views expressed in this paper are those of the authors only and should not be attributed to the Bank of Italy nor to the Eurosystem.

1. Introduction

Men are injured on the job at roughly twice the rate of women. In Italy, the annual probability of a workplace accident generating at least one week of absence is 2 percent for men and 1 percent for women, with the gap widest in manufacturing and construction. If labor markets compensate workers for bearing injury risk (Rosen 1986), gender differences in exposure to risk—through sorting across firms, allocation to tasks within firms, or differences in how risk is priced—could contribute to the overall gender wage gap. Despite a large literature documenting firm-level sources of gender pay disparities (Card, Cardoso, and Kline 2016), workplace injury risk has received little attention as a potential channel. This paper provides the first assessment of how non-fatal workplace injury risk contributes to gender differences in firm pay policies.

This paper. Our goal is to quantify the role of workplace injury risk in explaining gender differences in firm wage policies. The analysis proceeds in four steps. First, we document empirical facts on gender differences in injury rates and their relationship to firm pay in Italy. Second, we develop a search model with heterogeneous firms and workers that clarifies how wages reflect both productivity and compensation for risk, and how gender differences can arise through sorting and bargaining. Third, we bring the reduced-form results derived from the model to rich Italian administrative data and estimate gender-specific firm wage effects alongside firm-level injury effects within an Abowd, Kramarz, and Margolis (1999) [AKM] framework. Finally, guided by the model, we decompose the gender gap in firm effects into between- and within-firm gender differences in productivity and risk. We use this structure to interpret gender differences in risk-related pay and to express their magnitude in simple monetary terms. Altogether, our results highlight a modest but nontrivial role for workplace injury risk in explaining gender differences in firm pay policies, with the contribution concentrated in manufacturing and blue-collar occupations.

The gender injury gap. In the first step, we leverage matched employer-employee administrative data from the Italian Social Security Administration (INPS), available at the Bank of Italy, covering the Italian private sector from 2005 to 2019 and based on a representative sample of firms.¹ Italy has a substantial gender wage gap of 18 log points in our sample, making it an ideal context to study gender inequality. A distinctive advantage of this setting is that INPS records detailed information on individual-level workplace injuries, allowing us to observe which workers experience accidents and at which firms. We document that

¹The data include detailed information on labor contracts, work schedules, and labor input for all workers employed by firms in the representative sample, and they track workers' careers over time, including transitions to firms outside the sample.

injury rates are strongly negatively correlated with wages across the distribution and vary substantially across sectors, occupations, and firm sizes, with men facing higher risk than women in nearly all categories. At first glance, this negative wage–risk relationship may seem at odds with compensating differentials. The key is that the raw correlation reflects both sorting (higher-earning workers tend to work in safer jobs and at safer firms) and wage-setting (whether riskier workplaces pay a premium, holding worker composition fixed).

Model. In the second step, we develop a continuous-time search model in which firms differ in productivity and in characteristics that determine workplace injury rates. Workers differ by gender in their risk cost and in bargaining power. Upon meeting, workers and firms bargain over wages, with the contract reflecting both the productivity surplus of the match and compensation for injury risk.

The model delivers two key reduced forms. First, it yields an exact two-way additive wage decomposition that microfound the canonical AKM regression. The firm effect decomposes into a productivity component—reflecting Nash bargaining over match surplus—and a risk component—capturing both the worker’s share of safety cost savings and direct compensation for risk cost. Second, the model provides a structural interpretation for a Kitagawa-Oaxaca-Blinder (KOB) decomposition of the gender gap in firm effects, distinguishing sorting effects (men and women matching to different types of firms) from wage-setting effects (men and women receiving different returns to firm characteristics).

Injury risk and wage policies. In the third step, we estimate gender-specific AKM wage regressions to separate a portable worker component of wages from a firm pay-policy component, estimating the wage model separately by gender. We then estimate an analogous two-way fixed effects model for injury incidence and interpret the firm component as a measure of workplace risk that nets out worker composition. Standard mobility-event studies (Card, Heining, and Kline 2013) show no systematic pre-trends around job changes and broadly symmetric post-move adjustments, supporting the view that this firm risk measure captures persistent differences in workplace safety rather than transitory worker shocks.²

Decomposition. In the fourth step, we take the estimated firm pay-policy components to the firm level and relate them to two firm attributes: productivity (proxied by value added per worker) and workplace risk (measured by the firm component from the injury model). Both attributes are positively associated with firm pay policies, consistent with rent sharing and compensation for risk, but the associations are weaker for women.

²We deem this validation as also a novel contribution.

Women receive approximately 94 percent of men’s wage response to firm productivity and 86 percent of men’s wage response to workplace injury risk. Throughout, these responses are the objects being “priced”: they describe how firm pay policies load on firm attributes, holding worker composition fixed, and should not be read as the unconditional wage-risk correlation or as a direct measure of welfare. To gauge magnitudes, our estimates imply that reducing the annual probability of an observed injury by 1 percentage point is associated with a wage-equivalent value of about 440 euros per male worker-year and 250 euros per female worker-year.

Using these estimates, we implement a Kitagawa–Oaxaca–Blinder decomposition motivated by the theoretical model. In the dual-connected sample, the overall gender wage gap is 22 log points, of which about one-third is explained by differences in firm wage effects. Within this firm-effects gap, productivity-related channels account for approximately 61 percent, driven primarily by sorting across firms and, to a lesser extent, by within-firm wage-setting differences. Injury-related channels account for approximately 8 percent of the gap, evenly split between sorting and within-firm margins. We document substantial heterogeneity: the contribution of injury risk is concentrated in manufacturing, where it explains up to 17 percent of the firm-effects gap, primarily through within-firm margins that reflect differential allocation to risky tasks.

These patterns reflect genuine firm-level differences rather than occupational sorting (Goldin 2014). When we estimate firm-occupation effects and decompose the gap separately for blue-collar and white-collar workers, the within-firm component of the injury channel falls to essentially zero, while the between-firm component remains positive and meaningful. This indicates that, even conditional on broad occupation, women tend to work at firms offering lower injury-risk premia. Particularly, within blue-collar occupations the entire contribution of injury premia operates through the between-firm component: the results are not solely driven by men and women sorting into different occupations, but by differences in injury-related premia across firms.

Contribution to the literature. This paper contributes to three strands of the literature.

First, we contribute to the literature on the firm-level determinants of the gender wage gap.³ Card, Cardoso, and Kline (2016) decompose gender differences in firm effects into sorting and bargaining components, finding that firm-specific pay differentials explain about one-fifth of the gender wage gap in Portugal—a result extended to several other countries and institutional contexts (Bruns 2019; Casarico and Lattanzio 2024; Palladino, Roulet, and Stabile 2025; Palladino, Bertheau, Hijzen et al. 2025). Morchio and Moser (2024) develop an equilibrium search model showing that compensating differentials for nonpay amenities explain half of the gender pay gap in Brazil. We contribute by identifying a

³For reviews, see Blau and Kahn (2017); Olivetti, Pan, and Petrongolo (2024); Casarico and Lattanzio (2025).

specific, measurable disamenity—workplace injury risk—and quantifying its contribution to gender differences in firm pay policies. No single amenity is expected to explain a large share of the gender wage gap: Le Barbanchon, Rathelot, and Roulet (2021) estimate that gender differences in commute valuation account for 9 to 15 percent of the residualized wage gap through a structural search model, and Wiswall and Zafar (2018) find that gender differences in preferences for earnings growth, job security, and hours flexibility together explain about one quarter of the early-career gap.⁴ Our finding that a single disamenity accounts for 8 percent of the firm-effects gap—rising to 17 percent in manufacturing, or 5.4 percent of the total manufacturing wage gap—is of a comparable order of magnitude to these single-amenity benchmarks, while being the first to be structurally identified and decomposed within the AKM framework.

Second, we contribute to the literature on compensating wage differentials for job risk (Rosen 1974, 1986).⁵ Hersch (1998) shows that industry-level risk measures obscure wage-risk trade-offs for women, motivating our use of individual-level injury data. Garen (1988) addresses the endogeneity of job risk arising from worker selection into safer jobs. Guardado and Ziebarth (2019) show that because workers can supply safety to firms, individual-level risk is negatively correlated with wages, which downward-biases estimates of compensating differentials. On the firm side, Lavetti and Schmutte (2025) show that hedonic models fail to identify risk premia unless they incorporate firm wage-setting. Extending this approach to gender, Lavetti and Schmutte (2023) find no gender differences in the pricing of fatal risk in Brazil once firm effects are included. We build on these insights by addressing both sources of unobserved heterogeneity within an AKM framework. Guided by our search model, we isolate firm-level risk from worker heterogeneity to estimate gender-specific premia for non-fatal injury risk. In contrast to the null finding for fatalities, we document that women receive significantly smaller non-fatal risk premia than men at comparable firms.

Third, we contribute to the measurement of firm-level injury risk. While Picchio and Van Ours (2017) estimate worker and firm fixed effects for injury severity using Italian data, we estimate an AKM decomposition for injury incidence and link the resulting firm risk measure to firm pay policies, allowing for a productivity–risk decomposition of the gender gap in firm effects.

Outline. The remainder of the paper proceeds as follows. Section 2 describes the Italian institutional setting and the administrative data. Section 3 develops the theoretical model and derives the structural wage decomposition. Section 4 presents the empirical frame-

⁴Mas and Pallais (2017) estimate that the average worker’s willingness to pay for working from home is 8 percent of wages, with women—particularly those with young children—valuing it more.

⁵For a recent survey, see Lavetti (2023). On the theoretical side, Kerndler (2023) shows that frictions reduce optimal safety provision by shortening expected job durations.

work. Section 5 reports the main results and heterogeneity analyses. Section 6 provides an interpretation of the findings in currency units and relates them to welfare considerations. Section 7 concludes.

2. Facts on the injuries gap

We begin by describing the institutional setting for workplace injuries in Italy, introducing the administrative data and our injury measure, and documenting the descriptive patterns that motivate the model and empirical strategy.

2.1. Injuries' legislative framework

In Italy, a work injury is defined as an event occurring “for cause and on the occasion of work,” requiring a direct or indirect causal link between the worker’s activity and the harmful event. Although this notion originally referred to sudden and violent accidents, its scope has widened over time to cover a broader range of work-related harms, including physiological or psychological reactions to fatigue and stress.

Following the 2000 reform (D.Lgs. 38/2000), the definition was further extended to include commuting injuries (“*infortuni in itinere*”), thereby enlarging the set of events eligible for protection.

Occupational injuries are covered by a mandatory insurance system administered by the National Institute for Insurance against Occupational Accidents and Diseases (INAIL), which provides medical care and income replacement. Employers bear the cost of the first three days of absence, after which INAIL pays 60 percent of the worker’s average daily wage from day 4 to day 90 and 75 percent thereafter until recovery. Employers must report injuries to INAIL within two days. While the Italian Social Security Institute (INPS) does not pay injury benefits, it records the corresponding periods of absence through notional social-security contributions, ensuring that injury-related nonwork spells are credited toward workers’ pension histories.

2.2. Data

Matched employer-employee data. We use matched employer-employee administrative data from INPS archives, available at the Bank of Italy, covering the years 2005 to 2019. The dataset includes a representative sample of firms and the universe of multinational companies operating in the private sector in Italy. For each firm in the representative sample, we observe the full workforce and follow all employees over time, including those who change jobs or move across firms. In our final analysis, we focus on firms in the representative sample, excluding multinationals not drawn into it.

The worker-level data provide detailed information on employment relationships. For each job spell, we observe annual labor earnings and the number of weeks worked (both raw and in full-time equivalent terms). The data also record contract type (permanent, fixed-term, or seasonal), work schedule (full-time or part-time), occupation (blue-collar, white-collar, manager, apprentice, or other), and municipality of work. We observe workers' entry dates, as well as separation dates and reasons for termination where applicable.

The dataset also includes demographic characteristics of workers—gender, year of birth, migrant status, and province of residence—along with detailed firm-level information. For all firms, we observe the sector of activity and the dates of establishment and, if applicable, closure. For incorporated firms, the INPS data are linked to balance-sheet records from Cerved, providing measures such as value added and sales.

Injury data. Beyond standard employment information, the dataset contains records of events that give rise to notional social-security contributions, such as work-related injuries, sick leave, short-time work compensation, and parental or maternity leave. We focus on work-related injuries, which are recorded separately, and for each we observe the year of occurrence and the corresponding notional contribution. A key feature of our data is that INPS records work-related injuries only when they last at least seven consecutive days, as the week is the standard unit of computation for pension benefits. Shorter injuries, though compensated by INAIL, do not appear in the social-security archives. Consequently, the injury events we observe therefore correspond to accidents severe enough to trigger both INAIL compensation and INPS notional-contribution coverage. There are, however, only minor differences in the sectoral distribution of injuries between those lasting less than 7 days and those lasting at least 7 days. Using publicly available INAIL data for the period 2014–2018, Figure C1 reports the distribution of injuries across sectors by duration. The figure shows that restricting attention to longer-lasting injuries does not materially alter the composition of injuries under analysis, as the two distributions are very similar.

Sample selections and restrictions. Throughout the analysis, we rely on two complementary versions of these data. First, we use the full employment biographies of workers from 2005 to 2019, which include all job spells regardless of whether they occur inside or outside the representative firm sample. Second, we construct a firm-level panel by aggregating these worker histories, restricting attention to firms that (i) employ at least one man and one woman over the 2005–2019 period, (ii) belong to the largest connected component for both genders (see below for details), and (iii) have non-missing balance-sheet information.

We impose several sample restrictions. First, we limit the analysis to workers aged 20 to 60 (dropping around 3 percent of total person-year observations in the raw data). Second, we drop observations with missing or zero weeks worked and those with missing, zero or

negative earnings (3 percent of initial sample size). Third, we retain a single observation per worker-year, selecting the contract with the highest annual earnings; any remaining duplicate worker-year observations are removed at random (dropping 17.7 percent of initial observations).

Description of worker-level outcomes. Our worker-level analysis focuses on two main outcomes. The primary measure of pay is the log of full-time equivalent (FTE) weekly earnings, defined as the ratio between annual earnings—which include all forms of pecuniary compensation, grossed up to reflect labor income taxes and employee social-security contributions—and FTE weeks worked, computed as actual weeks worked multiplied by the ratio of monthly hours to the standard full-time hours specified in the relevant collective agreement.⁶ The second outcome is an indicator for whether a worker experiences at least one injury in a given year. As described above, this captures work-related accidents that generate notional social-security contributions, and thus reflects injuries lasting at least seven consecutive days.

