ISSN 1936-5330



Cautious Careers: Job Mobility under Incomplete Markets

Alex Clymo, Piotr Denderski, Yusuf Mercan, Benjamin Schoefer

August 2026
RWP 26-09
https://doi.org/10.18651/RWP2026-09



Cautious Careers: Job Mobility under Incomplete Markets

Alex Clymo
PSE

Piotr Denderski
University of Leicester & Institute of Economics, Polish Academy of Sciences

Yusuf Mercan
KC Fed

Benjamin Schoefer
UC Berkeley

August 2026

Job mobility is risky, workers are risk averse, and insurance markets are incomplete. This paper studies how these features curb and distort job-to-job transitions by making workers excessively cautious, placing too much weight on job safety over wage and productivity gains. We demonstrate this tradeoff by eliciting employed workers' wage-safety indifference curves in a custom, representative survey. On average, employed U.S. workers require a 1.63 percent pay raise to accept each additional percentage point of annual unemployment risk in a new job. We assess the macroeconomic consequences of our mechanism by embedding it into a general equilibrium search model. Jobs differ in both wages (productivity) and unemployment risk, and risk-averse workers self-insure against unemployment risk through a non-state-contingent bond while searching on and off the job. We find that a complete markets counterfactual would boost job mobility by 12 percent and productivity by 0.19 percent. We also highlight a new role for unemployment insurance—encouraging employed workers to accept risky but high-productivity offers, thereby increasing productivity (by 1.3 percent) and job creation, as well as job loss and unemployment.

This paper subsumes and merges a dissertation chapter titled "Unemployment Insurance and Worker Reallocation: The Experimentation Channel in Job-to-Job Mobility" (2018) and previous conference drafts titled "Job Mobility and Unemployment Risk" and "Job Search and Experimentation in the Standard Incomplete Markets Model" (first presented in March 2019). We are grateful to Jake Bradley, Carlos Carrillo-Tudela, Pierre Cahuc, Kyle Herkenhoff, Axel Gottfries, André Luduvice, Matthew McKernan, Guido Menzio, Benjamin Moll, Giuseppe Moscarini, Peter Rupert, Ayşegül Şahin, Alireza Sepahsalari, and Ludo Visschers. We thank Christina Kulp, of the Federal Reserve Bank of Kansas City Research Library, for her support in coordinating our survey with NORC. We also thank participants of the 2019, 2021, 2024, and 2025 SED Meetings, the 2019 Search and Matching Workshop at the University of Bristol, Search and Matching Annual Conference in Oslo 2019, Midwest Macro 2019, 3rd Dale T. Mortensen Conference in Aarhus, St. Louis Fed-WashU Macro Labor Conference, Recent Developments in the Macroeconomics of Labor Markets Conference in Istanbul, Summer Workshop on Macro and Finance Warsaw, SCE/CEF Conference Nice, Econometric Society European Meeting 2023, Theories and Methods in Macro (T2M) 2023, Royal Economic Society Annual Conference, Fed System EH Conference, Mannheim Workshop on Quantitative Macroeconomics 2024, 5th Labor Markets and Macroeconomic Outcomes Conference at Vanderbilt University, Barcelona Summer Forum 2023, and seminar attendees at HEC Montréal, University of Glasgow, University of Sussex, University of Melbourne, UC Berkeley, Paris School of Economics, Charles III University of Madrid, University of Manchester, Sciences Po, PSE-OFCE-ScPo Macro Group, and National University of Singapore for comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

1 Introduction

Besides allocating workers between employment and nonemployment, the labor market reallocates a comparable number of employed workers between different jobs and firms (Davis and Haltiwanger, 1999), and such direct job-to-job transitions are a crucial factor in wage growth through job ladders (Topel and Ward, 1992; Ozkan et al., 2023; Haltiwanger et al., 2018). Yet, climbing this job ladder is not risk-free, as new jobs may pay more but might be more likely to end in unemployment. While there is ample theoretical, empirical, and policy focus on how incomplete financial markets and risk aversion may affect the job choices of the *unemployed*, these features are typically either absent from or irrelevant in standard job ladder models—due to risk neutrality or jobs only differing in wages but not in stability or separation risk. Our paper fills this gap as the first dedicated comprehensive empirical and theoretical study of how risk aversion and market incompleteness shape the job mobility of the *employed*.

Our paper proposes, formalizes, and quantitatively examines the following basic idea: when employed workers encounter job opportunities with outside employers, some of those jobs may increase their wages but expose them to higher unemployment risk. If workers are risk averse and insurance markets are incomplete, then unemployment risk looms much larger in this tradeoff than in standard treatments with risk neutrality. As a result, workers stay put in—and actively seek to switch into—safe but unproductive jobs. At the aggregate level, this excessive caution is a source of misallocation that depresses productivity.[1] The mechanism also raises the question of policy intervention that may undo some of the distortions caused by financial constraints on job mobility. Specifically, unemployment insurance may gain a novel role as a subsidy for risky job taking of the employed, and hence catalyze labor market dynamism.

We first provide direct empirical evidence for this wage-job safety tradeoff. We do so by eliciting the indifference curve (IC) in job security and wage space. Integrating a custom module into the representative US AmeriSpeak Omnibus Survey, we present employed workers with a menu of hypothetical outside job offers akin to their current jobs, but differing in unemployment risk, and elicit associated reservation wages to accept those jobs. This tradeoff and the ICs that the workers' responses trace out directly map into our model. In the data, workers are willing to give up about 1.63% of wages to reduce annual unemployment risk by 1 pp. Consistent with the role of incomplete markets, this tradeoff is about 75% larger for borrowing-constrained households (as proxied by self-reported consumption declines upon job loss). We further show basic evidence for substantial dispersion in job risk even conditional on the wage and vice versa, so that these tradeoffs are consequential for labor market decisions.

Second, we formalize the mechanism and explore its workings and impact in an equilibrium search model with three key features: (i) both on-the-job and unemployed (random) search, (ii) job offers with heterogeneity in productivity and separation (unemployment) risk, and (iii)

[1] Throughout, we define "misallocation," "distortion," and "excessive" relative to *productive* (output) efficiency—the benchmark formalized in Appendix A, where an output-maximizing, effectively risk-neutral/perfectly-insuring planner dictates decisions.

risk aversion and incomplete markets through self-insurance in an interest-bearing but non-state-contingent asset subject to borrowing constraints. While there are existing models with some of these features *separately*, ours, with its first version dating to 2018, emerged as the first attempt to feature them *jointly* and comprehensively study their labor market impacts.

Importantly, we construct ICs of job safety and wages in the model. Intuitively, the IC's slope proxies the "price of risk": the difference between the values of the current job and unemployment. The average IC slope in the baseline model is 1.53. Market incompleteness tilts this IC in favor of job safety: a risk-neutral IC slope falls to 1.14, whereas a hand-to-mouth version rises to 2.65. Strikingly, we find that the (untargeted) baseline model IC slope (1.53)—the key nexus of our mechanism—closely matches the empirical counterpart from our survey (1.63). Additionally, the model is calibrated to deliver a realistic marginal propensity to consume and, though untargeted, generates a realistic earnings cost of job loss (Davis and von Wachter, 2011; Jarosch, 2023).

Third, with the model in hand, we conduct counterfactuals to show the quantitative aggregate relevance of the wage-job safety tradeoff. In doing so, we also endogenize labor demand and demonstrate a novel feedback between employed workers' willingness to accept (risky) outside offers and firms' incentives to create jobs.

Our first quantitative experiment studies the aggregate effects of asset market incompleteness in partial equilibrium. We contrast the steady state of our baseline model with risk-neutral and risk-averse but hand-to-mouth (HtM) counterfactuals. The resulting shifts in the IC slopes correspond to very different labor market flows. Under risk neutrality (equivalently, complete markets), the employer-to-employer (EE) transition rate increases by about 12%, coinciding with a rising employment-to-unemployment (EU) rate. These changes capture the higher willingness to accept risk, leading to an increase in both the unemployment rate and productivity. The HtM counterfactual (extreme market incompleteness) moves each of these outcomes in the opposite direction, lowering the EE and EU rates, unemployment, and productivity.

Our second experiment revisits the role of unemployment insurance (UI) in general equilibrium. We consider a dramatic reduction in UI generosity (comparable to falling back on US non-UI basic social assistance programs). Two new effects emerge. First, UI subsidizes the acceptance of risky jobs by the employed, so its withdrawal amplifies the distortion stemming from incomplete markets. Second, the reduced willingness of the employed to accept riskier outside offers *discourages* job creation: the equilibrium offer arrival rate *falls* as UI is reduced below a 50% replacement rate—the opposite of the standard homogeneous-risk model, where it *rises* monotonically as UI is reduced. Overall, after UI is reduced, aggregate labor productivity drops by about 1.3%. The EE margin—employed workers now less willing to take risks—accounts for essentially all of this effect, exceeding the standard channel that operates through the reduced wage-based selectivity of the unemployed. This productivity loss comes alongside a fall in unemployment—so, conversely, more generous UI raises productivity while subsidizing job loss and on net raising unemployment.

Finally, we study the business cycle implications of our mechanism. In a recession, a lower job-finding rate makes unemployment more costly and hence raises the price of risk, steepening

the wage-job safety ICs. This changes the *composition* of EE moves, as workers redirect toward safer, lower-paying jobs—i.e., a *climb to safety*. A countercyclical UI expansion dampens this repricing of risk and largely undoes the climb to safety motive, at the cost of higher unemployment.

The paper is organized as follows. We discuss related literature below. Section 2 presents empirical evidence on the wage-job safety tradeoff. Section 3 builds a job-ladder model featuring firm entry, workers' consumption-saving decisions under borrowing constraints, and heterogeneity in match productivity and unemployment risk. Section 4 calibrates the model and validates it against the survey-based ICs and other moments. In Sections 5 to 7, we study counterfactuals around market incompleteness and UI as well as business cycle implications. Section 8 concludes.

Related literature. There is a sizable literature that studies how wealth and liquidity constraints may distort the behavior of *unemployed* job seekers under homogeneous job separation risk, and how policies such as UI might alleviate some of the inefficiencies arising at the unemployment-to-employment margin.[2] In this paper, we focus on the implications of incomplete markets for on-the-job search, i.e., the job ladder.

There are two strands of the literature that focus on different subsets of our model's three key features. First, several papers combine on-the-job search and risk aversion but feature *homogeneous separation risk*. For example, Lise (2013) studies the consequences of the job ladder for savings behavior, and Griffy (2021) shows that wealthier agents search less intensively. Hubmer (2018) studies the role of the job ladder and wealth in explaining life-cycle earnings dynamics. Chaumont and Shi (2022) study how wealth affects the tradeoff between wages and job finding rates in a directed search model. Birinci, Karahan, Mercan, and See (2026) study the role of the job ladder in shaping inflation dynamics and monetary policy in the presence of incomplete asset markets.

Second, models with on-the-job search that do feature separation risk heterogeneity (e.g., Pinheiro and Visschers, 2015; Krolikowski, 2017; Jarosch, 2023) typically assume *risk neutrality* or, equivalently, *complete markets*. Pinheiro and Visschers (2015) focus on the compensating differentials mechanism (see our analogous Equation (11)) under wage posting. Jarosch (2023) also features risk neutrality with wage bargaining, and focuses on the effect of the risky job ladder on persistent earnings losses after job displacement.

We view our investigation as the first comprehensive empirical (including direct survey elicitations of the wage-job safety tradeoff) and theoretical (including general equilibrium) analysis of the three key features with a focus on their labor market implications: job security heterogeneity, incomplete markets with risk aversion, and on-the-job search. While the earliest version of this paper dates back to 2018, three complementary papers have examined more specific implications. The important contemporaneous work by Larkin (2026) proposes a partial-(labor-market-)equilibrium model that focuses primarily on the amplification of the consumption drop during the Great Re-

[2]For an incomplete list of such papers see Acemoglu and Shimer (1999); Marimon and Zilibotti (1999); Lentz and Tranæs (2005); Bloemen and Stancanelli (2001); Algan et al. (2003); Chetty (2008); Krusell et al. (2010); Ravn and Sterk (2017); den Haan et al. (2018); Herkenhoff (2019); Herkenhoff et al. (2024); Braxton et al. (2024); Griffy (2021); Ravn and Sterk (2021); Clymo et al. (2022); Eeckhout and Sepahsalari (2024); Huang and Qiu (2026).

cession. Specifically, his paper links the secular decline in the US job separation rate (see, e.g., Baksy et al., 2026) to the increasing share of illiquid assets in household portfolios. In more recent work, Caratelli (2025) and Pellegrini (2026) build incomplete-markets job ladder models in which EU risk declines with job tenure.[3] Both study how risk slows down the job ladder via precautionary job-keeping motives. Their studies focus on heterogeneity in household wealth (in particular following business cycle shocks) and the implications for income inequality (including from UI). We hope that the three features we dedicate our study to—risk aversion/incomplete markets, heterogeneity in job safety and on-the-job search—remain of interest to the literature and ultimately establish themselves as standard choices in the workhorse model of the labor market.

Finally, we point out novel effects of UI generosity on job creation and aggregate productivity, operating through subsidizing employed workers' willingness to switch into riskier but more productive jobs. Here, our channel also contrasts with the canonical productivity gains from UI for the *unemployed* explored in Acemoglu and Shimer (1999); Marimon and Zilibotti (1999); Acemoglu and Shimer (2000). While our main exposition prioritizes a positive analysis, we provide a thorough discussion of efficiency considerations in Appendix A.

2 Empirical evidence

We start by presenting descriptive evidence that in the US labor market, jobs exhibit substantial heterogeneity in unemployment risk even conditional on pay. We then present direct evidence on the tradeoff between job safety and pay in a custom survey eliciting the "price of risk": the reservation wages capturing employed respondents' willingness to accept job offers with varying separation risks, which trace out ICs.

2.1 Basic evidence for the safety-wage tradeoff

Our first analysis draws on the Current Population Survey (CPS) and the Survey of Consumer Expectations (SCE). We relegate details on sample selection and measurement to Appendix B.1.

EU-job tenure profiles. Heterogeneity in separation risk between potential jobs a given worker may accept is essential to our mechanism. An observable implication of this feature would be a downward-sloping relationship between EU separation rates and job tenure. Such a pattern may arise from compositional effects (as high separation risk jobs leave the sample early) or duration dependence (whereby jobs become stable over time). Because incomplete markets dampen job transitions regardless of which of the two features is present, we choose the former route. Of course, homogeneous separation risk would imply a flat tenure profile, so that the empirical gradient serves as a target to calibrate the heterogeneity in the model.

[3] Previous versions of our paper featured experience-goods or exogenously declining separation rates and yielded similar results to our current specification with permanent match-level heterogeneity in separation risk.

Figure 1 Panel (a) plots the average EU rate by job tenure in the CPS, showing raw averages (no controls) and composition-adjusted versions. The main takeaway is that EU risk falls sharply in tenure even when controlling for observable heterogeneity.[4] Worker controls hardly change this pattern. Adding in job controls leads to a flattening of the profile, especially once we control for the wage. It is precisely this joint determination of wages and unemployment risk that we study in Section 4, and our model matches this gradient to discipline the separation risk distribution.

Prevalence of high-wage, high-risk jobs. Our theory is based on the idea that workers may have to trade off job safety against wages—so that the two dimensions of the job ladder are imperfectly correlated. To show evidence consistent with this intuitive assumption, we construct unemployment risk and wages at the industry-occupation-state level and show the bin-specific, conditional distribution in an employment-weighted "fan chart," along with the regression line (Appendix Figure B.1 presents the underlying cell-level scatter plot).[5]

Figure 1 Panel (b) presents the results. Overall, the figure shows a negative relationship between these predicted EU rates and wages: on average, higher-wage jobs exhibit lower unemployment risk. At the same time, there is substantial heterogeneity in unemployment risk among jobs with similar wage levels, indicated by the wide percentile bands around the downward-sloping regression line. Of course, the figure is based on realized, equilibrium accepted jobs rather than the job offer distributions, and some dispersion may reflect worker heterogeneity or measurement error. Our model is consistent with both this average negative relationship and the dispersion.

Change in predicted EU probability upon EE switch. To further highlight separation risk heterogeneity in job choice, we study the effect of job switching on EU risk changes. We identify an EE transition in the CPS as a respondent who is employed in two adjacent calendar months and reports having changed employers between them (which hence includes short nonemployment spells within the month). A job's EU risk is predicted by a simple regression model in the CPS based on observables.[6] We then calculate the change in this predicted worker-level EU probabil-

[4]Specifically, the figure reports tenure fixed effects α_τ from the regression $\mathbb{I}(EU_{it}) = \alpha_\tau + \gamma_t + \beta X_{it} + \varepsilon_{it}$, where $\mathbb{I}(EU_{it})$ is an indicator for an EU separation between months t and $t+1$, γ_t are calendar time fixed effects and X_{it} are time-varying observables. The figure plots the tenure fixed effects α_τ, and across specifications additively adjusts them so that their employment-weighted averages coincide with the unconditional average rate over the tenure range (0 to 44 months). In X_{it}, we incrementally include *worker* controls (quartic in potential experience, number of (own and own aged 5 and younger) children in the household (linear), and dummies for education, race, gender, and marital status), *job controls* (industry, occupation, state fixed effects), and a *wage control* (log average real wage across (up to) two earnings observations in the Outgoing Rotation Groups).

[5]Specifically, we define a job as an industry-occupation-state combination and estimate $y_{it} = \gamma_t + \eta_j + \zeta_o + \omega_s + \beta X_{it} + \varepsilon_{it}$, where the left-hand side variable $y_{it} \in \{\mathbb{I}(EU_{it}), \log(w_{it})\}$ is either an indicator for an EU separation or the log real wage, γ_t are time fixed effects, $\eta_j, \zeta_o, \omega_s$ are industry, occupation and state dummies, respectively, and X_{it} are the *worker controls* defined in the context of the EU-tenure regressions described in Footnote 4. We predict job-specific EU rates and wages by fixing the time effects and worker controls in the above regression equation to a baseline value and only using the variation in industry-occupation-state dummies. We then run a linear regression of the predicted EU rate on the predicted log real wage at the cell level, weighted by the survey-weighted employment share of each cell, as described in the figure note.

[6]To predict EU risk, we regress the EU separation indicator on the set of worker, job, and wage controls discussed above, as well as calendar time fixed effects, drawing again on the CPS. We then predict the EU rate while fixing the

Figure 1: Motivating empirical evidence: unemployment risk, job tenure, job switching, and wages