2.3. Descriptive statistics

Sample composition. Table 1 reports means and standard deviations of the worker-level data. We first describe the full dataset (column "All"), which includes firms outside the representative sample where workers may move during their careers. Statistics are shown separately for men and women. Men are slightly older and earn higher annual earnings and weekly wages; the gap in log full-time-equivalent weekly earnings is 18 log points. Women are more likely to work part-time (37 vs. 9 percent) and less likely to hold permanent contracts (75 vs. 81 percent for men). They are also more frequently employed in the Center-North and in services. The average injury probability is 2 percent for men and 1 percent for women. Figure C2 in the Appendix shows that male injury rates declined over time, particularly between 2007 and 2013, while female injury rates remained fairly stable at around 1 percent throughout 2005-2019. The higher injury rate among men is also linked to their greater likelihood of experiencing more than one accident over their working lives: the share of men with multiple injuries is more than twice that of women (Figure C3).

In the representative sample (second column of Table 1, "BDI"), we observe the full workforce of each included firm, which implies a mechanical increase in average firm size (16 vs. 11 employees). Workers in this sample hold more stable positions and earn higher annual earnings, and the gender gap rises to 21 log points. The geographical distribution

⁶This measure corresponds to a wage rate in the absence of overtime. Missing information on hours of work may nonetheless introduce measurement error if men and women differ systematically in overtime or in the timing of hours worked (Goldin 2014). Despite these limitations, the dataset provides detailed longitudinal information on workers' careers, earnings, and employers, allowing us to track individuals over time.

TABLE 1. Descriptive statistics of the samples

<i>Variable</i>	All		BDI		DC-BDI	
	(1) Men	(2) Women	(3) Men	(4) Women	(5) Men	(6) Women
Age	39.70 (10.37)	38.33 (9.96)	41.45 (10.28)	40.12 (9.97)	41.48 (10.28)	40.11 (9.97)
Immigrant	0.12 (0.32)	0.11 (0.31)	0.09 (0.29)	0.09 (0.29)	0.09 (0.28)	0.09 (0.29)
Real annual earnings	26,558.24 (23,190.13)	17,733.50 (15,419.56)	30,434.50 (23,556.99)	20,198.03 (16,044.32)	30,886.36 (23,927.55)	20,320.27 (16,212.16)
Weeks worked	40.93 (16.02)	33.74 (17.24)	44.55 (13.70)	36.69 (16.33)	44.58 (13.69)	36.68 (16.37)
Log weekly earnings	6.26 (0.49)	6.08 (0.46)	6.36 (0.44)	6.15 (0.43)	6.37 (0.44)	6.15 (0.43)
Part-time	0.09 (0.29)	0.37 (0.48)	0.06 (0.25)	0.37 (0.48)	0.06 (0.25)	0.36 (0.48)
Permanent contract	0.81 (0.39)	0.75 (0.44)	0.87 (0.34)	0.82 (0.39)	0.86 (0.34)	0.82 (0.39)
Any event	0.26 (0.44)	0.27 (0.44)	0.30 (0.46)	0.32 (0.46)	0.30 (0.46)	0.32 (0.47)
Injury event	0.02 (0.14)	0.01 (0.10)	0.02 (0.14)	0.01 (0.11)	0.02 (0.14)	0.01 (0.11)
Sick leave	0.10 (0.30)	0.10 (0.30)	0.10 (0.30)	0.11 (0.31)	0.10 (0.30)	0.11 (0.31)
Centre-North	0.82 (0.38)	0.88 (0.33)	0.81 (0.39)	0.88 (0.32)	0.82 (0.38)	0.89 (0.32)
South	0.18 (0.38)	0.12 (0.33)	0.19 (0.39)	0.12 (0.32)	0.18 (0.38)	0.11 (0.32)
Industry	0.43 (0.50)	0.26 (0.44)	0.52 (0.50)	0.33 (0.47)	0.52 (0.50)	0.33 (0.47)
Construction	0.06 (0.25)	0.01 (0.10)	0.03 (0.17)	0.00 (0.07)	0.03 (0.16)	0.00 (0.07)
Knowledge-int. services	0.11 (0.31)	0.17 (0.37)	0.08 (0.28)	0.10 (0.30)	0.09 (0.28)	0.10 (0.30)
Other services	0.39 (0.49)	0.56 (0.50)	0.36 (0.48)	0.56 (0.50)	0.36 (0.48)	0.57 (0.50)
Firm size (employees)	11.44 (17.98)	11.67 (18.80)	16.37 (21.03)	16.45 (21.29)	18.51 (23.07)	18.49 (23.19)
Num of workers	7,934,782	5,186,451	5,096,377	3,170,060	4,850,105	3,073,710
Num of firms	1,209,960	1,043,359	98,506	96,348	76,165	76,165
Person-year obs.	68,516,680	41,175,796	31,985,368	17,736,132	30,296,530	17,084,024

Note: The table reports means and standard deviations of the worker-level dataset. The first two columns covers the full sample, including workers employed in firms inside the representative sample as well as those in firms outside it, where workers may move during their careers. The second two columns restricts the sample to workers employed in firms in the representative sample. The third two columns further narrows the sample to workers in firms that belong both to the representative sample and to the dual connected component, defined as the intersection of the largest connected components for male and female workers.

remains similar, while the sectoral composition becomes more concentrated in industry relative to the full dataset.

Having characterized the samples, we now turn to descriptive patterns in injury rates.

Injury rates. Figure 1 reports average injury rates by gender across worker and job characteristics. Men have higher injury rates at all ages, but the male–female gap narrows at older ages: injury risk declines for men, while it rises slightly for women (panel 1A). Injury risk is also negatively correlated with wages (panel 1B): high-wage workers are less likely to experience an accident, as they tend to be employed in lower-risk occupations. The gender gap in injury risk effectively disappears in the top decile of the weekly wage distribution.

For men, injury rates increase with firm size up to very large firms (250 employees or more), whereas for women they vary little across size categories (panel 1C). Injury probabilities are higher for both genders in blue-collar and apprentice occupations, and the male–female gap widens in these groups. Workplace accidents are also more common in industry and construction, where the gender gap is largest, whereas differences are smaller in services—particularly non–knowledge-intensive services—where men’s injury rates are closer to those in industry (panel 1E).


Figure C4 provides a more detailed breakdown by industry, showing especially high injury rates in several manufacturing sectors, in energy and gas utilities, in construction, and in health services. Injury rates are also higher in the North-East, reflecting that region’s more industry-intensive employment structure (panel 1F).

Figure C5 reports average injury probabilities by contract type and work schedule. Workers on permanent contracts are more likely to experience a work-related accident, as shown in panel C5A.⁷ Among men, full-time workers have higher injury rates than part-time workers, whereas among women the pattern is reversed, though the difference is small (panel C5B).

As a final piece of descriptive evidence, Figure C6 shows the distribution of firms by the number of workplace accidents recorded over 2005–2019. About one third of firms experience exactly one accident, and roughly one fifth fall into each of the categories of 2, 3–4, and 5–10 accidents. A sizeable 15 percent of firms report more than 10 accidents during the period, indicating a substantial concentration of workplace injuries in a relatively small group of firms.

⁷This result is consistent with Picchio and van Ours (2017), who show that work-related accidents occur more often among permanent full-time workers, though they are generally more severe among temporary workers.

FIGURE 1. Injury rate by worker and job characteristics, by gender



Note: The figure reports average annual injury rates by gender across worker and job characteristics, computed over the full sample (2005–2019). Panel A groups workers into five-year age bins. Panel B shows injury rates by decile of the full-time-equivalent wage distribution, computed separately by gender. Panel C uses the following firm-size classification: micro (<10 employees), small (10–49), medium (50–249), and large (250+). Panel D distinguishes five occupational categories: other, apprentice, executive, white-collar, blue-collar, and middle manager. Panel E groups two-digit NACE sectors into four broad categories: industry (mining and manufacturing), construction, knowledge-intensive services, and other services. Panel F reports rates by broad geographical area. Injury is defined as experiencing at least one work-related accident generating notional social-security contributions, corresponding to absences of seven or more consecutive days.

3. Theoretical model

Guided by the descriptive evidence, this section develops a search model in which wages reflect both productivity and compensation for injury risk, and derives the reduced-form decompositions we bring to the data. Technical derivations and proofs are in Appendix A.

3.1. Setting

Environment. Time is continuous. The economy is populated by a continuum of infinitely-lived firms and workers who discount the future at a common rate r .

Firms. Firms are heterogeneous in productivity p and injury risk λ , with type (p, λ) drawn from the joint distribution $\Gamma(p, \lambda)$. We treat (p, λ) as exogenous firm primitives. In reduced form, λ summarizes a firm's technology, safety culture, and regulatory compliance environment. Firms post vacancies at flow cost k_v .

Workplace risk. Workers face the risk of workplace accidents. A firm with injury rate λ experiences accidents at Poisson rate λ . Maintaining a lower injury rate is costly: a firm with rate λ incurs flow cost $c(\lambda)$, with $c'(\lambda) < 0$. When an accident occurs, the worker enters an “injured” state: productivity falls to $p^I = \alpha p$ for some $\alpha \in (0, 1)$, and the worker receives a reduced income stream $w_g^I(p, \lambda) \leq w_g(p, \lambda)$ (e.g., from sick pay or public disability insurance). The worker recovers to the healthy state at an exogenous rate η .

Workers. Workers are characterized by gender $g \in \{M, F\}$ and are either employed or non-employed. Preferences are linear in income, but workers dislike accident risk in a gender-specific way, with $\phi_g \geq 0$ capturing the worker's risk cost—a composite of risk preferences and effective exposure to hazardous tasks (reflecting fear, physical vulnerability, family responsibilities, task assignment, etc.). A non-employed worker i of gender g receives a flow payoff $b_{g,i}$, which we allow to be individual-specific. This parsimoniously captures persistent, portable worker heterogeneity (e.g. skill/ability, constraints, preferences) and will map directly into the worker component of the AKM decomposition in Section 3.4. An employed worker of gender g at a firm (p, λ) produces flow output p , receives wage $w_g(p, \lambda)$, and experiences flow payoff

$$u_g(p, \lambda) = w_g(p, \lambda) - \phi_g \lambda.$$

Matching. We assume no on-the-job search; jobs end at an exogenous rate δ , regardless of the worker's health state. The number of matches is governed by a constant-returns-to-scale matching function $M(U_g, V)$, where U_g is the measure of unemployed workers of

gender g and V is the measure of vacancies. Defining market tightness as $\theta_g \equiv V/U_g$, the job-finding rate for workers is $f(\theta_g)$ and the vacancy-filling rate is $q(\theta_g)$.

3.2. Value functions

An employed worker can be either healthy or injured. Let $W_g^H(p, \lambda)$ and $W_g^I(p, \lambda)$ be the value to a worker of gender g of being healthy and injured, respectively, while employed at a firm (p, λ) . Let $W_{g,i}^0$ be the value of non-employment for worker i of gender g .⁸ Similarly, let $J_g^H(p, \lambda)$ and $J_g^I(p, \lambda)$ be the firm's value of a filled job with a healthy or injured worker of gender g , respectively, and J^0 the value of a vacancy. The Hamilton-Jacobi-Bellman (HJB) equations that define these values in steady state are as follows.

Healthy employed worker. A healthy employed worker of gender g at firm (p, λ) produces p , receives wage $w_g(p, \lambda)$, and bears risk cost $\phi_g \lambda$ each instant. The job is destroyed at rate δ , moving the worker to non-employment, and an accident occurs at rate λ , moving the worker to the injured state. The HJB is

$$rW_g^H(p, \lambda) = [w_g(p, \lambda) - \phi_g \lambda] + \delta [W_{g,i}^0 - W_g^H(p, \lambda)] + \lambda [W_g^I(p, \lambda) - W_g^H(p, \lambda)].$$

Injured employed worker. An injured worker receives reduced income $w_g^I(p, \lambda)$, still bears risk cost $\phi_g \lambda$, faces the same exogenous destruction rate δ , and recovers to the healthy state at rate η . The HJB is

$$rW_g^I(p, \lambda) = [w_g^I(p, \lambda) - \phi_g \lambda] + \delta [W_{g,i}^0 - W_g^I(p, \lambda)] + \eta [W_g^H(p, \lambda) - W_g^I(p, \lambda)].$$

Unemployed worker. A non-employed worker i of gender g receives flow payoff $b_{g,i}$ and receives job offers (which start in the healthy state) at rate $f(\theta_g)$. All offers that generate non-negative surplus relative to non-employment are accepted. The HJB is

$$rW_{g,i}^0 = b_{g,i} + f(\theta_g) \mathbb{E}_{p,\lambda} [\max \{W_g^H(p, \lambda) - W_{g,i}^0, 0\}].$$

We interpret $W_{g,i}^0$ (equivalently $b_{g,i}$) as a reduced-form sufficient statistic for time-invariant, portable worker characteristics that affect wages through the outside option.

Participation after injury. We assume parameters and contracts are such that both health states are individually rational for workers in active matches, i.e.

$$W_g^H(p, \lambda) \geq W_{g,i}^0 \quad \text{and} \quad W_g^I(p, \lambda) \geq W_{g,i}^0$$

⁸We allow $W_{g,i}^0$ to vary across workers to capture time-invariant individual heterogeneity in non-employment values.

for all (p, λ, g) that arise in equilibrium. If $W_g^I(p, \lambda) < W_{g,i}^0$ for some match, the worker would strictly prefer to leave upon injury; in that case one can interpret separation at injury as part of the exogenous destruction process δ . For expositional clarity we abstract from this margin and restrict to matches satisfying the inequalities above.

Vacancy. Free entry drives the value of a vacancy to zero, $J^0 = 0$. We solve the model separately for each gender g , allowing for gender-specific market tightness θ_g and job-creation conditions. For a given gender g ,

$$k_v = q(\theta_g) \mathbb{E}_{p,\lambda} [J_g^H(p, \lambda)],$$

where the expectation is over successful matches, which begin with a healthy worker.