(a) EU rate-tenure profiles

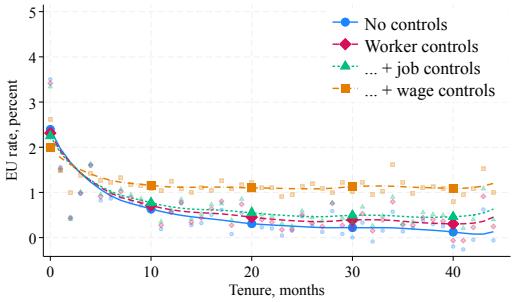

(b) Job-specific EU rates vs. wages

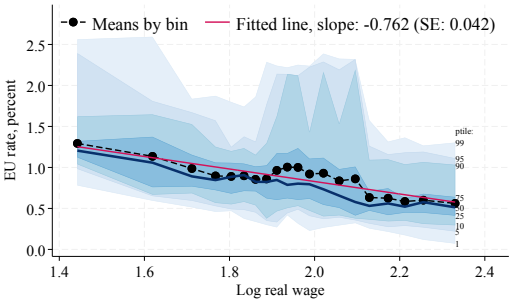

(c) Predicted ΔP(EU): stayers vs. switchers

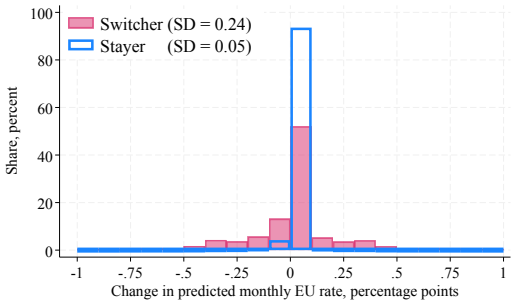

(d) Subjective ΔP(EU): stayers vs. switchers

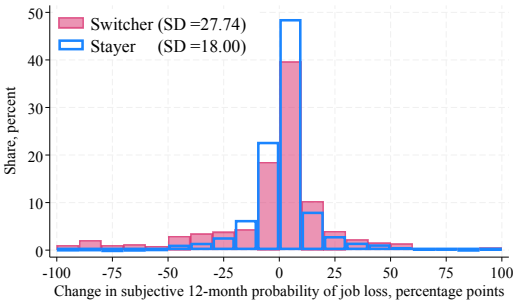

(e) Realized EE transitions vs. risk attitude

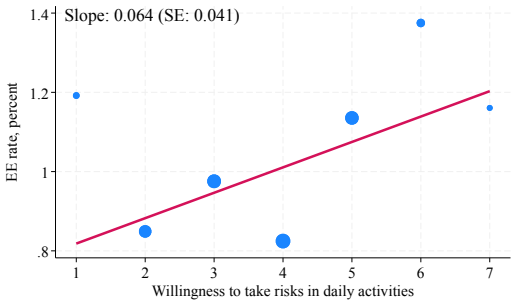

(f) Subjective quit probability vs. risk attitude

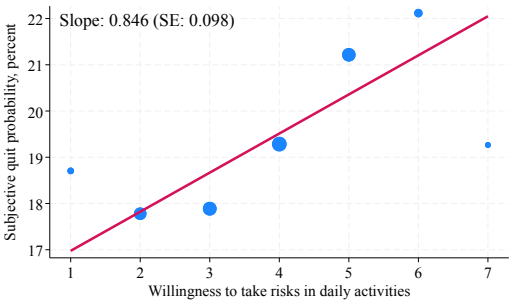

Notes: This figure plots motivating evidence on unemployment risk, job tenure, and job mobility. Panel (a) plots monthly EU rates by job tenure. It reports the tenure fixed effects from variants of the regression described in Footnote 4, rescaled so that for each specification, the employment-weighted EU rate matches the unconditional average EU rate for jobs with less than 45 months of tenure. The slight negative EU rate values arise from including time fixed effects in all specifications. Panel (b) plots the distribution of monthly job-specific (industry-occupation-state) EU rates against log real wages from the regressions described in Footnote 5, fixing time effects and worker observables at baseline values. We split the x-axis variable (wage) into 20 bins, each holding five percent of survey-weighted employment, and, within each bin, compute employment-weighted mean and percentiles of the y-axis variable (job-specific EU rate). The shaded bands span the symmetric percentile pairs labeled at the right edge, the dark blue line marks the median, and the dashed line connects the employment-weighted mean (i.e., the standard binned scatter plot). The red line is the OLS fit estimated on the underlying data (cell-level, weighted by employment shares), with a robust standard error. Appendix Figure B.1 presents the underlying cell-level scatter plot. Panel (c) shows histograms of the one-month change in predicted EU probabilities, and Panel (d) the distribution of the one-month within-worker change in the subjective 12-month job-loss probability, for job switchers and job stayers. In both panels, the sample is restricted to workers observed to switch employers at least once, so that the two distributions are computed over the same respondents. Panels (e) and (f) show weighted scatter plots of realized monthly EE transitions and subjective 12-month quit probabilities against self-reported willingness to take risks in daily activities (discrete values 1–7); circle sizes denote the share of observations at each value, and the regression line and its robust standard error are estimated on the underlying micro data. Panels (a), (b), (e), and (f) use the survey-provided weights; the histograms in Panels (c) and (d) are unweighted. Sources: Panels (a)–(c) use the CPS and Panels (d)–(f) the SCE; see Appendix B.1 for the analysis sample underlying each panel.

ity between two adjacent months a worker is in the sample, separately for job switchers (EE as defined above) and job stayers (the remaining continuously employed workers). We restrict the sample to workers with at least one switch, so that the two distributions are computed over the same respondent sample. Figure 1 Panel (c) plots the histogram of these predicted EU probability changes for job stayers (blue bars) and job switchers (red bars). Not surprisingly, the change in unemployment risk for job stayers is centered around zero. However, for job switchers, there is substantial variation in the change in unemployment risk, on *both* sides of zero. The standard deviation more than quadruples, from 0.05 to 0.24 percent.

Change in subjective job-loss risk upon EE switch. Moreover, we draw on the SCE to show that the subjective job-loss risk moves more with EE transitions than within an ongoing employment spell. Within each worker, we take the one-month difference of the subjective 12-month job-loss probability, and classify a worker-month as a job *switch* if an employer change is reported between the two months. Again, we restrict the sample to workers with at least one switch, so that the two distributions are computed over the same respondent sample. Panel (d) plots the histogram of the within-worker change of this job loss probability separately for job switchers and stayers.[7] Stayers' beliefs have a spike at zero (no change), and switchers perceive wider changes in separation risk, with fatter tails on both sides. The standard deviation increases by more than half, from 18 to 28 percent. (Unreported variance-ratio tests reject equality at the one percent significance level.)

Cautiousness and job mobility. We close our motivating facts with a suggestive illustration of our mechanism in action: risk-takers outside of the labor market also exhibit higher job mobility rates. We take advantage of the SCE's risk attitude questions, providing a measure of the willingness to take risks.[8] We then relate it to realized EE transitions—a respondent employed in two consecutive months reporting to have switched employers in between—as well as the subjective probability of a voluntary separation over the next year. Figure 1 Panels (e) and (f) report these results, respectively. Both relationships are positive. The gradient is steep and precisely estimated for the subjective quit probability, and positive but imprecisely estimated for realized monthly EE transitions, which are rare events. This pattern captures the narrative of the paper: job mobility can be risky, and workers' risk-taking capacity shapes their willingness to switch jobs, and, ultimately, climb the job ladder. To explore this further, our custom survey, presented next, isolates this precise tradeoff and some forces that shape it.

Taking stock. Our analysis shows that (i) unemployment risk declines with job tenure (consistent with compositional effects from permanent job heterogeneity), (ii) workers switch to both safer and riskier jobs (consistent with a two-dimensional job ladder), (iii) while higher-wage jobs are

calendar time effects to a baseline.

[7] Subjective annual separation risk changes would exceed an annualized version of the CPS objective job loss probability changes, reflecting the survey measure's scale amplification that resembles the level differences in, e.g., Hendren (2017). Our custom survey, reported on next, therefore relies primarily on a fixed specified separation risk menu.

[8] Namely, the SCE asks: "On a scale from 1 to 7, how would you rate your willingness to take risks in daily activities?"

generally safer, unemployment risk varies substantially across jobs with similar wages, and (iv) workers who are more willing to take risk (outside of the labor market) exhibit more job mobility. Our model presented in Section 3 and calibrated in Section 4 matches all these patterns.

2.2 Estimating the wage-job safety indifference curve: our own custom survey

We now provide direct evidence from a custom survey on how employed workers trade off job security and wages when considering outside job offers. We are not aware of existing research eliciting these preferences. Figure 2 collects the resulting indifference-curve slope between wages and unemployment risk for employed workers across the survey and model variants and hence serves as a summary graph and bridge throughout the paper.

Custom labor market survey. We fielded a custom survey in May–June 2025 as part of three consecutive bi-weekly waves of the nationally representative AmeriSpeak Omnibus survey administered by NORC at the University of Chicago.[9] Our final analysis sample contains 1,008 respondents. Appendix B.2 provides additional details, including the questionnaire, data cleaning procedures, sample selection choices (e.g., we drop a small share of "speeders" that complete the survey unrealistically fast), descriptive statistics, and our robustness checks.

The key survey questions. We elicit wage-job safety ICs as follows. We have two stepping stones. First, we elicit current pay and the probability that the current job leaves the respondent unemployed at some point in the next 12 months, δ_i:

> We now ask you about the unemployment risk in your current job. What do you think is the probability that this job leaves you unemployed at some point in the next 12 months? (For instance, you get dismissed, you quit, your firm closes and there are layoffs. . .). The numbers below are between 0 and 100 and are reported in percent.
>
> - 0 percent means you are sure you will not become unemployed at some point in the next 12 months.
> - 10 percent means that you think there is a 1 in 10 chance of you becoming unemployed at some point in the next 12 months.
> - 50 percent means it is equally likely that you will become unemployed or not become unemployed at some point in the next 12 months.
> - 100 percent means you are sure you will become unemployed at some point in the next 12 months.
>
> Pick a probability between 0 and 100 percent: [SLIDER, 0–100] %

Second, we elicit the reservation pay w_i^* for a job offer at another employer that is identical to the current job—in particular, carrying the *same layoff risk*.

> You described your current job as follows:

[9] The sample is designed to be representative of the US population, recruited by mail, phone and face-to-face interviews. Dennis (2019) describes the sampling, recruitment and survey administration. Each Omnibus sample is drawn from the AmeriSpeak panel. Interviews are conducted online and by phone. The Omnibus program is designed for shorter questionnaires such as ours and is, e.g., also used by the AP-NORC Center for Public Affairs Research, and in existing labor market research (e.g., Mui and Schoefer, 2025). NORC also runs the General Social Survey.

Figure 2: Summary of wage-job safety indifference curve slopes in the survey and the model

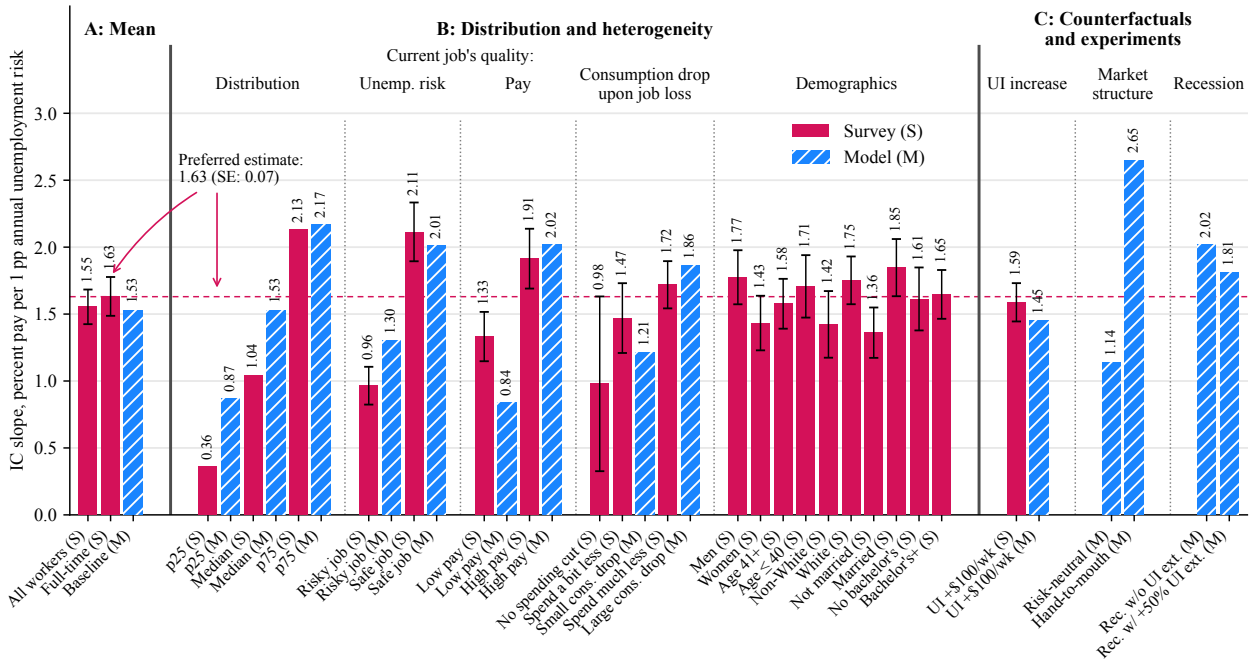

Notes: This figure plots the wage-job safety IC slopes in the survey and the model presented throughout the paper, expressed in additional percent of pay demanded per 1 pp of additional annual unemployment risk, organized into three groups: mean (A), distribution and heterogeneity (B), and counterfactuals and experiments (C). The red solid bars are survey estimates (S) and the blue hatched bars are model counterparts (M). For empirical means, whiskers span the 95% confidence interval, i.e., ± 1.96 standard errors (clustered at the respondent level). For group-level estimates formed as the base effect plus an interaction with a group indicator, the standard error of the sum accounts for the covariance between the two coefficients. In Group A, "All workers (S)" estimates Equation (1) on all workers (Table B.3 Panel A Column (1)); all other survey estimates are based on the full-time sample. "Full-time (S)" is our preferred estimate (Table 1 Column (1)), and "Baseline (M)" is the average model IC slope from the same regression on model-simulated data (Figure 7 Panel (a)). The horizontal dashed line marks this preferred estimate for reference. In Group B, the distribution bars report the 25th percentile, median, and 75th percentile of individual slopes in the survey and the model (Appendix Figure C.3); the survey dispersion partly reflects measurement error. The current job's unemployment-risk and pay bars are from Table 1 Columns (3) and (2): safe and risky jobs carry below- and above-median subjective job-loss probability, and low and high pay refer to below- and above-median annualized gross pay. The consumption-drop bars are from Table 1 Column (4), where the omitted group is those workers who report no change or even an increase in spending upon job loss. The model counterparts in these three blocks (Section 4.3) split workers by the median job separation rate, wage, and consumption decline upon job loss, respectively; the model has no counterpart to the no-spending-cut group. The demographic bars are from Appendix Table B.2 Columns (1)–(5), splitting full-time workers by gender, age, race, marital status, and holding a bachelor's degree. In Group C, the UI increase bars report the IC slope under a hypothetical $100-per-week benefit increase: the survey bar is the constant plus the treatment interaction from the post-UI-information re-elicitations (Appendix Table B.4 Column (3)), and the model bar applies the same benefit increase to the calibrated baseline model (Section 6.1). The risk-neutral and hand-to-mouth slopes correspond to the counterfactual ICs in Figure 7 Panel (b), which hold all other parameters at their baseline values. In the recession block, "Rec. w/o UI ext. (M)" applies to the baseline slope the 32% rise at the onset of the recession (Figure 9 Panel (c)), and "Rec. w/ +50% UI ext. (M)" applies the 18% rise that remains when the recession is paired with a +50% UI expansion (Section 7). Furthermore, Figure 3 Panel (b) provides robustness checks of our preferred estimate of 1.63, omitted here.

- Your [previously elicited pay frequency] pay is $[previously elicited current pay].
- The risk of your current job ending and you becoming unemployed at some point in the next 12 months is [previously elicited subjective unemployment risk]%.

Now imagine a job offer at another employer. If you do not accept this offer, you can still stay at your current job. Suppose the new job is comparable to your current job. In particular, the unemployment risk is the same as in your current job, i.e., there is a [previously elicited subjective unemployment risk]% probability that the new job ends and you become unemployed at some point in the next 12 months. However, the new job may have a different [previously elicited pay frequency] pay. Below, please indicate the lowest [previously elicited pay frequency] pay level (in $) this new job needs to offer for you to accept it:

This stepping stone is intended to provide an anchor for our subsequent hypothetical job offers (having netted out switching costs, as they are absent from our model).

Finally, we elicit the ICs as follows: we ask about the reservation pay, $\widetilde{w}_{ij}$, for hypothetical job offers j that again keep all job dimensions fixed to the current job's, but now vary the unemployment risk, $\widetilde{\delta}_j$, with a menu of 2, 5, 10, 20, and 50 percent (again over the next 12 months, which in the model we interpret to be a constant annual rate):

You stated that you would accept a new job that has:

- an unemployment risk of [previously elicited unemployment risk]% at some point in the next 12 months,
- at least a(n) [previously elicited pay frequency] pay of $[previously elicited reservation pay], and is otherwise comparable to your current job.

But what if the new job had a different level of unemployment risk? By unemployment risk, we again mean the probability that the new job leads to unemployment at some point in the next 12 months. The table below shows different levels of unemployment risk in this new job for you to consider: 2%, 5%, 10%, 20%, and 50%. For each risk level, please enter the lowest [previously elicited pay frequency] pay level at which you would be willing to switch to this new job.

Unemployment risk level	Lowest [pay freq.] pay at which I would accept this job
2% unemployment risk	$
5% unemployment risk	$
10% unemployment risk	$
20% unemployment risk	$
50% unemployment risk	$

The respondent's reservation pay against layoff risk traces out her wage-job safety IC—formally and quantitatively capturing our notion of caution in job mobility throughout the paper. Our baseline analysis sample contains the respondents that fill in all five risk levels and do so monotonically (which is the vast majority), but our results are robust to relaxing either restriction.[10]

[10] Among those who answer all 5 scenarios, 91.4% of them provide weakly monotonic responses in either direction. We explain below why we also permit wrongly signed monotonic answers—thereby only dropping "zig-zag" answers.

Regression specification. To estimate the job safety versus pay tradeoff empirically, we express each reservation pay $\widetilde{w}_{ij}$ as a percent deviation from the respondent's own answer at the 10 percent risk level, $\widetilde{r}_{ij} = 100\,(\widetilde{w}_{ij}/\widetilde{w}_{i,10\%} - 1)$, winsorize it at the 2% level (across respondents within each risk scenario j), and regress this reservation ratio on the unemployment risk in each of the five scenarios $\widetilde{\delta}_j$ including respondent fixed effects α_i:[11]

$$\widetilde{r}_{ij} = \alpha_i + \beta\,\widetilde{\delta}_j + \varepsilon_{ij}. \tag{1}$$

The coefficient β captures the *IC slope*: the percentage increase in reservation pay required to keep respondent i indifferent between two hypothetical jobs whose probability of job loss differs by one percentage point (pp) but are otherwise identical to the current job. Rather than a parameter, this reduced-form object reflects multiple factors such as the utility function, financial market completeness, the labor market environment, and heterogeneity. Our model features this richness, and the average IC slope β forms a quantitative benchmark for cautiousness in job switching.

The preferred estimate. We estimate Equation (1) on all and on full-time employees. The coefficient estimate is 1.55 for all workers and 1.63 for full-time workers, which is our preferred estimate. That is, workers require a 1.63% increase in pay to accept a one pp increase in the annual risk of job loss (out of a job otherwise similar to the respondent's current one). Table 1 reports the full regression results. The effect is precisely estimated, with a respondent-clustered standard error of just 0.074. Figure 3 Panel (a) traces out the underlying relationship between wages (percent) and annual separation risk (pp), showing means of each of the five risk levels while absorbing respondent fixed effects—*i.e., it plots the IC between wages and unemployment risk across jobs*.

Robustness. The coefficient plot in Figure 3 Panel (b) shows that this estimate is remarkably stable across robustness checks. These checks bring in "speeders," permit partial answers or non-monotonic reservation wage series in the main question (a likely sign of inattention), winsorize at 1% rather than 2%, trim rather than winsorize, and use two subsequently repeated IC elicitations that bring in UI considerations. Restricting to weakly positively monotonic ICs changes the result only modestly, while requiring a larger reservation wage for the 50% than the 2% risk level pushes the average IC slope to about 2, reaching about 2.5 if we require strict and positive monotonicity.[12] We also obtain somewhat steeper slopes when normalizing the reservation wages by the current job's reservation wage (rather than the common 10% risk variant). Detailed regression results for these robustness checks are in Appendix Table B.3.

[11] The particular normalization (division) by a respondent's "anchor" reservation pay is inconsequential for the slope; our baseline design does so by $\widetilde{w}_{i,10\%}$, the mid-point of the five common risk levels, but below we report robustness to using instead the separately elicited reservation wage w_i^* for switching into a job with the same risk level as the current one (while winsorizing the resulting pay ratio separately). We here prefer the midpoint normalization by $\widetilde{w}_{i,10\%}$ as it anchors all respondents at the same baseline level of 10% risk, and as the w_i^* concept was elicited on a separate screen with respect to the respondent's own job's subjective job loss probability.

[12] We prefer the baseline specification, which imposes a sign-agnostic weak monotonicity restriction to weed out careless answers while avoiding the positive bias that positive monotonicity entails in the presence of noisy answers.