Firm with a filled job. The value of a filled job for a firm with a worker of gender g depends on the worker's health state:

$$rJ_g^H(p, \lambda) = p - w_g(p, \lambda) - c(\lambda) - \delta J_g^H(p, \lambda) + \lambda [J_g^I(p, \lambda) - J_g^H(p, \lambda)]$$

and with an injured worker,

$$rJ_g^I(p, \lambda) = \alpha p - w_g^I(p, \lambda) - c(\lambda) - \delta J_g^I(p, \lambda) + \eta [J_g^H(p, \lambda) - J_g^I(p, \lambda)].$$

We maintain throughout that injuries are costly for firms: the output loss from an injured worker exceeds any wage savings, so that $J_g^H(p, \lambda) > J_g^I(p, \lambda)$. Subtracting the two HJBs, this requires

$$(1) \quad (1 - \alpha)p > w_g(p, \lambda) - w_g^I(p, \lambda)$$

for all matches in equilibrium. Under the state-independent wage normalization $w^I = w$ adopted below, condition (1) is automatically satisfied since $(1 - \alpha)p > 0$.⁹

3.3. Wages

Bargaining over the surplus. Upon meeting in the healthy state, the firm and the worker bargain once over a state-contingent wage contract $\omega_g(p, \lambda) \equiv ((w_g(p, \lambda), w_g^I(p, \lambda)))$, taking (p, λ) and the healthy-state Markov process as given and ruling out subsequent re-bargaining after shocks. Define the healthy-state continuation surplus of the match as

$$S_g^H(p, \lambda; \omega_g) \equiv [W_g^H(p, \lambda; \omega_g) - W_{g,i}^0] + [J_g^H(p, \lambda; \omega_g) - J^0],$$

⁹In extensions where firms can reduce λ through costly safety investments, the gap $J_g^H - J_g^I$ determines the marginal benefit of lowering injury risk.

which already internalizes expected transitions to injury and recovery through continuation values. The contract solves a generalized Nash program in continuation values:

$$\omega_g(p, \lambda) \in \arg \max_{(w, w^I)} \left[W_g^H(p, \lambda; w, w^I) - W_{g,i}^0 \right]^{\beta_g} \left[J_g^H(p, \lambda; w, w^I) - J^0 \right]^{1-\beta_g}$$

subject to participation $W_g^H(p, \lambda; w, w^I) \geq W_{g,i}^0$ and the HJB constraints that determine $\{W_g^H, W_g^I, J_g^H, J_g^I\}$ given (w, w^I) . This yields first-order conditions that implicitly pin down $(w_g(p, \lambda), w_g^I(p, \lambda))$.

Because wages are pure transfers between the worker and the firm, they drop out of S_g^H : the joint surplus $S_g^H(p, \lambda)$ depends on (p, λ) and the worker's risk cost ϕ_g , but not on the particular split (w_g, w_g^I) . In equilibrium Nash bargaining therefore implies the familiar surplus split

$$W_g^H - W_{g,i}^0 = \beta_g S_g^H, \quad J_g^H - J^0 = (1 - \beta_g) S_g^H,$$

with S_g^H defined at the match level and gender-specific through ϕ_g .

PROPOSITION 1. *Fix a firm-worker pair (p, λ, g) and the one-shot healthy-state bargaining protocol described above (no re-bargaining after shocks), with state-contingent contracted incomes $(w_g(p, \lambda), w_g^I(p, \lambda))$. In any steady state consistent with the HJB equations, the equilibrium healthy-state wage satisfies*

$$(2) \quad w_g(p, \lambda) = \underbrace{rW_{g,i}^0}_{\text{outside option flow}} + \underbrace{(r + \delta) \left[W_g^H(p, \lambda) - W_{g,i}^0 \right]}_{\text{flow value of match surplus}} + \underbrace{\lambda \left[W_g^H(p, \lambda) - W_g^I(p, \lambda) + \phi_g \right]}_{\text{injury risk compensation}}.$$

The wage has three components. First, the outside option component is the flow value of the worker's non-employment option. It sets the baseline: any acceptable job must at least match this value. Second, the value of match surplus converts the stock gain from being in the match (relative to non-employment) into a flow. Third, the total risk compensation accounts for the job risk. It has two pieces inside the bracket: the capital loss in continuation value if an injury actually occurs (lower productivity and lower sick pay); and the per-unit-hazard cost of being exposed to risk even when no accident occurs. Multiplying by the accident arrival rate converts these into an expected flow per unit time. The wage must therefore be high enough that, per unit of time, the worker is compensated for both: *i*) the expected drop in future value when an injury arrives, and *ii*) the fact that bearing risk in the workplace carries a cost of intensity ϕ_g per unit of hazard.

3.4. A two-way fixed effect reduced form for productivity and risk

We now exploit the linear structure of preferences introduced above to derive an exact AKM-style decomposition of wages. Fix a gender g with risk-cost parameter ϕ_g and consider a worker i with outside value $W_{g,i}^0$. Because wages are pure transfers between firm and worker, they cancel out of joint values; the only non-transferable component is the flow risk cost $\phi_g \lambda$ borne by the worker at risky firms.

Joint surplus. For a match (p, λ, g) , define the gender-specific joint continuation surpluses

$$S_g^H \equiv (W_g^H - W_{g,i}^0) + J_g^H, \quad S_g^I \equiv (W_g^I - W_{g,i}^0) + J_g^I.$$

These objects summarize the total gains from the match in each health state. The next lemma characterizes the surplus as a function of firm and worker primitives.

LEMMA 1 (Surplus characterization). *The healthy-state surplus is affine in productivity:*

$$(3) \quad S_g^H(p, \lambda; W_{g,i}^0) = \kappa_p(\lambda) p - \frac{1}{r + \delta} [c(\lambda) + \phi_g \lambda + r W_{g,i}^0],$$

where the productivity coefficient is

$$(4) \quad \kappa_p(\lambda) \equiv \frac{1}{r + \delta} \cdot \frac{(r + \delta + \eta) + \alpha \lambda}{(r + \delta + \eta) + \lambda}.$$

The surplus has an intuitive structure. Higher productivity p raises surplus, with the coefficient $\kappa_p(\lambda)$ capturing how accident risk attenuates productivity's contribution: at riskier firms (higher λ), workers spend more expected time in the low-productivity injured state, reducing the effective weight on p . Both safety costs $c(\lambda)$ and the worker's risk cost $\phi_g \lambda$ reduce surplus, discounted at the constant rate $r + \delta$ that reflects impatience and job destruction.

State-independent wages. To derive a wage-in-levels two-way fixed effect representation, we impose the *state-independent contracted wages* normalization:

$$w_g^I(p, \lambda) = w_g(p, \lambda).$$

With linear utility and full commitment, Nash bargaining pins down the division of joint match surplus but does not uniquely discipline the within-match insurance profile across health states. We therefore select the full-insurance profile $w^I = w$ as a convenient normalization that yields a transparent AKM mapping.¹⁰

¹⁰This normalization concerns the *contracted* wage rate (or the within-match transfer). In the data, realized earnings during injury may still fall because injured workers supply fewer hours, leave the workplace, or are

Under this normalization (and identical flow risk cost across health states), subtracting the injured-state worker HJB from the healthy-state worker HJB implies $W_g^H = W_g^I$.

PROPOSITION 2 (AKM representation). *Suppose wages are state-independent ($w_g^I = w_g$) and Nash bargaining with parameter $\beta_g \in (0, 1)$ determines the surplus split. Then for worker i of type g employed at firm j with characteristics (p_j, λ_j) , the wage admits an exact two-way additive decomposition:*

$$(5) \quad w_{ij} = \underbrace{(1 - \beta_g)rW_{g,i}^0}_{\alpha_i^g} + \underbrace{(r + \delta)\beta_g\kappa_p(\lambda_j)p_j - \beta_gc(\lambda_j) + (1 - \beta_g)\phi_g\lambda_j}_{\psi_j^g},$$

The decomposition (5) provides a structural microfoundation for a standard AKM regression. The *worker effect* $\alpha_i^g = (1 - \beta_g)rW_{g,i}^0$ is the portable component of earnings tied to the outside option, scaled by $(1 - \beta_g)$: workers with stronger bargaining power rely relatively less on their outside option and more on intra-match surplus sharing. Although we model worker heterogeneity through outside options, this is without loss for our empirical purpose: the worker fixed effect absorbs any time-invariant portable component of pay, including permanent skill/ability and other persistent worker attributes. Accordingly, we interpret $rW_{g,i}^0$ as a reduced-form index of portable heterogeneity and do not attempt to distinguish its underlying sources within the model.

The *firm effect* ψ_j^g aggregates three distinct components:

- (i) *Productivity pass-through:* $(r + \delta)\beta_g\kappa_p(\lambda_j)p_j$ captures how firm productivity transmits into wages through the Nash share β_g and the risk-adjusted surplus loading $\kappa_p(\lambda_j)$.
- (ii) *Safety-cost drag:* $-\beta_gc(\lambda_j)$ reflects that maintaining a lower injury rate is costly and reduces the surplus available for sharing; the worker bears fraction β_g of this cost.
- (iii) *Risk compensation:* $(1 - \beta_g)\phi_g\lambda_j$ is the portion of risk cost that appears directly in wages. Workers with higher risk cost ϕ_g command higher compensation at riskier firms. The coefficient $(1 - \beta_g)$ reflects that fraction β_g of the risk cost is already absorbed through reduced surplus S_g^H ; the remaining $(1 - \beta_g)$ must be compensated directly in wages.

REMARK 1 (Decomposing the firm effect). *Collecting terms, the firm effect admits the decomposition*

$$\psi_j^g = \underbrace{(r + \delta)\beta_g\kappa_p(\lambda_j)p_j}_{\text{productivity}} + \underbrace{-\beta_gc(\lambda_j)}_{\text{safety-cost}} + \underbrace{(1 - \beta_g)\phi_g\lambda_j}_{\text{risk-compensation}}.$$

In an empirical AKM regression, the estimated firm fixed effect aggregates productivity-driven and risk-related components of wage determination.

paid through external benefit systems subject to caps and eligibility rules. The model can therefore accommodate reduced observed income during injury even when the within-match wage schedule is normalized to be state-independent.

3.5. Decomposing the gender gap in firm effects

We now move from the structural wage regression to its estimable empirical counterpart. This way, we allow for a direct structural interpretation of the empirical specifications discussed in Section 4.

Pooling the risk-related components. The firm effect contains two terms related to injury risk: the safety-cost drag $-\beta_g c(\lambda_j)$ and the risk compensation $(1 - \beta_g)\phi_g \lambda_j$. Both depend solely on the firm's injury rate λ , so we can collect them into a composite risk component:

$$(6) \quad \psi_j^g = \underbrace{(r + \delta)\beta_g \kappa_p(\lambda_j) p_j}_{\text{productivity}} + \underbrace{R_g(\lambda_j)}_{\text{risk}},$$

where

$$R_g(\lambda) \equiv -\beta_g c(\lambda) + (1 - \beta_g)\phi_g \lambda.$$

Equation (6) is exact but is generally nonlinear in λ_j through both $\kappa_p(\lambda_j)$ and $R_g(\lambda_j)$. To solve this, we impose two approximations. First, we assume that variation in $\kappa_p(\lambda)$ across firms is limited, so that $\kappa_p(\lambda_j) \approx \bar{\kappa}_p$ for some constant $\bar{\kappa}_p > 0$. This holds when accident rates do not vary too widely across firms, or when injuries have modest effects on productivity (α close to 1). Second, we assume that $R_g(\lambda)$ is approximately linear. A sufficient condition for this is that safety costs are approximately linear in the injury rate: $c(\lambda) \approx c_0 - c_1 \lambda$ for constants $c_0 > 0$ and $c_1 > 0$ (consistent with $c'(\lambda) < 0$). Then

$$R_g(\lambda) \approx -\beta_g(c_0 - c_1 \lambda) + (1 - \beta_g)\phi_g \lambda = -\beta_g c_0 + [\beta_g c_1 + (1 - \beta_g)\phi_g] \lambda.$$

The constant $-\beta_g c_0$ is absorbed into the regression intercept, and the firm effect simplifies to

$$(7) \quad \psi_j^g \approx \gamma_g p_j + \delta_g \lambda_j,$$

where

$$(8) \quad \gamma_g \equiv (r + \delta)\beta_g \bar{\kappa}_p, \quad \delta_g \equiv \beta_g c_1 + (1 - \beta_g)\phi_g.$$

The coefficient γ_g captures productivity pass-through: the worker's share β_g of surplus, scaled by the discount-adjusted productivity loading. The coefficient δ_g captures the net effect of injury risk on wages, combining the two components discussed in Proposition 2. Both forces imply $\delta_g > 0$: wages are higher at riskier firms, consistent with a compensating

differential.¹¹

Appendix B provides a calibrated numerical example that validates the linear approximation.

3.6. Roadmap

The model delivers a direct mapping from theoretical objects to estimable quantities. Worker heterogeneity enters through outside options and shows up empirically as a portable worker component of wages, while firm heterogeneity shows up as a gender-specific firm wage component that combines productivity pass-through with compensation for workplace risk (Proposition 2). In Section 4, we recover these components using two-way fixed effects regressions for wages (estimated separately by gender) and, in parallel, for injury incidence; the firm component from the injury regression provides a firm-level measure of workplace risk. We then relate the gender-specific firm wage components to a firm-level productivity proxy based on value added per worker and to the firm risk measure, and use these relationships to decompose the gender gap in firm pay policies into productivity and risk channels, each split into within-firm wage-setting and between-firm sorting components.

4. Empirical model

In this section, we describe how we estimate the quantities introduced in the theoretical framework.

4.1. Gender-specific AKM wage decomposition

The theoretical model described in Section 3.4 derives a two-way fixed effects model of wages. We bring the model to the data by estimating, separately by gender, the following AKM regression:

$$(9) \quad w_{it} = \alpha_i + \psi_{J(i,t)}^{G(i)} + X'_{it} \beta^{G(i)} + \varepsilon_{it}.$$

Here, w_{it} denotes log full-time-equivalent weekly earnings of worker i at time t ; α_i are worker fixed effects; $\psi_{J(i,t)}^{G(i)}$ are firm fixed effects, specific for each gender $G(i) \in \{M, F\}$;

¹¹Since $\delta_g = \beta_g c_1 + (1 - \beta_g) \phi_g$, changes in bargaining power affect the injury-risk coefficient through two opposing forces: higher β_g raises safety-cost pass-through ($\beta_g c_1$) but lowers the weight on compensating differentials ($(1 - \beta_g) \phi_g$). The net effect is $\partial \delta_g / \partial \beta_g = c_1 - \phi_g$, so cross-group comparisons of δ_g are not signed without further restrictions.

X_{it} is a vector containing a third-order polynomial in age and year dummies, multiplied by a gender-specific parameter $\beta^{G(i)}$; and ε_{it} is an idiosyncratic error term.¹²

As discussed in Abowd, Creedy, and Kramarz (2002), identification of firm effects relies on worker mobility across firms. Accordingly, we restrict the analysis to the largest connected component in both the male and female samples, defined as the set of firms linked through worker mobility. This component includes all workers who have ever been employed by one of the firms in the set and all firms that have ever employed at least one of these workers. The largest connected components include 98.1 percent of male and 96.4 percent of female person-year observations in the original sample.