Figure 3: The tradeoff between pay and job safety: estimating indifference curves

(a) The wage-job safety indifference curve

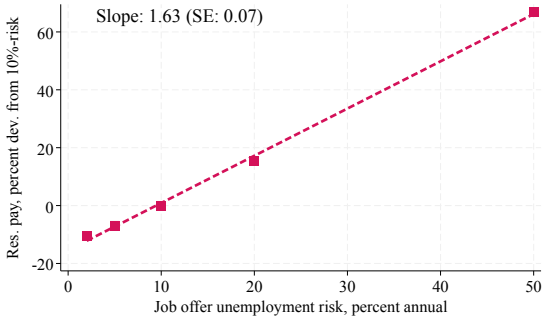

(b) IC slopes across specifications

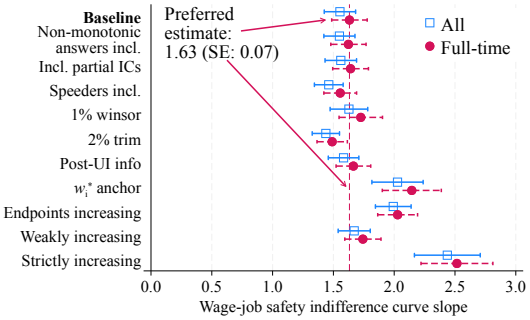

Notes: This figure plots the estimated indifference curves. Panel (a) plots a binned scatter of the reservation ratio $\widetilde{r}_{ij}$ against offered unemployment risk $\widetilde{\delta}_j$, absorbing respondent fixed effects, for the five risk levels. The line presents the linear fit. Panel (b) plots estimates of Equation (1) with 95% confidence intervals with standard errors clustered at the respondent level. Each row changes one restriction relative to the baseline. The vertical dashed line marks our preferred full-time estimate. Appendix B.4 reports the underlying estimates and defines each specification. Source: Own survey.

Heterogeneity in ICs. We further assess the heterogeneity in the IC slopes at the individual level. We start by letting β_i vary across respondents in Equation (1). The full-time sample yields a mean of 1.63, with a standard deviation of 2.14 (Appendix Figure B.3). Of course, this dispersion not only captures the true underlying heterogeneity in wage-job safety tradeoffs but also captures sampling error in the five-point regressions.

Model-consistent drivers of heterogeneity. We then examine economic drivers of this variation. We focus on proxies of heterogeneity active in our model: the value of the job (employment) compared to unemployment value. Hence, we split the sample by the median of gross pay and job loss risk in the current job (driving employment value), and by categories for elicited expected consumption drop following a job loss (driving unemployment value through liquidity constraints and the cost of unemployment).[13] We interact each indicator with the offered risk $\widetilde{\delta}_j$ in Equation (1), so that the worker fixed effect absorbs its level effect and the interaction shifts the IC slope. We do so among full-time job holders (but find similar results for all workers). Additionally, Appendix Table B.2 reports heterogeneity across broad demographic groups rather than proxies for mechanism-motivated heterogeneity.

Table 1 presents the results. Column (1) reiterates the baseline specification. Column (2) shows that workers in the upper half of the annual pay distribution require a 1.9% increase in pay to compensate for a 1 pp rise in EU risk, against 1.3% (1.91 − 0.58) for workers in the lower half. Column (3) indicates that workers in (below-median risk) safe jobs require a 2.1% increase in pay, more than *twice* the 0.96% (2.11−1.15) required by workers in risky, above-median risk jobs.

[13] We construct this measure by asking how respondents' spending would change upon job loss, on a five-point scale from "much less" to "much more than before." See Appendix B.2 for the full questionnaire (`FRKC6`).

Table 1: Heterogeneity in job safety–pay indifference curves

	(1) $\widetilde{r}_{ij}$	(2) $\widetilde{r}_{ij}$	(3) $\widetilde{r}_{ij}$	(4) $\widetilde{r}_{ij}$	(5) $\widetilde{r}_{ij}$
$\widetilde{\delta}_j$	1.632***	1.914***	2.114***	0.979***	1.974***
	(0.074)	(0.114)	(0.112)	(0.333)	(0.356)
$\widetilde{\delta}_j\times$ Low pay		−0.582***			−0.537***
		(0.147)			(0.145)
$\widetilde{\delta}_j\times$ High EU risk			−1.149***		−1.124***
			(0.133)		(0.134)
$\widetilde{\delta}_j\times$ Spend much less				0.740**	0.456
				(0.345)	(0.327)
$\widetilde{\delta}_j\times$ Spend a bit less				0.491	0.224
				(0.359)	(0.345)
N (job offers)	4135	4005	4135	4040	3910
Within R^2	0.352	0.367	0.394	0.356	0.411
N (respondents)	827	801	827	808	782
Individual FE	Yes	Yes	Yes	Yes	Yes
Full-time only	Yes	Yes	Yes	Yes	Yes

Notes: This table presents estimates of Equation (1) including interactions between the offered job risk and person-level indicators. Low pay indicates below-median annualized gross pay and high EU risk indicates above-median subjective job loss probability. The omitted consumption category is workers who would not cut spending (Footnote 13). Standard errors in parentheses are clustered at the respondent level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: Own survey.

Intuitively, both high-pay and safe jobs are valuable relative to unemployment, which raises the compensation their holders demand for additional risk. Finally, Column (4) shows that incomplete markets—and the resulting imperfect insurance—also shape their indifference curves: respondents who anticipate they would cut consumption substantially following job loss demand about 75% higher compensation (0.74 on a base of 0.98) to accept jobs with greater layoff risk. Column (5) attempts a multivariate regression with all three dimensions, showing robustness of effect sizes or directions (in the case of the consumption drop variables), with a loss in precision.

Further validation: stability of individual-level estimates across elicitation. Finally, to further validate the survey ICs, we show that they are stable across reelicitations within individuals. To do so, we use two subsequent questions in the same survey, one after an information treatment clarifying the average US UI replacement rate and the other in the context of a hypothetical $100 boost to UI weekly benefits (discussed in Section 6.1 with questionnaire provided in Appendix B.2). Appendix Figure B.4 scatter plots these two sets of respondent-level β_i IC slopes against their baseline values, and documents their near-perfect stability.[14]

[14]Econometrically, the estimated regression slope between the two elicitations, 0.9, captures the share of stable variation across individuals and can be viewed as a measurement error correction (as in applied micro applications of "split sample IVs," see, e.g., Drenik et al., 2023; Schoefer and Ziv, 2024). We discuss the potential effects of the

Taking stock. Our custom survey yields two findings, which Figure 2 collects in one place, alongside upcoming model counterparts. First, workers demand higher wages to compensate for higher job-loss risk. The average wage premium is about 1.63% of the current pay for a 1 pp increase in risk across a 12-month horizon—a finding robust to, and possibly even underestimated relative to, alternative specifications. Second, workers' tradeoff between job safety and wages is higher for those in jobs that are already safe and highly paying, at the top of the ladder. Demand for safety is also higher among workers with higher anticipated consumption drop upon job loss, a proxy for limited self-insurance. Furthermore, in Section 6.1, we report on a model- and survey-based experiment that hypothetically increases UI generosity, which lends some support to higher UI encouraging job mobility by subsidizing risky job taking and by flattening the ICs. The figure also previews the quantitative performance of the model we develop next. The blue hatched bars report the model counterpart of each estimate, taken from the analysis in Sections 4 through 7. The model, despite not targeting any IC slope, reproduces the average slope very closely and matches the direction of every heterogeneity cut. Incomplete markets prove essential for this fit, as risk-neutral and hand-to-mouth variants miss the average slope in opposite directions.

3 A model of incomplete markets and a productivity-safety job ladder

Our model features a frictional labor market with a two-dimensional job ladder (offer distribution) along match productivity and safety, where workers can partially self-insure against idiosyncratic labor market shocks using a risk-free bond.

3.1 Environment

Time is continuous and indexed by $t \in [0, \infty)$, and there are three types of agents: workers, entrepreneurs, and the government. The economy is stationary.

Preferences. There is a unit mass of risk-averse workers, whose preferences over consumption flow are described by a strictly increasing, twice-differentiable strictly concave utility function $u(c)$. Workers can self-insure against labor income fluctuations using a risk-free bond a with real return r. Risk-neutral entrepreneurs own the firms. The workers and entrepreneurs discount the future at rates $\rho \geq 0$ and $\rho_f \geq 0$, respectively.

The labor market. Workers and firms meet randomly in a frictional labor market. Firms enter freely to post vacancies for single-worker matches at a flow cost of κ. When a worker and firm meet, they draw a match productivity-separation rate pair (z, δ) from CDF $F(z, \delta)$ with support $\left[\underline{z}, \overline{z}\right] \times \left[\underline{\delta}, \overline{\delta}\right] \subset \mathbb{R}^2_+$. Match characteristics are observed before each party accepts or rejects the offer without recall. Productivity z is constant during the match unless it drops to zero at rate δ to destroy the match. Besides these exogenous separations, either party can unilaterally terminate

information treatment and hypothetical UI increase contexts below.

the match. The worker-firm pair agrees on a wage $w(z)$ to split their bilateral match surplus, which we describe in Section 3.3.

Let $g^u(a)$ denote the endogenous mass of unemployed workers with assets a, and $g^e(a, z, \delta)$ that of employed workers further indexed by their job characteristics (z, δ). Thus $n = \int g^e(a, z, \delta) da dz d\delta$ is total employment, $u = \int g^u(a) da$ is total unemployment, and $u + n = 1$ is the worker mass normalized to one. Search efficiency of the unemployed is normalized to $s_u = 1$, and s_e denotes the relative search efficiency of the employed. Aggregate search effort is $S = u + s_e(1-u)$; the aggregate mass of vacancies is V, and labor market tightness is $\theta = V/S$. Worker-firm contacts are governed by a matching function $m = M(V, S)$, which is increasing in both arguments and concave, and exhibits constant returns to scale; hence, contact rates are functions of tightness θ alone. Workers meet a vacancy at rate $\lambda(\theta) = m/S$ per unit of search effort and vacancies meet workers at rate $q(\theta) = m/V$. Under random search, the type of worker a vacancy meets depends on the worker distribution across employment states and job characteristics.

Government. The government levies progressive income taxes to finance flat UI benefits b and exogenous spending G. Both labor and UI income are taxed; capital income is not. Specifically, given gross income y, the after-tax income is $\widehat{y} = \tau y^{1-\Upsilon}$, where τ governs the average level and $\Upsilon \geq 0$ governs the progressivity of taxes, as in Benabou (2002) and Heathcote et al. (2014).[15]

3.2 Worker choices

Workers make two choices given the job offer arrival rate $\lambda(\theta)$: how much to consume and save, and which job offers to accept. For what follows, let the value of an employed worker in a job with productivity z and separation rate δ holding assets a be $W(a, z, \delta)$ and the value of an unemployed worker with assets a be $U(a)$.

Consumption and saving. Assets a held at instant t pay an endogenous instantaneous interest rate r. Given consumption c, savings evolve according to the law of motion

$$\dot{a} \equiv \frac{da}{dt} = \tau y^{1-\Upsilon} + ra - c, \tag{2}$$

where gross income is either a wage (described in Section 3.3) or UI benefits:

$$y = \begin{cases} w(z) & \text{if employed in job with productivity } z \\ b & \text{if unemployed.} \end{cases}$$

[15] Parameter τ is inversely related to the average tax rate. For example, under $\Upsilon = 0$, i.e., no progressivity in the tax schedule, the proportional tax rate becomes $1 - \tau$.

The worker also faces a borrowing constraint $a \geq \underline{a}$. Since consumption flow is finite, this constraint only binds when $a = \underline{a}$, implying $c \leq \tau y^{1-\Upsilon} + r\underline{a}$ (equivalently $\dot{a} \geq 0$).[16]

Job acceptance sets. Workers accept job offers which (weakly) increase their value relative to their current state. We define job offer acceptance sets as job productivity-separation rate combinations $(\widetilde{z}, \widetilde{\delta})$ that a worker with assets a accepts for two cases: $A_U(a)$ for the unemployment-to-employment (UE) margin, i.e., a currently unemployed job seeker presented with a job offer, and $A(a, z, \delta)$ for the EE margin, i.e., a currently employed worker with (z, δ) considering an outside offer:

$$A_U(a) = \left\{(\widetilde{z}, \widetilde{\delta}) : W(a, \widetilde{z}, \widetilde{\delta}) \geq U(a)\right\} \quad \text{and} \quad A(a, z, \delta) = \left\{(\widetilde{z}, \widetilde{\delta}) : W(a, \widetilde{z}, \widetilde{\delta}) \geq W(a, z, \delta)\right\}.$$

Unemployment value. The value of unemployment for a worker with assets a solves the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\rho U(a) = \max_c \; u(c) + U_a(a)\dot{a} + \lambda(\theta) \int_{A_U(a)} \left[W\left(a, \widetilde{z}, \widetilde{\delta}\right) - U(a)\right] dF(\widetilde{z}, \widetilde{\delta}) \tag{3}$$

subject to the borrowing constraint $a \geq \underline{a}$, where U_a denotes the derivative of $U(a)$ and $\dot{a}$ follows Equation (2) with $y = b$. The first two terms capture the worker's consumption-saving problem. The last term is the value change from a job that the worker accepts.[17]

Employment value. The value of employment for a worker with assets a in a job with characteristics (z, δ) who does not wish to quit to unemployment is:

$$\begin{aligned} \rho W(a, z, \delta) = \max_c \; & u(c) + W_a(a, z, \delta)\dot{a} + \delta \left[U(a) - W(a, z, \delta)\right] \\ & + s_e \lambda(\theta) \int_{A(a,z,\delta)} \left[W(a, \widetilde{z}, \widetilde{\delta}) - W(a, z, \delta)\right] dF(\widetilde{z}, \widetilde{\delta}) \end{aligned} \tag{4}$$

subject to the borrowing constraint $a \geq \underline{a}$. Here, the first two terms again capture the consumption-saving problem, where $W_a(a, z, \delta)$ is the partial derivative of the value function with respect to a and $\dot{a}$ follows Equation (2) with $y = w(z)$. The third term captures exogenous job destruction at rate δ, and the last term captures on-the-job search.[18] Clearly, $W(a, z, \delta)$ is increasing in assets and productivity, but decreasing in unemployment risk.

[16] The borrowing constraint implies that there are parameter restrictions needed to ensure that consumption remains positive: UI benefits and wages must be sufficiently large, $\tau y^{1-\Upsilon} > -r\underline{a}$ for $y \in \{b, w\}$. Since r is small and $\underline{a}$ is ultimately calibrated to equal quarterly earnings, this constraint is very loose.

[17] Implicit in this notation is that a firm never chooses not to accept a match that a worker would accept. This result follows from the wage setting we assume, as discussed further below.

[18] Furthermore, the worker has the option to quit to unemployment at any time. Since such jobs do not prevail in steady state, we omit this option for notational brevity. However, we account for unemployment quits in our quantitative exercises that involve aggregate shocks, where the option might become binding.

Indifference curves: the wage-job safety tradeoff. We now derive ICs between the wage and job safety. These ICs formalize our notion of caution in job mobility, and moreover permit us to bridge the empirical and model workers' tradeoffs.

Let $c_U^*(a)$ denote optimal consumption for an unemployed worker with assets a and $c^*(a, z, \delta)$ denote the consumption of an employed worker with job (z, δ). Further, let the Hamiltonians at these optimal consumption levels be denoted by $H_U(a) = u(c_U^*(a)) + U_a(a)\dot{a}$ and $H(a, z, \delta) = u(c^*(a, z, \delta)) + W_a(a, z, \delta)\dot{a}$, respectively, where $\dot{a}$ follows Equation (2) given the optimal consumption choice in each Hamiltonian.

For *unemployed* workers, reservation jobs satisfy the following indifference condition:

$$H(a, \widetilde{z}, \widetilde{\delta}) - H_U(a) = (1 - s_e)\lambda(\theta) \int_{A_U(a)} \left[W\left(a, \widehat{z}, \widehat{\delta}\right) - U(a) \right] dF(\widehat{z}, \widehat{\delta}). \tag{5}$$

For any unemployed worker with assets a, Equation (5) implicitly defines the minimum productivity $\widetilde{z}$ required for her to accept a job with risk $\widetilde{\delta}$, denoted by $z^u(a, \widetilde{\delta})$. The term on the right-hand side captures the change in the value of search when a worker moves from unemployment to employment. When $s_e < 1$, the worker only moves to jobs that sufficiently improve consumption and asset accumulation relative to unemployment. When $s_e > 1$, the worker is willing to accept a job with a lower Hamiltonian as she is compensated with a higher on-the-job offer arrival rate. For the reservation job, where $W = U$, the term that involves the value loss upon exogenous job separation in Equation (4) drops out. Therefore, unemployed workers do not directly care about the $\widetilde{\delta}$ of a reservation job offer, but care only about its impact on consumption via $H(a, \widetilde{z}, \widetilde{\delta})$.

For *employed* workers, reservation jobs satisfy the following indifference condition:

$$H(a, \widetilde{z}, \widetilde{\delta}) - H(a, z, \delta) = \left(\widetilde{\delta} - \delta\right) \times \left[W(a, z, \delta) - U(a) \right]. \tag{6}$$

Here, the on-the-job search terms cancel out as the reservation job offer yields an identical value to the current job. For any worker with assets a currently in a job with characteristics (z, δ), Equation (6) implicitly defines the minimum productivity $\widetilde{z}$ required for her to accept an offer to move to a new job with risk $\widetilde{\delta}$, which we denote by $z^r(a, z, \delta, \widetilde{\delta})$.

Equation (6) crystallizes the central insight. What determines the productivity-safety tradeoff in the new job ($\widetilde{z}$ vs. $\widetilde{\delta}$) is the *value gap*: the difference between current job value $W(a, z, \delta)$ and unemployment value $U(a)$. This value gap is measured in units of utility, making it hard to interpret. We instead measure the value of a job by its equivalent variation in wealth: the one-time reduction $\text{PR}(a, z, \delta)$ that leaves an employed worker indifferent between keeping the job and becoming unemployed, $W\left(a - \text{PR}(a, z, \delta), z, \delta\right) = U(a)$. Linearizing W in wealth and using the first order condition $W_a(a, z, \delta) = u'(c^*(a, z, \delta))$, we obtain the first-order approximation $\text{PR}(a, z, \delta) \approx \left[W(a, z, \delta) - U(a)\right] / u'(c^*(a, z, \delta))$, i.e., the value gap priced by the worker's marginal utility of consumption. With a slight abuse of notation—treating this approximation as an equality—we

define the worker's *price of risk* as:

$$\text{PR}(a, z, \delta) = \frac{W(a, z, \delta) - U(a)}{u'(c^*(a, z, \delta))}. \tag{7}$$

Since $W(a, z, \delta)$ is increasing in productivity and decreasing in risk, ceteris paribus, workers in higher-wage and lower-risk jobs have a higher price of risk. Because a higher separation rate exposes the worker to losing the entire value gap at a faster rate, the price of risk is naturally related to the slope of the wage-job safety IC. Our survey evidence in the previous section documents exactly this heterogeneity in the ICs, which we pin down next.

The IC slopes can be formally derived by differentiating Equation (6) with respect to $\widetilde{\delta}$—thereby deriving a marginal rate of substitution. To better connect ICs to the data, we express them in wage terms instead of productivity. Let $\widehat{w}(z) = \tau w(z)^{1-\Upsilon}$ (and similarly $\widehat{w}^r(a, z, \delta, \widetilde{\delta}) = \tau w(z^r(a, z, \delta, \widetilde{\delta}))^{1-\Upsilon}$) denote (reservation) wages net of taxes. Using the definition of $H(a, \widetilde{z}, \widetilde{\delta})$, the slope of the net wage-risk IC around $\widetilde{\delta}$ is:[19]

$$\frac{\partial \widehat{w}^r(a, z, \delta, \widetilde{\delta})}{\partial \widetilde{\delta}} = \frac{W(a, z, \delta) - U(a) - \frac{\partial^2 W(a, \widetilde{z}, \widetilde{\delta})}{\partial \delta \partial a} \dot{a}}{u'(c^*(a, \widetilde{z}, \widetilde{\delta})) + \frac{\partial^2 W(a, \widetilde{z}, \widetilde{\delta})}{\partial \widehat{w} \partial a} \dot{a}}. \tag{8}$$

Evaluated at the worker's own job—$\widetilde{\delta} = \delta$ and $\widetilde{z} = z$, where the reservation job coincides with the current job—Equation (8) gives us what we call the worker's *local IC slope*. We see immediately that the local IC slope is closely related to the price of risk. Indeed, it equals $\text{PR}(a, z, \delta)$ *exactly* for workers who have remained in their current job long enough to reach their job-specific optimal asset position and so have zero asset drift, $\dot{a} = 0$. See also Appendix A.2 for our elaboration on the hand-to-mouth case.