For firm effects to be consistently estimated, firm-to-firm mobility must be conditionally random. This requires that the error term is not systematically correlated with workers' mobility decisions. Potential violations arise if idiosyncratic productivity shocks to firms induce workers to seek alternative jobs, or if worker-specific shocks trigger mobility across firms. A long-standing literature has shown that this assumption is plausible in the context of Italian matched employer–employee data (Casarico and Lattanzio 2024; Macis and Schivardi 2016; Devicienti, Fanfani, and Maida 2019).

Wage AKM validity. A second concern in AKM estimation is the well-known limited-mobility bias, which primarily affects the estimation of higher-order moments of the fixed-effects distribution. When worker mobility across firms is limited, the estimation errors in worker and firm effects are negatively correlated, leading to a downward bias in estimated covariances between worker and firm effects (Andrews, Gill, Schank et al. 2008). To address this issue, we recover firm effects using a split-sample strategy, following Palladino, Babet, and Godechot (Forthcoming). Specifically, we randomly split workers within each firm into two groups and estimate the AKM model separately on each subsample. The covariance between worker and firm effects across the two splits provides an unbiased estimate of higher-order moments.

Normalization. Firm effects are identified only up to an additive constant, and this constant generally differs across estimation samples. As a result, comparing male and female firm effects requires an explicit normalization to recover meaningful within-firm differences across genders. We follow a standard approach in the literature, grounded in the theoretical framework outlined above. Specifically, firm pay policy is normalized to zero in firms that generate no surplus, where the scope for rent sharing is minimal. Such low-surplus firms have been identified in the literature either through regression-based methods that isolate firms with low value added per worker, or by using the hospitality sector as a proxy for low value added. Card, Cardoso, and Kline (2016) show that these

¹²Following Card, Cardoso, Heining et al. (2018) and the evidence for Italy in Casarico and Lattanzio (2024), we normalize the age profile to be flat at age 50 and exclude the linear age term from the AKM specification.

normalization strategies are largely equivalent in the Portuguese context, while Casarico and Lattanzio (2024) provides similar evidence for Italy using data from the same source as in this study. For these reasons, and for its simplicity relative to value-added-based normalizations, we normalize firm effects to have zero mean in the hospitality sector (NACE section I).

4.2. AKM injury decomposition

We estimate an analogous two-way fixed effects model for injury probabilities. The resulting firm effects proxy workplace risk at the employer level, measuring the average injury probability after accounting for worker composition. To this end, we estimate the following AKM regression:

$$(10) \quad i_{it} = \theta_i + \lambda_{J(i,t)} + X'_{it}\eta + v_{it}.$$

The dependent variable i_{it} is an indicator for worker i experiencing a workplace injury at time t . θ_i and $\lambda_{J(i,t)}$ denote worker and firm fixed effects, respectively; X_{it} is defined as above; and v_{it} is an idiosyncratic error term. In this case, we do not estimate separate regressions by gender, as our objective is to recover an overall firm-level measure of injury risk, consistent with the theoretical framework.

Injury AKM validity. While the validity of the AKM wage decomposition is well established, much less evidence exists on the applicability of the AKM framework to injury probabilities. As in the wage case, consistent estimation of fixed effects requires worker mobility across firms to be conditionally exogenous, meaning that mobility decisions should not be systematically correlated with the error term. We assess the plausibility of this assumption using descriptive tests adapted from the canonical AKM wage literature. In particular, mobility should not be driven by idiosyncratic worker health shocks, firm-level technological changes that alter workplace risk, or persistent worker–firm match components.

To this end, we implement a mover design and examine injury dynamics for workers transitioning across quartiles of average coworker injury probabilities. Figure C7, panel A, reports average individual injury probabilities for workers moving from the bottom to the top quartile and vice versa. If mobility were induced by health shocks or by firm-level changes that make workplaces inherently more or less risky, we would expect to observe systematic pre-trends in injury probabilities prior to the move. Instead, the figure shows no such pre-trends: injury probabilities are flat before mobility. Post-move dynamics are also relatively flat, suggesting that workers do not, on average, sort into safer or riskier firms in response to changing injury risk. At the same time, transitions from low- to high-

risk coworker environments are associated with higher individual injury probabilities, and the reverse holds for moves in the opposite direction, as expected in the absence of strong match effects.

Consistent with this interpretation, panel B shows that moves within the same quartile display similarly flat dynamics. The symmetry of these patterns further suggests that worker–firm match effects play a limited role in explaining mobility decisions related to injury risk. Panels C and D reinforce this conclusion by comparing changes in injury probabilities for symmetric moves between the bottom and top quartiles. Panel C presents raw changes, while panel D reports adjusted changes after residualizing injury probabilities with respect to a quadratic polynomial in age, broad occupation-by-tenure dummies, and year fixed effects. In both cases, and especially for adjusted outcomes, the responses are symmetric, providing additional evidence against systematic sorting driven by match-specific considerations.

Finally, Figure C8 reports average AKM residuals by deciles of firm and worker effects. In most cells, residuals lie between -0.03 and 0.03, with larger values appearing only at the extreme deciles of worker effects. Overall, residuals are small across most of the distribution, particularly beyond the lower tail, supporting the suitability of the AKM framework for modeling individual injury probabilities.

Beyond consistency, limited-mobility bias may also affect the estimation of higher-order moments of the fixed-effects distributions in the injury setting. As in the wage case, this bias can distort variance and covariance components. Accordingly, we apply the same split-sample procedure to the injury regressions, which allows us to correct higher-order moments and obtain unbiased estimates of the variance components.

4.3. Decomposing the gender gap in firm wage policies

After estimating gender-specific firm wage policies and firm injury policies using the full worker-level data, we restrict the analysis to firms belonging to the representative sample in the dual connected set, defined as the intersection of the largest connected components for male and female employees. This restriction also implies focusing on dual-gender firms, namely firms employing at least one male and one female worker. Descriptive statistics for this sample are reported in column (5)–(6) of Table 1. The restriction to the dual connected set does not materially affect worker characteristics, which remain very close to those reported in column (3)–(4) for the full set of firms in the BdI sample.

We therefore collapse the data to the firm level by averaging outcomes across years and estimate the empirical counterpart of equation (7):

$$(11) \quad \widehat{\psi}_j^g = \gamma_g NS_j + \delta_g \widehat{\lambda}_j + \omega_j,$$

where the dependent variable $\widehat{\psi}_j^g$ denotes the gender-specific firm wage policy for firm j and gender g . NS_j is the firm's net surplus, which proxies productivity p_j in equation (7) of the model. Following Card, Cardoso, and Kline (2016), net surplus is defined as $NS_j = \max(S_j - \mathbb{E}[S_j | \text{low-surplus}], 0)$, that is, net surplus equals observed surplus when it exceeds the average surplus in low-surplus sectors, and zero otherwise. We measure S_j using log value added per worker (or, as a robustness check, log sales per worker), averaged across years at the firm level. To account for differences in firm size, regressions are weighted by the total number of person-year observations for each firm over the sample period 2005-2019.

The estimated coefficients and variables in equation (11) are then used to implement a Kitagawa–Oaxaca–Blinder decomposition of the gender gap in average firm effects (Kitagawa 1955; Oaxaca 1973; Blinder 1973). The gender gap in average firm wage policies can be expressed as

$$\mathbb{E}[\psi_j^M | g = M] - \mathbb{E}[\psi_j^F | g = F].$$

Substituting equation (11) into this expression and adding and subtracting appropriate terms yields the following decomposition:

$$(12) \quad \mathbb{E}[\psi_j^M | M] - \mathbb{E}[\psi_j^F | F] = \underbrace{(\gamma_M - \gamma_F) \mathbb{E}[NS_j | M]}_{\text{Productivity: within}} + \underbrace{\gamma_F (\mathbb{E}[NS_j | M] - \mathbb{E}[NS_j | F])}_{\text{Productivity: between}} \\ + \underbrace{(\delta_M - \delta_F) \mathbb{E}[\lambda_j | M]}_{\text{Injury: within}} + \underbrace{\delta_F (\mathbb{E}[\lambda_j | M] - \mathbb{E}[\lambda_j | F])}_{\text{Injury: between}}.$$

The four components have distinct economic interpretations.

The *within-firm productivity* component, $(\gamma_M - \gamma_F) \mathbb{E}[NS_j | M]$, captures the portion of the gap arising from men and women receiving different returns to firm productivity within the same firm.¹³

The *between-firm productivity* component, $\gamma_F (\mathbb{E}[NS_j | M] - \mathbb{E}[NS_j | F])$, reflects gender differences in sorting across firms with different productivity levels.

The *within-firm injury* component, $(\delta_M - \delta_F) \mathbb{E}[\lambda_j | M]$, measures the contribution of differential compensation for injury risk within the same firm.¹⁴

Finally, the *between-firm injury* component, $\delta_F (\mathbb{E}[\lambda_j | M] - \mathbb{E}[\lambda_j | F])$ captures the contribution of gender-specific sorting into firms with different injury risk levels.

¹³From equation (8) in the theoretical framework, $\gamma_M - \gamma_F = (r + \delta) \bar{\kappa}_P (\beta_M - \beta_F)$; this term is therefore positive when men have greater bargaining power than women ($\beta_M > \beta_F$).

¹⁴From equation (8), $\delta_M - \delta_F = (\beta_M - \beta_F) c_1 + (1 - \beta_M) \phi_M - (1 - \beta_F) \phi_F$. This term depends on gender differences in bargaining power, risk cost, and their interaction.

Sources of sorting. The “between” components in (12) require that men and women work at different firms (or, more generally, at different parts of the firm distribution). In the model, such sorting can arise endogenously from gender differences in risk cost ϕ_g . Consider a worker’s decision to accept a job offer from firm (p, λ) , where p denotes productivity and λ is the firm’s injury rate. The match is acceptable if and only if $S_g^H(p, \lambda) \geq 0$, which requires

$$(r + \delta)\kappa_p(\lambda) p \geq c(\lambda) + \phi_g \lambda + rW_{g,i}^0.$$

Workers with higher risk cost ϕ_g face a tighter acceptance constraint at high- λ firms: for a given productivity level p , there exists a threshold $\bar{\lambda}_g(p)$ above which the match is not individually rational. Accordingly, if $\phi_F > \phi_M$, this acceptance channel predicts that women are less likely to match with high-risk firms than men, implying $\mathbb{E}[\lambda_j | F] < \mathbb{E}[\lambda_j | M]$; conversely, if $\phi_F < \phi_M$, the same channel predicts the opposite ordering.¹⁵

Similarly, if outside options differ ($W_F^0 \neq W_M^0$), the acceptance constraint shifts for all firm types; for a given risk level, this changes the set of productivity levels p that satisfy the constraint, potentially generating productivity sorting as well. The decomposition (12) thus provides a framework for quantifying how much of the gender gap in firm effects is due to differential returns (within) versus differential sorting (between), separately for productivity and injury risk channels.

5. Results

This section reports the estimated wage and injury decompositions and the resulting gender-gap breakdown, with particular attention to how injury-risk premia vary across sectors and occupations.

5.1. Estimation of wage worker and firm effects

We estimate equation (9) on the largest connected components of male and female employees. Table 2 reports descriptive statistics and variance decompositions, which provide a benchmark for comparison with the existing literature. Panel A presents plug-in (biased) estimates, while Panel B reports split-sample bias-corrected measures of dispersion.¹⁶

The standard deviation of wages is larger for men than for women. For both genders, the dispersion of worker effects is roughly twice that of firm effects. The correlation between worker and firm effects is positive for both genders—0.18 for men and 0.08 for

¹⁵In our empirical implementation, the moment restrictions in the linearized setting imply $\phi_F < \phi_M$ throughout the admissible range of c_1 . We return to the interpretation of this restriction—and to the role of outside options and offer-side differences in shaping observed sorting—in Section 5 and Appendix B.5.

¹⁶As discussed above, the AKM model is estimated on all person-year observations to fully exploit worker mobility. All statistics are then reported for the subsample of firms in the representative sample, for which full employment information is observed and firm effects are more reliably estimated.

TABLE 2. Standard deviation and variance decomposition of wage AKM effects

	Men		Women	
	(1) Level	(2) Share (%)	(3) Level	(4) Share (%)
<i>A. Plug-in</i>				
SD Y	0.4428		0.4254	
SD worker effects	0.3105		0.2817	
SD firm effects	0.1784		0.1949	
Corr worker–firm effects	0.1763		0.0849	
Var Y	0.1961	100.0	0.1810	100.0
Var worker effects	0.0964	49.2	0.0794	43.9
Var firm effects	0.0318	16.2	0.0380	21.0
Cov worker–firm effects	0.0097	5.0	0.0046	2.6
<i>B. Split-sample correction</i>				
SD worker effects	0.2995		0.2583	
SD firm effects	0.1706		0.1789	
Corr worker–firm effects	0.2140		0.1924	
Var worker effects	0.0897	45.8	0.0667	36.9
Var firm effects	0.0291	14.9	0.0320	17.7
Cov worker–firm effects	0.0109	5.6	0.0089	4.9

Note: The table reports variance decompositions from the AKM wage model estimated separately for men and women. Panel A presents plug-in estimates, while Panel B reports split-sample bias-corrected estimates that address limited-mobility bias in the estimation of worker and firm effects. The model is estimated on the full sample of workers and firms, separately by gender, but statistics are computed on workers and firms belonging to the representative sample. For each gender, the table reports the standard deviation (SD) of wages, worker effects, and firm effects, as well as the correlation between worker and firm effects. The lower part of each panel reports the corresponding variance components. “Var Y” denotes the variance of log wages. “Var worker effects” and “Var firm effects” report the variance attributable to worker and firm fixed effects, respectively, while “Cov worker–firm effects” reports their covariance. “Level” reports the variance contribution in log-point units, and “Share” reports the percentage contribution to the total variance of wages. Shares may not sum exactly to 100 due to rounding.

women—but these correlations are downward biased. After split-sample correction, the worker–firm correlation increases to 0.21 for men and 0.19 for women, suggesting that limited-mobility bias is more severe in the female sample, consistent with lower job mobility among women (Casarico and Lattanzio 2024).

After bias correction, worker effects account for about 46 percent of wage variance for men and 37 percent for women. Firm effects and the covariance between firm and worker effects together explain roughly 20.5 percent of wage variance for men and 22.6 percent for women. These magnitudes are in line with existing evidence for Italy (Casarico and Lattanzio 2024; Palladino, Bertheau, Hijzen et al. 2025).