The price of risk as a near-sufficient statistic for the IC slope. In general, the IC slope in Equation (8) departs from the price of risk, $\text{PR}(a, z, \delta)$, in two ways, and we preview here that in the calibrated model, these adjustments are small—so that the PR term introduced in Equation (7) serves as a near-sufficient statistic for the precise IC slope defined in Equation (8). First, nonzero asset drift adds two correction terms since Equation (6) shows that the worker is compensated for the risk differential through the whole Hamiltonian. The asset drift terms capture how a change in risk can be compensated by an expected change in future value rather than just a change in current consumption. Quantitatively this departure is small in our baseline model (see end of Section 4.3), as detailed in Appendix C.2 and shown in Appendix Figure C.1 Panel (a). The figure plots this local IC slope directly against the price of risk. The scatter plot is tightly concentrated around the 45-degree line with a slope of 0.98 and an R^2 of 0.995. Second, moving *along* the IC induces changes in all variables which depend on the risk in the new job, $\widetilde{\delta}$. This gives the IC its

[19] In the quantitative analysis we focus on gross wages, consistent with our survey questionnaire. To ease notation, it is taken as implied in this expression that $\widetilde{z} = z^r(a, z, \delta, \widetilde{\delta})$, and $\dot{a}$ is evaluated at the optimal policy at state $(a, \widetilde{z}, \widetilde{\delta})$.

curvature, which we discuss further in the hand-to-mouth and risk-neutral cases studied below. This departure is quantitatively modest as well. Appendix Figure C.1 Panel (b) replaces the local derivative with the slope estimated across the five risk levels we elicit in the survey, which span 2 to 50 percent per year. The linear regression coefficient of this concept with the PR term is 0.95 with an R^2 of 0.996. Panel (c) makes the same comparison in the units of the survey regression itself, to which we return in Section 4.3—which also details the underlying calibration choices.

Market incompleteness and job search. To understand how workers trade off wages and job safety, it is useful to consider two extreme cases of market incompleteness, namely hand-to-mouth (HtM) workers (corresponding to a lack of any insurance) and risk-neutral (RN) workers (capturing perfect consumption insurance).[20] For ease of exposition, we set income taxes to zero ($\tau = 1, \Upsilon = 0$, so that $\widehat{w}(z) = w(z)$) and focus directly on reservation wages rather than productivities. For HtM workers Equation (5) becomes

$$u(w^u) - u(b) = (1 - s_e)\lambda(\theta) \int_{A_U} \left[W_{\text{HtM}}\left(\widehat{z}, \widehat{\delta}\right) - U_{\text{HtM}} \right] dF(\widehat{z}, \widehat{\delta}), \tag{9}$$

since value functions and job acceptance sets no longer depend on assets. Given that the option value of search on the right-hand side is a scalar, this condition defines a unique reservation wage for unemployed workers w^u independent of the unemployment risk $\widehat{\delta}$ of a specific job offer the worker receives. Similarly, with HtM workers, Equation (6) defines a reservation wage for employed workers $w^r(z, \delta, \widetilde{\delta})$ that satisfies

$$u(w^r(z, \delta, \widetilde{\delta})) - u(w(z)) = \left(\widetilde{\delta} - \delta\right) \left[W_{\text{HtM}}(z, \delta) - U_{\text{HtM}} \right]. \tag{10}$$

The perfect consumption insurance case corresponds to Equations (9)–(10) under linear utility, $u(c) = c$.[21] In this case, the EE reservation wage is analogous to Pinheiro and Visschers (2015):

$$w^r(z, \delta, \widetilde{\delta}) - w(z) = \left(\widetilde{\delta} - \delta\right) \left[W_{\text{RN}}(z, \delta) - U_{\text{RN}} \right]. \tag{11}$$

Similarly, the wage-safety indifference curve slopes depend on the degree of self-insurance. Equation (8) in the HtM case becomes $\partial w^r(z, \delta, \widetilde{\delta})/\partial \widetilde{\delta} = (W_{\text{HtM}}(z, \delta) - U_{\text{HtM}})/u'(w^r(z, \delta, \widetilde{\delta}))$ and under RN it is simply $\partial w^r(z, \delta, \widetilde{\delta})/\partial \widetilde{\delta} = W_{\text{RN}}(z, \delta) - U_{\text{RN}}$. Therefore, for these models the lack of asset drift terms means that the local IC slope is always exactly equal to their model-specific price of risk. It also directly follows that, for a given (z, δ), the indifference curves are linear in $(w^r, \widetilde{\delta})$ space under risk neutrality and are strictly convex when workers are hand-to-mouth. Intuitively,

[20] This exercise is important for broader theoretical reasons. Jarosch (2023) shows that a decentralized equilibrium of a random search model with risk-neutral workers and jobs heterogeneous along productivity and EU risk is socially efficient only in the case of workers capturing the full surplus from the job, as otherwise the EE margin is inefficient. In Appendix A we show that risk aversion induces inefficiency along the UE and EE margins even under this assumption.

[21] The perfect insurance case can be rationalized by assuming workers can sign state-contingent insurance contracts with perfectly competitive, risk-neutral insurers who share their discount rate. Those insurers, in turn, get to observe and dictate job acceptance decisions to workers.

the partial self-insurance baseline we consider is in between these two extremes, and we use our quantitative model to assess the cost of limits to self-insurance for job mobility. We can use these EE reservation wages to build model counterparts to the individual indifference curves we estimate empirically. Section 5.1 quantitatively compares the indifference curves from the baseline model to those implied by the RN and HtM variants of the model. There, we show that our baseline model's average indifference curve slope is consistent with the empirical estimate in Section 2.2 and lies between these extreme cases.

3.3 Firm values and wage determination

We introduce additional notation for workers' job acceptance decisions to define firm value functions.[22] Let $\mathbf{1}^e(\widetilde{z}, \widetilde{\delta}|a, z, \delta)$ denote whether an offer $\left(\widetilde{z}, \widetilde{\delta}\right)$ is accepted by an employed worker with (a, z, δ). Let $\mu^e(a, z, \delta)$ denote the fraction of those job offers accepted. Then:

$$\mathbf{1}^e(\widetilde{z}, \widetilde{\delta}|a, z, \delta) = \begin{cases} 1 \text{ if } W\left(a, \widetilde{z}, \widetilde{\delta}\right) \geq W(a, z, \delta) \\ 0 \text{ otherwise} \end{cases} \quad \text{and} \quad \mu^e(a, z, \delta) = \int_{\widetilde{z},\widetilde{\delta}} \mathbf{1}^e(\widetilde{z}, \widetilde{\delta}|a, z, \delta) dF(\widetilde{z}, \widetilde{\delta}). \quad (12)$$

Similarly, let $\mathbf{1}^u(\widetilde{z}, \widetilde{\delta}|a)$ indicate whether an unemployed worker with assets a accepts an offer $\left(\widetilde{z}, \widetilde{\delta}\right)$ and let $\mu^u(a)$ denote the fraction of those offers accepted:

$$\mathbf{1}^u(\widetilde{z}, \widetilde{\delta}|a) = \begin{cases} 1 \text{ if } W\left(a, \widetilde{z}, \widetilde{\delta}\right) \geq U(a) \\ 0 \text{ otherwise} \end{cases} \quad \text{and } \mu^u(a) = \int_{\widetilde{z},\widetilde{\delta}} \mathbf{1}^u(\widetilde{z}, \widetilde{\delta}|a) dF(\widetilde{z}, \widetilde{\delta}). \quad (13)$$

Production value. We assume that asset markets are segmented so that the risk-neutral entrepreneurs, who own the firms, finance vacancies using only their own equity. Profits are therefore discounted at rate ρ_f. The value of a firm employing a worker with assets a in a match with productivity z and unemployment risk δ satisfies the following HJB equation:

$$\rho_f J(a, z, \delta) = z - w(z) + J_a(a, z, \delta)\dot{a} - (\delta + s_e \lambda(\theta)\mu^e(a, z, \delta))J(a, z, \delta), \quad (14)$$

where $J_a(a, z, \delta)$ is the partial derivative of the value function with respect to assets and asset drift $\dot{a}$ follows Equation (2) according to the worker's optimal consumption choice. The firm keeps the residual output after paying wages. The match terminates either exogenously at rate δ or because the worker quits to another job at rate $s_e \lambda(\theta)\mu^e$. The firm value depends on assets through their effect on worker decisions.[23] Here, we assume a free-entry condition, discussed below, which pushes the firm's value of the vacant job to zero.

[22] In the numerical solution we replace these binary acceptance indicators by smoothed probabilities, which renders the HJBs differentiable in model parameters. Appendix C.3 derives the smoothed equations and describes how the smoothing scale is calibrated automatically from the offer-grid coarseness.

[23] A firm with unconstrained access to borrowing would discount the future at the risk-free interest rate r instead of ρ_f. We have experimented with both options and found very similar results.

Free entry. Due to random matching, a vacant job meets a given type of worker in proportion to that type's share of total search effort. Therefore, conditional on contact, the probability of meeting an unemployed worker with assets a is $g^u(a)/S$, and the probability of meeting an employed worker with characteristics (a, z, δ) is $s_e g^e(a, z, \delta)/S$. Free entry equates the cost and expected benefit of vacancy posting, accounting for worker acceptance decisions:

$$\kappa = q(\theta)\frac{1}{S}\int_{\widetilde{z},\widetilde{\delta}}\left[\int_a J(a,\widetilde{z},\widetilde{\delta})\mathbf{1}^u(\widetilde{z},\widetilde{\delta}|a)g^u(a)da + s_e\int_{a,z,\delta} J(a,\widetilde{z},\widetilde{\delta})\mathbf{1}^e(\widetilde{z},\widetilde{\delta}|a,z,\delta)g^e(a,z,\delta)dadzd\delta\right]dF(\widetilde{z},\widetilde{\delta}). \tag{15}$$

Wage determination. In Appendix C.4 we describe a game yielding the following wage rule:

$$w(z) = \psi z + (1-\psi)\chi, \tag{16}$$

where ψ captures worker bargaining power and $\chi \leq \underline{z}$ is a disagreement payoff.

3.4 The asset market, government budget constraint and equilibrium

Asset market. The risk-free bond is traded by households and the government, and is in zero net supply. Letting A_g denote government savings and $A_h = \int a g^e(a,z,\delta)dadzd\delta + \int a g^u(a)da$ denote private household savings, market clearing requires that $A_h + A_g = 0$.

Government budget constraint. The government is subject to a budget constraint where tax revenues and interest income from savings cover government spending and UI outlays:

$$G + b\cdot u = \int \left(w(z) - \tau w(z)^{1-\Upsilon}\right) g^e(a,z,\delta)dadzd\delta + \left(b - \tau b^{1-\Upsilon}\right)u + rA_g. \tag{17}$$

We can now define the steady-state equilibrium.