5.2. Estimation of injury worker and firm effects

We now turn to the estimation of equation (10), where injury probability is the outcome. Table 3 reports dispersion measures and variance decompositions, with plug-in estimates in Panel A and split-sample bias-corrected estimates in Panel B. The plug-in results suggest that the standard deviation of worker effects is about 2.5 times larger than that of firm effects and that the correlation between worker and firm effects is large and negative.

Split-sample correction mitigates limited-mobility bias and substantially reduces the dispersion of worker effects, while the magnitude of the worker–firm correlation also declines. In terms of variance components, Panel B shows that person effects account for about 5 percent of the variance in injury probability, while firm effects explain roughly 1 percent—both markedly smaller than implied by the plug-in estimates.¹⁷ The covariance term contributes only marginally, indicating little systematic sorting of more injury-prone workers into riskier firms.

Figure C9 shows the joint distribution of worker and firm injury effects.

5.3. The determinants of the gender gap in firm wage policies

Regression results. We next estimate firm-level rent-sharing models. We collapse the data to the firm level and estimate equation (11), regressing gender-specific firm wage effects on net surplus (excess log value added per worker) and firm-level injury effects. Regressions are weighted by the total number of person–year observations associated with firm j . Table 4 reports the coefficient estimates. Panel A presents OLS results, while Panel B reports 2SLS estimates in which injury effects are instrumented with their split-sample counterparts; in this case, the number of observations is smaller because not all firms appear in both split samples. The 2SLS estimates address potential attenuation bias arising from measurement error. The firm-level effects used in the analysis are generated regressors, whose estimated variance reflects both true cross-firm heterogeneity and sampling noise (Andrews, Gill, Schank et al. 2008; Kline, Saggio, and Solvsten 2020). To assess the quantitative importance of this measurement error, we rely on split-sample estimates of the firm effects, constructed on two non-overlapping subsamples, and use one estimate to instrument the other (see, e.g., Goldschmidt and Schmieder 2017; Gerard, Lagos, Severnini et al. 2021; Drenik, Jäger, Plotkin et al. 2023).

Columns (1)–(2) report estimates with male firm effects as the outcome, while columns (3)–(4) focus on female firm effects. Columns (2) and (4) additionally control for a quadratic polynomial in firm size,¹⁸ five macro-area dummies, and 38 macro-sector dummies corre-

¹⁷The implied reliability of the firm injury effect—measured as the ratio of bias-corrected to plug-in variance—is $0.00018/0.00041 \approx 0.44$. This level of measurement error motivates the 2SLS strategy in Table 4, where the split-sample instrument yields first-stage F -statistics of 229 and 389 (without and with controls), well above conventional weak-instrument thresholds.

¹⁸Firm size is measured as average employment over time.

TABLE 3. Standard deviation and variance decomposition of injury AKM effects

	(1) Level	(2) Share (%)
<i>A. Plug-in</i>		
SD Y	0.13116	
SD worker effects	0.05169	
SD firm effects	0.02022	
Corr worker–firm effects	-0.20325	
Var Y	0.01720	100.0%
Var worker effects	0.00267	15.5%
Var firm effects	0.00041	2.4%
Cov worker–firm effects	-0.00021	-1.2%
<i>B. Split-sample correction</i>		
SD worker effects	0.02921	
SD firm effects	0.01328	
Corr worker–firm effects	-0.04185	
Var worker effects	0.00085	5.0%
Var firm effects	0.00018	1.0%
Cov worker–firm effects	-0.00002	-0.1%

Note: The table reports variance decompositions from the AKM model for injury probability. Panel A presents plug-in estimates, while Panel B reports split-sample bias-corrected estimates that address limited-mobility bias in the estimation of worker and firm effects. The model is estimated on the full sample of workers and firms, but statistics are computed on workers and firms belonging to the representative sample. For each panel, the table reports the standard deviation (SD) of injury probability, worker effects, and firm effects, as well as the correlation between worker and firm effects. The lower part of each panel reports the corresponding variance components. “Var Y” denotes the variance of the injury outcome. “Var worker effects” and “Var firm effects” report the variance attributable to worker and firm fixed effects, respectively, while “Cov worker–firm effects” reports their covariance. “Level” reports the contribution of each component in probability units, and “Share” reports its percentage contribution to the total variance of injury probability. Shares may not sum exactly to 100 due to rounding.

sponding to the OECD groupings of NACE Rev. 2 sectors. Figure C10 provides a graphical illustration of these relationships.

We find a positive and statistically significant association between net surplus and both male and female firm wage effects. Coefficients are larger for men, implying a relative bargaining power of 0.94 (0.124/0.132) based on the controlled 2SLS estimates. In other words, women receive about 94 percent of the wage response that men obtain with respect to firm surplus, in line with the cross-country evidence in Palladino, Bertheau, Hijzen et al. (2025) and consistent with Caldwell, Haegele, and Heining (2026), who document that firms engaging in individual wage negotiation have a 3 percentage point higher gender wage gap.

We also find a positive association between firm-level injury effects and firm wage effects, again stronger for men than for women. In the 2SLS estimates, a one–percentage–point increase in the injury firm effect (i.e., a 0.01 increase in the annual average firm-level

TABLE 4. Firm-level rent- and risk-sharing models

Dependent variable:	Male firm effect		Female firm effect	
	(1)	(2)	(3)	(4)
<i>A. OLS</i>				
Excess log VA/L	0.160*** (0.012)	0.133*** (0.005)	0.148*** (0.013)	0.124*** (0.005)
Injury firm effect	0.843*** (0.172)	0.654*** (0.096)	0.743*** (0.183)	0.514*** (0.118)
R-sq.	0.322	0.489	0.238	0.380
Observations	76,165	76,165	76,165	76,165
<i>B. 2SLS</i>				
Excess log VA/L	0.158*** (0.012)	0.132*** (0.005)	0.146*** (0.013)	0.124*** (0.006)
Injury firm effect	1.676*** (0.338)	1.416*** (0.198)	1.557*** (0.349)	1.220*** (0.251)
F-stat	228.9	388.6	228.9	388.6
R-sq.	0.292	0.468	0.215	0.366
Observations	75,061	75,061	75,061	75,061
Controls	No	Yes	No	Yes

Note: The table reports cross-sectional regression results for firms included in the dual connected sample. Panel A shows OLS estimates, panel B shows 2SLS estimates in which the firm injury effect is instrumented with a split-sample strategy. *Excess log VA/L* is the mean log value added per worker in deviation from the average in the low surplus sector (food and accommodation). The mean is computed over each firm-specific timespan. *Injury firm effect* is the firm fixed effect from the AKM model estimated on worker-level injury incidence. Controls include a second-order polynomial in firm size, 5 macro-area dummies, and 38 broad sector dummies. All regressions are weighted by the firm-level number of person-year observations. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

conditional injury probability) is associated with a 1.4 percent increase in male firm wage effects, compared with a 1.2 percent increase for female firm wage effects (Table 4, columns 2 and 4). For a one-standard deviation increase in injury firm effects (0.0176 in the firm-level sample), the implied effects correspond to increases of 2.5 percent for male firm effects and 2.1 percent for female firm effects, respectively. This implies that women receive about 86 percent (1.220/1.416) of the wage response to workplace risk embedded in firm pay policies. In contrast to the null gender gap for fatality risk in Lavetti and Schmutte (2023) once accounting for firm effects and measurement, we find a meaningful gender difference in the firm-level wage-risk gradient for non-fatal injuries.

Table D1 shows that using excess log sales per worker as an alternative measure of net surplus yields very similar relative estimates across genders.

TABLE 5. KOB decomposition of firm wage effects gap

	(1)	(2)	(3)	(4)
Coefficients estimated with:	OLS		2SLS	
	Level	Share	Level	Share
Wage gap	0.220		0.220	
Wage firm effects gap	0.072	100	0.072	100
Productivity gap	0.045	61.8	0.044	61.1
Within firm	0.010	13.3	0.010	13.1
Between firm	0.035	48.6	0.035	48.0
Risk gap	0.003	3.5	0.006	8.1
Within firm	0.001	1.6	0.003	4.2
Between firm	0.001	1.9	0.003	3.9

Note: The table reports a Kitagawa–Oaxaca–Blinder decomposition of the gender gap in firm pay policies for firms in the dual connected sample, as outlined in equation (12). The overall wage gap is measured in log points. Columns (1)–(2) use OLS coefficients, while columns (3)–(4) use 2SLS coefficients from the firm-level rent- and risk-sharing regressions. In the 2SLS specification, firm-level injury effects are instrumented with their split-sample counterparts. “Wage firm effects gap” denotes the portion of the overall wage gap attributable to differences in firm wage effects between men and women and is normalized to 100 percent in the share columns. The productivity gap captures differences related to net surplus (excess log value added per worker) and is decomposed into within-firm (differences in rent sharing within firms) and between-firm (sorting across firms with different productivity levels) components. The risk gap captures differences related to firm-level injury effects and is similarly decomposed into within- and between-firm components. “Level” reports the contribution of each component in log points, while “Share” reports its percentage contribution to the firm wage effects gap. Shares may not sum exactly to 100 due to rounding.

5.4. Decomposition

After estimating the rent- and risk-sharing models, we use the resulting coefficients to decompose the gender gap in firm pay policies, as formalized in equation (12). Table 5 reports the Kitagawa–Oaxaca–Blinder decomposition for all firms in the dual connected sample, using OLS and 2SLS estimates from Table 4 in columns (1)–(2) and (3)–(4), respectively. The overall firm-level wage gap amounts to 22 log points, of which about one third is explained by gender differences in firm wage effects, consistent with the evidence in Casarico and Lattanzio (2024).

Most of the gender gap in firm pay policy is driven by productivity-related channels. Sorting of women into lower-productivity firms accounts for 48 percent of the gap, while differences in rent sharing within firms explain an additional 13 percent. Differences in firm-level injury effects account for 8 percent of the wage premium gap, evenly split between between-firm and within-firm components. The between-firm component reflects differential sorting of men and women across firms with different injury risk profiles, whereas the within-firm component captures differences in the assignment of men and women to more or less risky tasks and occupations within the same firm.

Structural interpretation. The finding that $\delta_F/\delta_M = 0.86$ has a direct structural interpretation in the model. Importantly, this exercise speaks to *pricing* (how wages load on firm attributes) and therefore informs the *within*-firm components in (12); sorting affects the *between* components only through gender differences in the conditional means $\mathbb{E}[NS_j | g]$ and $\mathbb{E}[\lambda_j | g]$.

From equation (8), the injury-risk coefficient is $\delta_g = \beta_g c_1 + (1 - \beta_g)\phi_g$, where ϕ_g captures the *effective* marginal risk cost of firm-level injury risk. Combining $\gamma_F/\gamma_M = 0.94$ (which implies $\beta_F/\beta_M = 0.94$) with $\delta_F/\delta_M = 0.86$ therefore restricts the set of $(c_1, \phi_F/\phi_M)$ pairs consistent with the estimates.¹⁹

Figure 2 visualizes this restriction. Panel A maps the estimated coefficient ratios into an *identification locus* linking the local safety-cost slope c_1 and the implied ratio ϕ_F/ϕ_M . Rather than identifying a single value of ϕ_F/ϕ_M , the estimates pin down a set of admissible combinations.

As a simple benchmark, in the limiting case $c_1 \rightarrow 0$ the restriction yields the closed-form bound

$$\frac{\phi_F}{\phi_M} \leq \frac{\delta_F}{\delta_M} \cdot \frac{1 - \beta_M}{1 - \beta_F},$$


which under our baseline calibration (Appendix B) implies $\phi_F/\phi_M \lesssim 0.81$.²⁰ Panel B clarifies magnitudes: empirically relevant injury rates vary by only a few percentage points, so the implied change in ψ from moving λ from 1 to 5 percent is mechanically smaller than the change induced by typical productivity dispersion, even when δ_g is non-negligible.

This bound may appear at odds with survey evidence that women tend to report higher risk aversion than men (Croson and Gneezy 2009). A key distinction is that ϕ_g reflects the worker's effective risk cost at a given firm—a composite of preferences and within-firm exposure—and therefore can differ from stated risk preferences. Consistent with an interpretation based on differential exposure within firms, the within-firm component of the injury channel becomes negligible once conditioning on broad occupation (Table 7), suggesting that women at high- λ firms may be assigned to safer tasks than men. Finally, our firm-level injury measure captures severe injuries generating at least one week of absence; if injury severity or reporting differs by gender, the estimated δ_g coefficients may not map one-to-one into compensating differentials for overall injury risk exposure.

¹⁹In the model, ϕ_g also shapes the acceptance cutoff in risk, hence the direction of sorting predicted by a pure acceptance channel. When $\phi_F/\phi_M < 1$, that channel alone would predict women to be (weakly) more willing to accept high- λ firms, so any empirical tendency for women to be concentrated in lower- λ firms must reflect differences in outside options and/or in the offer distribution across firm types, or the fact that ϕ_g embeds gender differences in effective exposure within firms.

²⁰This bound is obtained when the safety-cost channel is shut down ($c_1 \rightarrow 0$); for $c_1 > 0$ the implied ϕ_F/ϕ_M is typically smaller. Appendix B reports implied values for a range of c_1 .

FIGURE 2. Model implications for identification and magnitudes



Note: Left panel: implied ϕ_F/ϕ_M from $\delta_g = \beta_g c_1 + (1 - \beta_g) \phi_g$ given $\delta_F/\delta_M = 0.86$ and $\beta_F/\beta_M = 0.94$. Right panel: implied changes in firm effects from productivity dispersion ($p_{90} - p_{10}$) and from shifting λ from 1 to 5 percent, with the injury effect decomposed into safety-cost pass-through and risk compensation.

5.5. Heterogeneity

Heterogeneity by sector. The contribution of injury-related channels to the gender gap in firm pay policies is likely to vary across sectors for two reasons. First, workplace injury risk differs substantially by sector (Figure 1E). Second, men are disproportionately employed in manufacturing and construction relative to women (Table 1). To account for these differences, we compute the decomposition in equation (12) separately for manufacturing and services.²¹ Table 6 reports the results, with Panel A focusing on manufacturing and Panel B on services.

In both sectors, most of the gap in firm pay policies is explained by gender differences in firm productivity, largely driven by between-firm components. The within-firm component is more pronounced in services, consistent with women in services being, on average, less educated and more negatively selected relative to their counterparts in manufacturing, leading to larger within-firm differences in wage setting.²² By contrast, the contribution of injury premia to the gap in firm pay policies is entirely driven by manufacturing: it accounts for 6.3 percent of the gap using OLS estimates and rises to 17.4 percent under 2SLS. Most of this contribution arises from within-firm components, pointing to gender differences in the allocation to riskier tasks within firms.

In services, injury premia make a zero or negative contribution to the gap in firm

²¹We define “manufacturing” as NACE sections A–F, including industry and construction, and “services” as sections G–U.

²²The weekly wage gap between women in services and manufacturing is 17 percent, compared with 13 percent among men. The gap rises to 42 percent when annual earnings are used as the measure of pay, compared with 24 percent for men.

TABLE 6. KOB decomposition of firm wage effects gap by sector

	(1)	(2)	(3)	(4)
Coefficients estimated with:	OLS		2SLS	
	Level	Share	Level	Share
<i>A. Manufacturing</i>				
Wage gap	0.185		0.185	
Wage firm effects gap	0.056	100	0.056	100
Productivity gap	0.040	72.3	0.040	72.1
Within firm	0.004	7.4	0.004	7.3
Between firm	0.036	64.9	0.036	64.8
Risk gap	0.003	6.3	0.010	17.4
Within firm	0.002	3.6	0.007	12.1
Between firm	0.002	2.7	0.003	5.3
<i>B. Services</i>				
Wage gap	0.198		0.198	
Wage firm effects gap	0.060	100	0.060	100
Productivity gap	0.047	79.3	0.047	78.4
Within firm	0.015	24.3	0.015	24.6
Between firm	0.033	55.1	0.032	53.8
Risk gap	-0.000	-0.2	-0.003	-4.8
Within firm	0.000	0.2	-0.002	-3.8
Between firm	-0.000	-0.5	-0.001	-1.0

Note: The table reports a Kitagawa–Oaxaca–Blinder decomposition of the gender gap in firm pay policies, as defined in equation (12), separately by broad sector. Panel A presents results for manufacturing firms, while Panel B reports results for service-sector firms; both panels are based on firms in the dual connected sample. The overall wage gap is measured in log points. Columns (1)–(2) use OLS coefficients, and columns (3)–(4) use 2SLS coefficients from the firm-level rent- and risk-sharing regressions. In the 2SLS specification, firm-level injury effects are instrumented with their split-sample counterparts. “Wage firm effects gap” denotes the portion of the overall wage gap attributable to gender differences in firm wage effects and is normalized to 100 percent in the share columns. The productivity gap captures differences related to net surplus (excess log value added per worker) and is decomposed into within-firm (differences in rent sharing within firms) and between-firm (sorting across firms with different productivity levels) components. The risk gap captures differences related to firm-level injury effects and is similarly decomposed into within- and between-firm components. “Level” reports the contribution of each component in log points, while “Share” reports its percentage contribution to the firm wage effects gap. Shares may not sum exactly to 100 due to rounding.

pay policies. In the 2SLS estimates, this negative contribution is driven by within-firm components, suggesting that women in services may be relatively more exposed to risky tasks—such as in health, retail, or education—despite the overall lower level of injury risk in these activities. Finally, the between-firm component of the injury gap, which plays a role in the full sample, is primarily driven by manufacturing, while it is close to zero in services, consistent with the idea that sorting across firms by injury risk largely reflects

differences in sectoral allocation rather than within-sector sorting.

Firms vs. occupations. Men and women are unevenly distributed across occupations (Goldin 2014), raising the concern that the effects we document may simply reflect sorting into occupations with different injury risks rather than differences across firms. In particular, the between-firm component associated with injury-related premia could capture differential allocation of men and women across occupations rather than genuine firm-level differences. To address this concern, we follow an approach similar in spirit to Card, Cardoso, and Kline (2016) and estimate AKM models analogous to equations (9) and (10), replacing firm effects with firm–occupation effects.

Specifically, we define two broad occupational groups—blue-collar (including apprentices) and white-collar (including middle managers and executives)—and treat employment in different broad occupations within the same firm as distinct units in the AKM framework.²³ We then estimate gender-specific firm–occupation wage effects and overall firm–occupation injury effects, and decompose the gender gap in firm–occupation wage effects as in equation (12), separately for blue- and white-collar workers. Results are reported in Table 7, Panels A and B, respectively.

Consistent with the sectoral analysis, injury premia play a role only in blue-collar occupations, which exhibit higher injury risk (Figure 1D). In contrast, for white-collar occupations the contribution of injury premia is zero or negative, mirroring the patterns observed in services. Importantly, within blue-collar occupations the entire contribution of injury premia operates through the between-firm component, which is positive and economically meaningful. This indicates that the results are not driven by men and women sorting into different occupations in the labor market, but by genuine differences in injury-related premia across firms: even after accounting for occupational sorting, women tend to work at employers characterized by lower injury-risk premia. The within-firm component is zero or negative, reinforcing the interpretation that the positive within-firm component documented in the aggregate analysis reflects differential task allocation within firms.

6. Magnitudes and incidence

This section provides a transparent benchmark that translates the estimated cross-firm wage–risk gradient into currency units, and clarifies how these wage-based magnitudes relate to welfare in the model.

²³Card, Cardoso, and Kline (2016) classify occupations as “mostly male” or “mostly female” based on female shares across detailed ISCO occupations. Because we lack detailed occupational information, we focus on a broad blue- versus white-collar distinction.

TABLE 7. KOB decomposition of firm-occupation wage effects gap by broad occupation

	(1)	(2)	(3)	(4)
Coefficients estimated with:	OLS		2SLS	
	Level	Share	Level	Share
<i>A. Blue-collar</i>				
Wage gap	0.255		0.255	
Wage firm effects gap	0.115	100	0.115	100
Productivity gap	0.048	42.1	0.048	41.5
Within firm	0.004	3.1	0.004	3.2
Between firm	0.045	39.0	0.044	38.3
Risk gap	0.004	3.9	0.006	5.6
Within firm	0.000	0.4	-0.000	-0.3
Between firm	0.004	3.5	0.007	5.9
<i>B. White-collar</i>				
Wage gap	0.314		0.314	
Wage firm effects gap	0.085	100	0.085	100
Productivity gap	0.042	49.1	0.041	48.9
Within firm	0.015	17.8	0.015	17.6
Between firm	0.026	31.3	0.026	31.3
Risk gap	-0.001	-1.4	-0.003	-3.9
Within firm	-0.001	-1.2	-0.003	-3.4
Between firm	-0.000	-0.2	-0.000	-0.5

Note: The table reports a Kitagawa–Oaxaca–Blinder decomposition of the gender gap in firm pay policies, as defined in equation (12), separately by broad occupational group. Panel A presents results for blue-collar occupations, while Panel B reports results for white-collar occupations. The decomposition is based on AKM models estimated at the firm–occupation level and is restricted to the dual connected sample. The overall wage gap is measured in log points. Columns (1)–(2) use OLS coefficients, and columns (3)–(4) use 2SLS coefficients from the firm–occupation–level rent- and risk-sharing regressions. In the 2SLS specification, firm–occupation injury effects are instrumented with their split-sample counterparts. “Wage firm effects gap” denotes the portion of the overall wage gap attributable to gender differences in firm–occupation wage effects and is normalized to 100 percent in the share columns. The productivity gap captures differences related to net surplus (excess log value added per worker) and is decomposed into within-firm (differences in rent sharing across occupations within firms) and between-firm (sorting across firms with different productivity levels) components. The risk gap captures differences related to firm–occupation injury effects and is similarly decomposed into within- and between-firm components. “Level” reports the contribution of each component in log points, while “Share” reports its percentage contribution to the firm–occupation wage effects gap. Shares may not sum exactly to 100 due to rounding.

6.1. Wage-equivalent valuation of workplace safety

Let λ_j denote the firm injury effect (in annual probability units) and let δ_g denote the estimated coefficient on λ_j in the firm-level wage regression for gender g (Table 4). Since the wage outcome is in logs, a small change in risk $d\lambda$ implies a change in the firm wage

effect of approximately $d\psi_g \approx \delta_g d\lambda$. Converting this log change into euros, the implied change in annual earnings is

$$\Delta y_g(d\lambda) = \bar{y}_g (\exp(\delta_g d\lambda) - 1) \approx \bar{y}_g \cdot \delta_g \cdot d\lambda,$$

where $\bar{y}_g$ is an annual earnings level and the approximation uses a first-order expansion. We report magnitudes in absolute value for a risk reduction, i.e. using $|d\lambda|$.

We operationalize $\bar{y}_g$ using the DC-BDI sample means in Table 1, which report average annual earnings of 30,886.36 euros for men and 20,320.27 euros for women.

Combining these values with the controlled 2SLS estimates in Table 4 ($\delta_M = 1.416$, $\delta_F = 1.220$) yields three interpretable benchmarks:

- *1pp risk reduction (per worker-year)*. A one-percentage-point reduction in the annual probability of an observed injury corresponds to a wage-equivalent value of about 437 euros per male worker-year ($0.01 \times 30,886.36 \times 1.416$) and 248 euros per female worker-year ($0.01 \times 20,320.27 \times 1.220$).
- *Per expected injury avoided*. Expressed per expected injury avoided (100 percent of risk reduction), the implied wage-equivalent values are $\bar{y}_g \delta_g$, i.e. about 43,735 euros for men ($30,886.36 \times 1.416$) and 24,791 euros for women ($20,320.27 \times 1.220$).
- *One-standard-deviation improvement in firm safety*. A one-standard-deviation decrease in the firm injury effect is 0.01328 (Table 3, split-sample correction). This corresponds to wage-equivalent values of about 581 euros per male worker-year ($30,886.36 \times 1.416 \times 0.01328$) and 329 euros per female worker-year ($20,320.27 \times 1.220 \times 0.01328$). In log terms, this is $\delta_g \times 0.01328$, i.e. about 0.0188 and 0.0162 log points for men and women (roughly 1.9 and 1.6 percent of annual earnings), respectively.

These calculations are intended as a transparent benchmark that converts the estimated wage-risk gradient into currency units for the specific injury concept observed in the administrative data.

Relation to the gender gap in firm pay policies. Table 5 implies that injury-related channels account for 0.006 log points of the 0.072 log-point gap in firm pay policies in the dual connected sample (8.1 percent). In euro terms, 0.006 log points correspond to about 0.6 percent of annual earnings, i.e. roughly 120–185 euros per year for a worker earning 20–31k euros annually. The injury channel is concentrated in manufacturing (0.010 log points; 17.4 percent) and is near zero (and slightly negative) in services.

Equalizing the pricing of workplace risk: an upper bound on gap closure. Interpreted as a returns-equalization counterfactual, the 0.006 log-point injury component is the amount by which the firm-pay-policy gap would mechanically shrink if women faced the same wage-risk gradient as men, holding sorting across firms fixed.

Placing this channel in the context of the *overall* gender wage gap, 0.006 log points represent about 2.7 percent of the total gap in log weekly wages (0.006/0.220). This should be interpreted as an upper bound on the *mechanical* wage-gap closure attributable to differential risk premia embedded in firm pay policies; it is not a welfare statement, since reductions in risk can raise worker welfare even if compensating wages fall.

In monetary terms, 0.006 log points correspond to about a 0.6 percent earnings change. This is roughly $\bar{y}_F(\exp(0.006) - 1) \approx 122$ euros per woman-year (and $\bar{y}_M(\exp(0.006) - 1) \approx 186$ euros per man-year).

Finally, the mechanical closure is larger in manufacturing, where Table 6 implies that injury-related channels account for 0.010 log points of an overall wage gap of 0.185 log points, i.e. about 5.4 percent of the manufacturing wage gap (0.010/0.185).

6.2. Incidence and interpretation in the model

The wage-equivalent magnitudes above should be interpreted as risk prices implied by the estimated cross-firm wage-risk gradient, not as a direct estimate of risk cost. In the model, the reduced-form injury coefficient satisfies Equation (8) and therefore, δ_g generally combines (i) pass-through of safety-related costs into wages through bargaining and (ii) wage compensation for risk cost.

An implication is that a policy or shock that reduces injury risk may reduce the compensating wage premium (wages may fall), while worker welfare can nonetheless increase because risk cost falls. Without additional assumptions (or separate evidence) on safety costs and bargaining parameters, the wage-risk gradient alone does not pin down the welfare gain from safety improvements. Our structural-interpretation analysis (Figure 2) maps the estimated coefficient ratios into an identification locus for $(c_L, \phi_F/\phi_M)$, clarifying which combinations of safety-cost slopes and relative risk cost are consistent with the reduced-form estimates.

7. Conclusion

This paper examines how workplace injury risk contributes to gender differences in firm pay policies. We develop a search model that microfound an AKM wage decomposition with distinct productivity and risk components, and bring the implied reduced forms to Italian administrative data with individual-level injury records.

We find that women receive approximately 94 percent of the wage response to firm productivity and 86 percent of the wage response to workplace injury risk that men obtain at comparable firms. These estimates summarize differences in how firm attributes map into wages by gender; they do not represent unconditional wage gaps, nor a direct welfare valuation of risk.

Injury-related channels account for about 8 percent of the gender gap in firm wage effects overall, rising to 17 percent in manufacturing. A key result is that once we condition on broad occupation, the within-firm component of the injury channel becomes negligible, while the between-firm component remains. This pattern indicates that the injury channel primarily reflects differences in effective exposure driven by within-firm task allocation and sorting across firms with different risk and pay policies, rather than differential pricing of the same risk within firms.

Some limitations apply. Our injury measure captures only accidents generating at least seven days of absence, potentially missing less severe events. The model abstracts from on-the-job search and endogenous safety investment, both of which could enrich the analysis. Finally, while we document gender differences in the wage–risk gradient at the firm level, our framework does not fully disentangle the roles of bargaining power, preferences, and gender-specific exposure to risk at the task level. The occupation-conditioning results suggest that gender differences in effective exposure are central.

Future work could incorporate occupational choice and within-firm task assignment explicitly, and examine how institutional features—such as collective bargaining or safety regulations—mediate the relationship between injury risk, exposure, and wages.

References

- Abowd, John M., Robert H. Creedy, and Francis Kramarz. 2002. “Computing Person and Firm Effects Using Linked Longitudinal Employer–Employee Data.” Tech. Rep. 2002-06, Center for Economic Studies, U.S. Census Bureau. <https://ideas.repec.org/p/cen/tpaper/2002-06.html>.
- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. “High wage workers and high wage firms.” *Econometrica* 67 (2): 251–333.
- Andrews, Martyn J., Len Gill, Thorsten Schank, and Richard Upward. 2008. “High wage workers and low wage firms: negative assortative matching or limited mobility bias?” *Journal of the Royal Statistical Society Series A: Statistics in Society* 171 (3): 673–697.
- Blau, Francine D. and Lawrence M. Kahn. 2017. “The Gender Wage Gap: Extent, Trends, and Explanations.” *Journal of Economic Literature* 55 (3): 789–865. 10.1257/jel.20160995.
- Blinder, Alan S. 1973. “Wage Discrimination: Reduced Form and Structural Estimates.” *Journal of Human Resources* 8 (4): 436–455.
- Bruns, Benjamin. 2019. “Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap.” *American Economic Journal: Applied Economics* 11 (2): 74–113. 10.1257/app.20160664.
- Caldwell, Sydney, Ingrid Haefele, and Jörg Heining. 2026. “Bargaining and Inequality in the Labor Market.” *The Quarterly Journal of Economics* 10.1093/qje/qjaf049.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics* 36 (S1): S13–S70. 10.1086/694153.

- Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *The Quarterly Journal of Economics* 131 (2): 633–686. 10.1093/qje/qjw004.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *The Quarterly journal of economics* 128 (3): 967–1015.
- Casarico, Alessandra and Salvatore Lattanzio. 2024. "What firms do: Gender inequality in linked employer-employee data." *Journal of Labor Economics* 42 (2): 325–355.
- Casarico, Alessandra and Salvatore Lattanzio. 2025. "Exploring Gender Inequality: Firm Contribution and Policy Effects." *LSE Public Policy Review* 3 (4): 6.
- Croson, Rachel and Uri Gneezy. 2009. "Gender differences in preferences." *Journal of Economic literature* 47 (2): 448–474.
- Devicienti, Francesco, Bernardo Fanfani, and Agata Maida. 2019. "Collective bargaining and the evolution of wage inequality in Italy." *British Journal of Industrial Relations* 57 (2): 377–407.
- Drenik, Andres, Simon Jäger, Pascuel Plotkin, and Benjamin Schoefer. 2023. "Paying Outsourced Labor: Direct Evidence from Linked Temp Agency–Worker–Client Data." *The Review of Economics and Statistics* 105 (1): 206–216.
- Garen, John. 1988. "Compensating wage differentials and the endogeneity of job riskiness." *the Review of Economics and Statistics* pp. 9–16.
- Gerard, Francois, Lorenzo Lagos, Edson Severnini, and David Card. 2021. "Assortative Matching or Exclusionary Hiring? The Impact of Employment and Pay Policies on Racial Wage Differences in Brazil." *American Economic Review* 111 (10): 3418–3457.
- Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091–1119.
- Goldschmidt, Deborah and Johannes F. Schmieder. 2017. "The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure." *Quarterly Journal of Economics* 132 (3): 1165–1217. 10.1093/qje/qjx008.
- Guardado, José R and Nicolas R Ziebarth. 2019. "Worker investments in safety, workplace accidents, and compensating wage differentials." *International Economic Review* 60 (1): 133–155.
- Hersch, Joni. 1998. "Compensating differentials for gender-specific job injury risks." *The American Economic Review* 88 (3): 598–607.
- Kerndler, Martin. 2023. "Occupational safety in a frictional labor market." *Labour Economics* 83: 102388.
- Kitagawa, Evelyn M. 1955. "Components of a Difference Between Two Rates." *Journal of the American Statistical Association* 50 (272): 1168–1194.
- Kline, Patrick, Raffaele Saggio, and Mikkel Solvsten. 2020. "Leave-Out Estimation of Variance Components." *Econometrica* 88 (5): 1859–1898. 10.3982/ECTA16410.
- Lavetti, Kurt. 2023. "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications." *Journal of Economic Perspectives* 37 (3): 189–212. 10.1257/jep.37.3.189.
- Lavetti, Kurt and Ian M Schmutte. 2023. "Gender differences in sorting on wages and risk." *Journal of Econometrics* 233 (2): 507–523.
- Lavetti, Kurt and Ian M Schmutte. 2025. "Estimating Compensating Wage Differentials with Endogenous Job Mobility." *Mimeo*.
- Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. 2021. "Gender Differences in

- Job Search: Trading off Commute against Wage.” *The Quarterly Journal of Economics* 136 (1): 381–426. 10.1093/qje/qjaa033.
- Macis, Mario and Fabiano Schivardi. 2016. “Exports and wages: Rent sharing, workforce composition, or returns to skills?” *Journal of Labor Economics* 34 (4): 945–978.
- Mas, Alexandre and Amanda Pallais. 2017. “Valuing Alternative Work Arrangements.” *American Economic Review* 107 (12): 3722–3759. 10.1257/aer.20161500.
- Morchio, Iacopo and Christian Moser. 2024. “The gender pay gap: Micro sources and macro consequences.” Tech. rep., National Bureau of Economic Research.
- Oaxaca, Ronald. 1973. “Male–Female Wage Differentials in Urban Labor Markets.” *International Economic Review* 14 (3): 693–709.
- Olivetti, Claudia, Jessica Pan, and Barbara Petrongolo. 2024. “The Evolution of Gender in the Labor Market.” NBER Working Paper 33153, National Bureau of Economic Research. 10.3386/w33153.
- Palladino, Marco G., Damien Babet, and Olivier Godechot. Forthcoming. “In the Land of AKM: Explaining the Dynamics of Wage Inequality in France.” *Journal of Labor Economics*.
- Palladino, Marco G., Antoine Bertheau, Alexander Hijzen, Astrid Kunze, Cesar Barreto, Dogan Gülümser, Marta Lachowska, Anne Sophie Lassen, Salvatore Lattanzio, Benjamin Lochner, et al. 2025. “Firms and the Gender Wage Gap: A Comparison of Eleven Countries.” Working Paper 181, VAT Institute for Economic Research. <https://ideas.repec.org/p/fer/wpaper/181.html>.
- Palladino, Marco G., Alexandra Roulet, and Mark Stabile. 2025. “Narrowing industry wage premiums and the decline in the gender wage gap.” *Labour Economics* 94: 102693. 10.1016/j.labeco.2025.102693.
- Picchio, Matteo and Jan C Van Ours. 2017. “Temporary jobs and the severity of workplace accidents.” *Journal of safety research* 61: 41–51.
- Picchio, Matteo and Jan C. van Ours. 2017. “Temporary jobs and the severity of workplace accidents.” *Journal of Safety Research* 61: 41–51. 10.1016/j.jsr.2017.02.004.
- Rosen, Sherwin. 1974. “Hedonic prices and implicit markets: product differentiation in pure competition.” *Journal of political economy* 82 (1): 34–55.
- Rosen, Sherwin. 1986. “The theory of equalizing differences.” *Handbook of labor economics* 1: 641–692.
- Wiswall, Matthew and Basit Zafar. 2018. “Preference for the Workplace, Investment in Human Capital, and Gender.” *The Quarterly Journal of Economics* 133 (1): 457–507. 10.1093/qje/qjx035.

Appendix A. Model

A.1. Derivation of the wage equation

Proof of Proposition 1. Start from the healthy worker's HJB (suppressing (p, λ) to lighten notation):

$$rW_g^H = [w_g - \phi_g \lambda] + \delta (W_{g,i}^0 - W_g^H) + \lambda (W_g^I - W_g^H).$$

Rearrange to isolate w_g :

$$w_g = \phi_g \lambda + rW_g^H + \delta (W_g^H - W_{g,i}^0) + \lambda (W_g^H - W_g^I).$$

Add and subtract $rW_{g,i}^0$ inside the term rW_g^H :

$$rW_g^H = rW_{g,i}^0 + r(W_g^H - W_{g,i}^0),$$

so that

$$w_g = \phi_g \lambda + rW_{g,i}^0 + r(W_g^H - W_{g,i}^0) + \delta (W_g^H - W_{g,i}^0) + \lambda (W_g^H - W_g^I).$$

Collect the terms multiplying $W_g^H - W_{g,i}^0$:

$$w_g = \phi_g \lambda + rW_{g,i}^0 + (r + \delta) (W_g^H - W_{g,i}^0) + \lambda (W_g^H - W_g^I).$$

Finally, group the two risk-related terms:

$$\phi_g \lambda + \lambda (W_g^H - W_g^I) = \lambda (W_g^H - W_g^I + \phi_g).$$

Substituting this into the expression for w_g yields

$$w_g = rW_{g,i}^0 + (r + \delta) (W_g^H - W_{g,i}^0) + \lambda (W_g^H - W_g^I + \phi_g),$$

which is (2). □

A.2. Derivation of the surplus characterization

Proof of Lemma 1. Substitute the HJBs into the surplus definitions and use $W_g^H + J_g^H = S_g^H + W_{g,i}^0$ to obtain the system

$$(A1) \quad \begin{aligned} (r + \delta + \lambda)S_g^H &= (p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 + \lambda S_g^I, \\ (r + \delta + \eta)S_g^I &= (\alpha p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 + \eta S_g^H. \end{aligned}$$

Writing $A \equiv r + \delta + \lambda$ and $B \equiv r + \delta + \eta$, the system becomes

$$\begin{pmatrix} A & -\lambda \\ -\eta & B \end{pmatrix} \begin{pmatrix} S_g^H \\ S_g^I \end{pmatrix} = \begin{pmatrix} (p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 \\ (\alpha p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 \end{pmatrix}.$$

The determinant is $D = AB - \eta\lambda = (r + \delta)^2 + (r + \delta)(\lambda + \eta) > 0$.

Solving for S_g^H by Cramer's rule:

$$\begin{aligned} D \cdot S_g^H &= B \left[(p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 \right] + \lambda \left[(\alpha p - c(\lambda) - \phi_g \lambda) - rW_{g,i}^0 \right] \\ &= (B + \alpha\lambda)p - (B + \lambda) \left[c(\lambda) + \phi_g \lambda + rW_{g,i}^0 \right]. \end{aligned}$$

Using $B + \lambda = r + \delta + \eta + \lambda$ and $D = (r + \delta)(r + \delta + \eta + \lambda)$, we have

$$\frac{B + \lambda}{D} = \frac{1}{r + \delta}.$$

Hence

$$S_g^H = \frac{B + \alpha\lambda}{D} p - \frac{1}{r + \delta} \left[c(\lambda) + \phi_g \lambda + rW_{g,i}^0 \right],$$

which establishes (3) with $\kappa_p(\lambda) = (B + \alpha\lambda)/D$.

To obtain the alternative form (4), note that

$$\kappa_p(\lambda) = \frac{B + \alpha\lambda}{D} = \frac{(r + \delta + \eta) + \alpha\lambda}{(r + \delta)(r + \delta + \eta + \lambda)} = \frac{1}{r + \delta} \cdot \frac{(r + \delta + \eta) + \alpha\lambda}{(r + \delta + \eta) + \lambda}. \quad \square$$

A.3. Derivation of the AKM representation

Proof of Proposition 2. Under state-independent wages, the healthy worker's HJB can be rearranged as

$$(A2) \quad w_g(p, \lambda) = rW_{g,i}^0 + (r + \delta) \left(W_g^H - W_{g,i}^0 \right) + \phi_g \lambda,$$

where the first term is the flow value of the outside option, the second converts the stock gain from the match into a flow at rate $r + \delta$, and the third compensates for the direct flow risk cost from accident risk.

Nash bargaining implies $W_g^H - W_{g,i}^0 = \beta_g S_g^H$. Substituting this and the surplus expression from Lemma 1 into (A2):

$$\begin{aligned} w_g(p, \lambda) &= rW_{g,i}^0 + (r + \delta) \beta_g \left[\kappa_p(\lambda) p - \frac{c(\lambda) + \phi_g \lambda + rW_{g,i}^0}{r + \delta} \right] + \phi_g \lambda \\ &= rW_{g,i}^0 + (r + \delta) \beta_g \kappa_p(\lambda) p - \beta_g c(\lambda) - \beta_g \phi_g \lambda - \beta_g rW_{g,i}^0 + \phi_g \lambda \end{aligned}$$

$$= (1 - \beta_g) r W_{g,i}^0 + (r + \delta) \beta_g \kappa_p(\lambda) p - \beta_g c(\lambda) + (1 - \beta_g) \phi_g \lambda.$$

For worker i at firm j , this gives (5) with $\alpha_i^g = (1 - \beta_g) r W_{g,i}^0$ and

$$\psi_j^g = (r + \delta) \beta_g \kappa_p(\lambda_j) p_j - \beta_g c(\lambda_j) + (1 - \beta_g) \phi_g \lambda_j. \quad \square$$

Appendix B. Numerical illustration of the model

This appendix provides a calibrated numerical example that (i) illustrates the economics behind the firm effect in Proposition 2 and (ii) checks that the linear specification used in the empirical analysis is a good approximation over the empirically relevant range of injury risk.

Throughout, the (exact) firm effect for gender g can be written as the sum of three interpretable components:

$$(B1) \quad \psi_j^g = \underbrace{(r + \delta) \beta_g \kappa_p(\lambda_j) p_j}_{\text{productivity pass-through}} - \underbrace{\beta_g c_j}_{\text{safety-cost pass-through}} + \underbrace{(1 - \beta_g) \phi_g \lambda_j}_{\text{risk compensation}}.$$

where $c_j \equiv c(\lambda_j)$ denotes the firm's safety-cost term. The linear approximation in equation (7) replaces $\kappa_p(\lambda_j)$ by a constant $\bar{\kappa}_p$ and uses a local linearization of safety costs in λ .

B.1. Calibration

We choose a baseline calibration that matches key time-scale moments in the Italian labor market and reproduces the empirical firm-level coefficient ratios in Table 4. Table B1 reports the parameter values used throughout this appendix.

The job destruction rate $\delta = 0.10$ implies an expected job duration of $1/\delta \approx 10$ years. The recovery rate $\eta = 6$ implies that injured workers return to full productivity after roughly two months on average, consistent with the injury-duration threshold used to construct our injury measure.

B.2. Stability of the productivity coefficient

A first ingredient behind the linear approximation in equation (7) is that the productivity loading $\kappa_p(\lambda)$ varies little across firms. Recall from equation (4) that

$$\kappa_p(\lambda) = \frac{1}{r + \delta} \cdot \frac{(r + \delta + \eta) + \alpha \lambda}{(r + \delta + \eta) + \lambda}.$$

TABLE B1. Baseline calibration

Parameter	Description	Value	Source/Target
<i>Structural parameters</i>			
r	Discount rate	0.05	Standard
δ	Job destruction rate	0.10	≈ 10 -year avg. tenure
η	Recovery rate	6.0	≈ 2 -month avg. injury duration
α	Injured productivity	0.6	40% productivity loss
<i>Gender-specific parameters</i>			
β_M	Bargaining power (men)	0.50	Normalization
β_F	Bargaining power (women)	0.47	$\gamma_F/\gamma_M = 0.94$
ϕ_M	Risk disutility (men)	1.00	Normalization
ϕ_F	Risk disutility (women)	0.75	Illustrative baseline (see Table D.4)

Note: The recovery rate $\eta = 6$ implies an average injury duration of about two months ($1/\eta \approx 0.17$ years). Gender-specific parameters are disciplined by the empirical ratios $\gamma_F/\gamma_M = 0.94$ and $\delta_F/\delta_M = 0.86$ from Table 4. Interpreting δ_F/δ_M in terms of ϕ_F/ϕ_M requires an assumption on the local slope c_1 of the (locally) linearized safety-cost schedule.

TABLE B2. Stability of $\kappa_p(\lambda)$

Injury rate λ	$\kappa_p(\lambda)$	Deviation from $\kappa_p(0)$
0%	6.667	—
1%	6.662	−0.06%
2%	6.658	−0.13%
3%	6.654	−0.19%
5%	6.645	−0.32%

Note: Computed using baseline parameters from Table B1. The empirically relevant range of firm-level injury rates is approximately 1–5 percent.

Because empirically relevant injury rates are only a few percentage points per year, λ is small relative to $(r + \delta + \eta)$ (which is dominated by η under the calibration). As a result, $\kappa_p(\lambda)$ is close to $\frac{1}{r+\delta}$ and its cross-firm variation is of order $\lambda/(r + \delta + \eta)$.

Table B2 reports $\kappa_p(\lambda)$ over the empirically relevant range of injury rates.

Across $\lambda \in [0.01, 0.05]$, $\kappa_p(\lambda)$ varies by less than 0.35 percent. Quantitatively, this implies that cross-firm wage variation coming from productivity p_j loads on an approximately common coefficient $(r + \delta)\beta_g \bar{\kappa}_p$, with negligible distortion from the λ -dependence of $\kappa_p(\lambda)$. The “constant- κ_p ” approximation is quantitatively innocuous.

B.3. Wage decomposition at a representative firm

To illustrate the mechanics of (B1), Table B3 reports the three components of the firm effect for a representative firm with productivity $p = 1.3$, injury rate $\lambda = 0.02$, and safety cost $c = 0.03$.

The decomposition separates three economically distinct channels. First, the produc-

TABLE B3. Wage decomposition at a representative firm

Component	Men	Women	Ratio
Productivity: $(r + \delta)\beta_g \kappa_p p$	0.649	0.610	0.94
Safety cost: $-\beta_g c$	-0.015	-0.014	0.94
Risk compensation: $(1 - \beta_g)\phi_g \lambda$	0.010	0.008	0.80
<i>Firm effect</i> ψ_j^g	0.644	0.604	0.94

Note: Firm characteristics: $p = 1.3$, $\lambda = 0.02$, $c = 0.03$. The ratios in the last column summarize how gender differences in bargaining power β_g and risk cost ϕ_g map into each component.

tivity component scales one-for-one with bargaining power, so gender differences in β_g translate directly into differences in the productivity pass-through. Second, the safety-cost component also scales with β_g , reflecting how much of the safety-cost burden workers internalize through bargaining. Third, the risk-compensation component depends on $(1 - \beta_g)\phi_g$, so it combines bargaining power and risk cost; this is the object that the empirical coefficient δ_g captures in the linear projection.

Because empirically relevant injury rates are small, the level of the firm effect is dominated by the productivity component, so the overall firm-effect ratio closely tracks the bargaining-power ratio even though risk compensation differs across genders.

B.4. Linear approximation validation

We validate the linear approximation used in the empirical analysis with a simulation in which the *exact* firm effects implied by the model are known. We draw 2000 firms with heterogeneous productivity p_j and injury risk λ_j , compute exact firm effects from (B1), and then estimate by OLS the same linear projection used in the data:


$$\psi_j^g = a_g + \gamma_g p_j + \delta_g \lambda_j + \varepsilon_j^g.$$

The exercise asks a simple question: does the linear projection recover the model-implied composite loadings on productivity and injury risk over the empirically relevant support?

In the simulation, productivity is log-normal and injury risk is centered at a few percentage points per year and truncated to a plausible support. Safety costs are constructed to be decreasing in λ_j —firms with lower injury rates spend more on safety—with a small idiosyncratic component. This delivers a transparent environment in which the nonlinear theoretical object ψ_j^g can be compared directly to its linear approximation.

Figure B1 summarizes the key mechanics. Panel A shows that the productivity loading $\kappa_p(\lambda)$ varies negligibly over $\lambda \in [0.01, 0.05]$, so replacing $\kappa_p(\lambda)$ with a constant $\bar{\kappa}_p$ has little quantitative bite. Panels B and C show scatter plots of the exact firm effects against p_j and λ_j , together with the corresponding linear fits. Panel D reports the implied KOB

FIGURE B1. Numerical illustration of the model



Note: Panel A reports $\kappa_p(\lambda)$ and highlights the empirically relevant range $\lambda \in [0.01, 0.05]$. Panels B and C plot simulated firm effects against productivity and injury risk, together with linear fits. Panel D reports the KOB decomposition of the simulated gender gap in firm effects implied by the model.

decomposition of the simulated gender gap in firm effects.

The estimated coefficients closely match their theoretical counterparts implied by the approximation. Let $\bar{\kappa}_p$ denote the average of $\kappa_p(\lambda_j)$ over simulated firms. Then the theory implies $\gamma_g = (r + \delta)\beta_g\bar{\kappa}_p$, and we obtain:

$$\hat{\gamma}_M = 0.499 \quad \text{vs.} \quad (r + \delta)\beta_M\bar{\kappa}_p = 0.499,$$

$$\hat{\gamma}_F = 0.469 \quad \text{vs.} \quad (r + \delta)\beta_F\bar{\kappa}_p = 0.469.$$

The linear projection fits essentially perfectly. To quantify the remaining nonlinear component, the 95th-percentile absolute approximation error is 0.0020 for men and 0.0019 for women (about 1.27 and 1.27 percent of the cross-firm standard deviation of ψ), and the maximum absolute error is 0.0045 and 0.0043 (i.e., 2.88 and 2.88 percent of the SD, respectively).

Similarly, the model implies that the composite loading on injury risk is $\delta_g = \beta_g c_1 + (1 - \beta_g)\phi_g$, where c_1 is the local slope of the safety-cost schedule used to generate c_j in the simulation. Consistent with this prediction:

$$\begin{aligned}\hat{\delta}_M &= 0.606 \quad \text{vs.} \quad \beta_M c_1 + (1 - \beta_M)\phi_M = 0.650, \\ \hat{\delta}_F &= 0.497 \quad \text{vs.} \quad \beta_F c_1 + (1 - \beta_F)\phi_F = 0.538.\end{aligned}$$

Remaining differences reflect finite-sample noise and the idiosyncratic component in simulated safety costs.

B.5. Parameter restrictions implied by empirical findings

The firm-level empirical coefficients in Table 4 map directly into restrictions on the structural parameters. Using equation (8),

$$\gamma_g = (r + \delta)\beta_g \bar{\kappa}_p, \quad \delta_g = \beta_g c_1 + (1 - \beta_g)\phi_g,$$

where $c_1 > 0$ is the slope of the locally linearized safety-cost schedule around the relevant range of λ .

Productivity response ratio. The finding $\gamma_F/\gamma_M = 0.94$ pins down the bargaining-power ratio:

$$\frac{\beta_F}{\beta_M} = 0.94,$$

because $\bar{\kappa}_p$ cancels in the ratio. Interpreted through the model, women capture a smaller share of match rents, which proportionally reduces their productivity pass-through.

Injury-risk response ratio. The finding $\delta_F/\delta_M = 0.86$ restricts a *combination* of bargaining power and risk cost:

$$\frac{\beta_F c_1 + (1 - \beta_F)\phi_F}{\beta_M c_1 + (1 - \beta_M)\phi_M} = 0.86.$$

Unlike the productivity response, this restriction does not point-identify ϕ_F/ϕ_M without an assumption on c_1 , because δ_g loads both on safety-cost pass-through ($\beta_g c_1$) and on risk compensation $((1 - \beta_g)\phi_g)$.

Table B4 reports the implied ratio ϕ_F/ϕ_M under alternative values of c_1 .

Across the range of plausible c_1 values, matching $\delta_F/\delta_M = 0.86$ requires $\phi_F/\phi_M < 1$ in this model: conditional on bargaining power being lower for women (as implied by γ_F/γ_M), the observed *lower* wage response to workplace injury risk for women can only be reconciled if the effective marginal risk cost of firm-level injury risk is lower for women.

TABLE B4. Implied risk-cost ratio given $\delta_F/\delta_M = 0.86$


Safety cost slope c_1	Implied ϕ_F/ϕ_M
0.0	0.81
0.5	0.77
1.0	0.74
2.0	0.66

Note: Computed assuming β_M and β_F as in Table B1 and normalizing $\phi_M = 1$. For each value of c_1 , the table reports the ϕ_F/ϕ_M that satisfies $\delta_F/\delta_M = 0.86$.

This pattern admits several (not mutually exclusive) interpretations. First, ϕ_g captures the worker's risk cost—reflecting both preferences and effective exposure to injury risk within the firm (e.g., task assignment)—rather than global risk preferences alone; women may be disproportionately allocated to safer tasks within high- λ firms, weakening the mapping from firm-level λ to expected utility losses. Second, within-firm task allocation may generate compositional effects that are not captured by a single firm-level injury rate. Third, gender differences in how injury risk enters wage bargaining (e.g., through beliefs, outside options, or institutional features) may matter in ways not captured by the symmetric Nash bargaining benchmark.


Appendix C. Additional figures

FIGURE C1. Distribution of injuries across sectors in INAIL data




Note: The figure shows the distribution of workplace injuries recorded in INAIL over the period 2014–2018, by NACE Rev.2 section and injury duration. For each sector, one bar reports the share of injuries in that sector over total injuries lasting less than 7 days, while the other bar reports the share of injuries in that sector over total injuries lasting 7 days or more. Injuries with a duration of at least 7 days correspond to those reported in INPS administrative records and constitute the set of injuries used in the empirical analysis.

FIGURE C2. Injuries over time




Note: The figure reports total injuries (in thousands, left axis) and the injury rate (injuries as a percentage of total employment, right axis), shown separately for male and female workers over time. Source: Authors' calculations based on the INPS-BdI sample.

FIGURE C3. Repeated injury risk




Note: The figure reports the share of workers who experience more than one injury over their careers, relative to the total number of workers in the INPS-BdI sample during the period 2005–2019.

FIGURE C4. Injury rate by detailed industry and gender




Note: The figure reports the average injury rate by sector for male and female workers. Sectors are grouped according to the OECD-38 classification of NACE Rev. 2 industries.

FIGURE C5. Injury rate by contract and work schedule type, by gender




Note: The figure reports average injury rates by contract type and work schedule for male and female workers in Panels A and B, respectively.

FIGURE C6. Injuries distribution




Note: The figure reports the distribution of injury occurrences across firms. Each bar represents the share of firms in the BdI-INPS sample that experience the number of injuries shown on the horizontal axis over the period 2005–2019..

FIGURE C7. Exogenous mobility tests for injuries




Note: The figure reports exogenous mobility tests assessing the validity of the AKM model for injury probability. The tests exploit worker mobility across firms, grouping firms into quartiles based on average coworker injury probabilities. Panel A reports average injury probabilities from two years before to one year after the move for workers moving from first-quartile firms to any quartile and for workers moving from fourth-quartile firms to any quartile. Panel B presents the corresponding profiles for workers moving within the same quartile. Panels C and D report changes in injury probabilities for symmetric moves, plotting changes for workers moving from quartile q to q' against those for workers moving from q' to q . Panel C shows raw changes, while Panel D shows adjusted changes after controlling for a quadratic in age, broad occupation-by-tenure dummies, and year fixed effects.

FIGURE C8. Distribution of residuals by (injury) firm and worker effect deciles



Note: The figure reports a surface plot of average AKM residuals from equation (9), by deciles of worker and firm effects.


FIGURE C9. Joint distribution of (injury) firm and worker effect deciles




Note: The figure reports a surface plot of the share of workers in deciles of AKM worker and firm effects from equation (9).

FIGURE C10. Correlations of gender-specific wage firm effects with excess log value added per worker and injury firm effects


A. VA/L (norm.) vs. wage FE (no controls)




B. Injury FE vs. wage FE (no controls)



C. VA/L (norm.) vs wage FE (controls)



D. Injury FE vs. wage FE (controls)



Note: The figure reports binned scatter plots of the relationship between firm wage effects and firm characteristics. Panels A and B show the unconditional relationships between firm wage effects and, respectively, excess log value added per worker and firm injury effects. Panels C and D report the corresponding relationships after controlling for a second-order polynomial in firm size, 5 macro-area dummies, and 38 broad sector dummies. Excess log VA/L is defined as mean log value added per worker, computed over each firm-specific timespan, in deviation from the average in the low-surplus sector (food and accommodation). Firm injury effects are the firm fixed effects from the AKM model estimated on worker-level injury data. The sample is restricted to firms in the dual connected set, and all relationships are weighted by the firm-level number of person-year observations.

Appendix D. Additional tables

TABLE D1. Firm-level rent- and risk-sharing models, sales per worker

Dependent variable:	Male firm effect		Female firm effect	
	(1)	(2)	(3)	(4)
<i>A. OLS</i>				
Excess log sales/L	0.098*** (0.008)	0.081*** (0.004)	0.091*** (0.008)	0.074*** (0.004)
Injury firm effect	0.847*** (0.183)	0.680*** (0.099)	0.746*** (0.198)	0.539*** (0.119)
R-sq.	0.241	0.457	0.179	0.353
Observations	76,165	76,165	76,165	76,165
<i>B. 2SLS</i>				
Excess log sales/L	0.096*** (0.008)	0.080*** (0.004)	0.090*** (0.008)	0.074*** (0.004)
Injury firm effect	1.684*** (0.336)	1.481*** (0.204)	1.564*** (0.355)	1.282*** (0.254)
F-stat	229.3	386.3	229.3	386.3
R-sq.	0.211	0.434	0.155	0.338
Observations	75,061	75,061	75,061	75,061
Controls	No	Yes	No	Yes

Note: The table reports cross-sectional regression results for firms included in the dual connected sample. Panel A shows OLS estimates, panel B shows 2SLS estimates in which the firm injury effect is instrumented with a split-sample strategy. *Excess log sales/L* are mean log sales per worker in deviation from the average in the low surplus sector (food and accommodation). The mean is computed over each firm-specific time-span. *Injury firm effect* if the firm effect from the AKM model estimated on worker-level data on injury probability. Controls include a second-order polynomial in firm size, 5 macro-area dummies, and 38 broad sector dummies. All regressions are weighted by the firm-level number of person-year observations. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.