Definition (Steady-state equilibrium)**.** *A steady-state equilibrium is a collection of value functions $U(a), W(a,z,\delta), J(a,z,\delta)$, and associated job acceptance and consumption policy functions, wage function $w(z)$, ergodic worker distributions $g^u(a)$, $g^e(a,z,\delta)$, labor market tightness θ and real interest rate r such that, given exogenous model parameters:*

- *Worker values $U(a)$ and $W(a,z,\delta)$ solve Equations (3) and (4) yielding job acceptance and consumption policy functions,*
- *firm value $J(a,z,\delta)$ solves Equation (14),*
- *$w(z) = \psi z + (1-\psi)\chi$ is the outcome of the game described in Appendix C.4,*
- *labor market tightness θ satisfies the free-entry condition in Equation (15),*

- *the distributions* $(g^u(a), g^e(a, z, \delta))$ *follow the adjoint of the infinitesimal generator of the HJB equations (Achdou et al., 2022) implied by the policy functions and offer arrival rate,*
- *the government adjusts taxes, spending, or debt to balance its budget in Equation (17),*
- *the real interest rate clears the bond market so that* $A_h + A_g = 0$.

4 Quantitative analysis: the basics

We now describe our calibration strategy and validate its quantitative performance vis-à-vis key empirical targets, and most importantly, our survey evidence on employed workers' actual tradeoffs between safety and wages in job mobility. Then, in the remaining sections, we use the calibrated model to show qualitatively new and quantitatively important implications of the wage-job safety tradeoff for micro and macro job mobility and productivity dynamics. We relegate details on the solution, simulation, and calibration of the model to Appendix C. Throughout, the slopes of the indifference curves between wages and unemployment risk capture caution in job mobility; Figure 2 hence includes the model counterparts alongside the empirical, survey slopes.

4.1 Calibration

We solve and simulate the model in steady state, dividing the calibration discussion into pre-determined and estimated parameters.

Calibration strategy. There are 20 parameters in the model. We set 7 parameters externally and jointly estimate the remaining 13 to be consistent with an equal number of targets. The model period is set to one month, consistent with the frequency of observations from the CPS. Table 2 summarizes the model parameters and their values, as well as the model fit. The model provides an almost perfect match to the empirical targets.

Pre-determined parameters. We assume a constant relative risk aversion (CRRA) utility function, $u(c) = c^{1-\sigma}/(1-\sigma)$, and set $\sigma = 2$, as is standard in the literature. We specify the matching function as $M(V, S) = \alpha S^{\eta} V^{1-\eta}$, implying a worker contact rate $\lambda(\theta) = \alpha\theta^{1-\eta}$ and a firm contact rate $q(\theta) = \alpha\theta^{-\eta}$. We set the matching function elasticity to $\eta = 0.5$, as is standard. We set tax progressivity to $\Upsilon = 0.18$ following Heathcote et al. (2017), and discuss τ among the estimated parameters below. We set the firm's discount rate ρ_f to an annualized 5%, consistent with standard values from the real business cycle literature.

We set the strike payment to $\chi = 0$, resulting in $w(z) = \psi z$. Hence, wages and productivity move one for one in logs. We set $\psi = 0.9$.[24] While the precise value of the bargaining power over

[24]This generates an employer share of 10% of match output consistent with Hagedorn and Manovskii (2008) style calibration targets for quantitative macro-labor models and recent estimates in the sequential-auction bargaining literature. For instance, Jarosch (2023) estimates a model with Cahuc, Postel-Vinay, and Robin (2006) type bargaining on German data in a closely related context (i.e., match-specific separation risk and its consequences for job loss) and recovers

match surplus remains an open question, we view this calibration choice as conservative for the productivity (and productive efficiency) implications of our mechanism. This is because the resulting tight link between worker-facing wages and match productivity (i) if anything compresses the underlying match productivity distribution relative to wages and (ii) with a two-dimensional job ladder, the worker's choices more closely align with match productivity than otherwise, and hence also coincide more with the choices of a social planner valuing productive efficiency (Appendix A).

We set $\rho_{z,\delta} = 0$, i.e., the wage and risk of *offers* are uncorrelated (further details on the distributional assumptions below). This choice is justified, and turns out to be conservative, in the following check, summarized in Figure 4 Panel (a). This figure plots the model and empirical relationship between log wages and the EU risk, where the latter is at the industry-occupation-state cell level and CPS-based, repeating Figure 1 Panel (b). Both the model and the data exhibit a negatively sloped best-fit line (in the model, solely due to workers sorting into high-wage and low-risk jobs).[25] Hence, matching the empirical, flatter slope would actually require us to assume a *positive* correlation between wage and risk in the offer distribution. This would further boost the importance of the wage-risk tradeoff for EE moves at the heart of our paper, making our simplest possible assumption of uncorrelated draws conservative.[26]

Estimated parameters. We estimate the remaining 13 parameters jointly by matching an equal number of empirical moments. By way of informal identification, we organize our discussion by pairing each parameter with its most closely related empirical target moment.

We assume that match-specific productivity z and unemployment risk δ are log-normally distributed with means (μ_z, μ_δ) and standard deviations $(\sigma_z, \sigma_\delta)$ (and are uncorrelated, as described above). We choose μ_z to normalize the expected value of z for an offer to 1. The parameter σ_z governs the dispersion of *offered* wages, which we discipline by matching the ratio of the mean wage to the 5th percentile in the wage distribution. Our empirical target is 1.6, consistent with standard values in the literature (i.e., the mean-min ratio; see Hornstein et al., 2011).

To discipline EU risk heterogeneity, we choose the mean μ_δ to match the average monthly EU rate in the CPS discussed below. Then, we target the downward-sloping EU-tenure hazard profiles presented in Figure 1 Panel (a), controlling for worker characteristics, with σ_δ. In particular, we target a difference of 1.3 pp in the mean EU rate for tenures at 0–4 months relative to tenures at 40–44 months. The model explains the downward-sloping EU hazard through composition effects as shown in Figure 4 Panel (b). Riskier jobs end earlier, because they carry higher intrinsic EU risk or workers quit them sooner.

---

a worker bargaining power of 0.928, uniform across hires from unemployment and poaching, implying a substantial passthrough from productivity to wages.

[25] Specifically, regressing δ on $\log w$ in the model distribution of jobs yields a coefficient of -1.71. We then construct its empirical counterpart in the CPS, where each observation corresponds to an industry-occupation-state cell with an associated average wage and EU risk—after netting out worker observables such as education. In the data, we obtain a substantially smaller (in magnitude) regression coefficient of -0.76.

[26] However, we caveat that our coarse definition of jobs in the CPS likely masks within-cell heterogeneity in job risk. As a robustness check we have also considered $\rho_{z,\delta} = -0.2$ and found that this change does not quantitatively affect our results, and moves the wage-EU correlation further from that in the CPS data.

Table 2: Calibrated model parameters

Parameter	Description	Value	Source/Target		
Pre-determined:					
ρ_f	Firm discount rate	0.0043	5% annual		
σ	Risk aversion	2	Standard		
Υ	Tax progressivity	0.18	Heathcote et al. (2017)		
η	Matching elasticity	0.5	Standard		
ψ	Bargaining weight	0.9	Profit share of 0.1		
χ	Strike payment	0	See text		
$\rho_{z,\delta}$	Correlation EU risk and productivity	0	See text		
Estimated:				Target	Model
ρ	Discounting	0.0130	Average quarterly MPC	0.200	0.199
$\underline{a}$	Borrowing limit	-4.240	$1\times$ quarterly income (Kaplan et al., 2018)		
b	Unemployment benefit	0.566	Average replacement rate	0.400	0.400
s_e	Search efficiency on the job	1.31	Average EE rate	0.026	0.025
μ_δ	Unemployment risk mean parameter	-3.73	Average EU rate	0.013	0.013
σ_δ	Unemployment risk std parameter	0.954	EU tenure slope	0.013	0.013
μ_z	Mean of marginal density of z_0	-0.031	Normalize average z_0 draw	1.000	1.000
σ_z	Std. dev. of marginal density of z_0	0.285	Mean-Min ratio	1.600	1.600
κ	Vacancy posting cost	0.085	Unemployment rate	0.050	0.049
α	Match efficiency	0.257	Normalize labor market tightness θ	1	1
τ	Linear tax shifter	0.85	Personal income tax	0.200	0.200
A_g	Government saving	3.346	Asset market clearing $\rightarrow$ annual r	0.02	0.02
G	Wasteful gov. spending	0.252	Government budget balance (incl. debt)	0	0

Notes: This table presents the model parameters and their values. For estimated parameters, the right-most columns report the empirical target and the corresponding model moment. The unemployment rate, EU rate, and EE rate targets are 1996–2007 averages of aggregate monthly series (the EE rate from Fujita et al., 2024); the EU tenure slope is computed from our CPS sample; the MPC target follows Auclert et al. (2020), the mean-min ratio Hornstein et al. (2011), and the borrowing limit Kaplan et al. (2018).

We further target aggregate labor market moments, taking each moment as the average of the corresponding monthly series over 1996–2007, a period when the labor market was broadly in balance, with details in Appendix B.1.[27] Given an average monthly EU rate of 1.3% to discipline μ_δ, we choose the vacancy posting cost κ such that the worker offer arrival rate per unit of search efficiency, $\lambda(\theta)$, is consistent with an unemployment rate of $u = 0.05$. This target pins down market tightness θ through the free-entry condition. The implied UE rate is not separately targeted, which follows from the steady-state flow balance $UE = EU(1-u)/u$. We choose matching efficiency α to normalize labor market tightness to $\theta = 1$. Finally, we set the relative employed search efficiency, s_e, to match a monthly EE transition rate of 2.6% (computed by Fujita et al., 2024, who correct for biases induced by CPS survey design changes over time).

The value of unemployment benefits b is set to generate a replacement rate of 40% at the mean wage. We choose the amount of wasteful government spending, G, to balance the government's budget, given tax revenues, savings, and UI expenditures. We target a monthly risk-free rate of

[27] The EU-tenure gradient discussed above is the only labor market target computed from our own CPS micro sample. The aggregate moments instead come from published economy-wide series and are therefore not restricted to our micro sample of household heads aged 25 to 55.

Figure 4: EU rate by wage and tenure

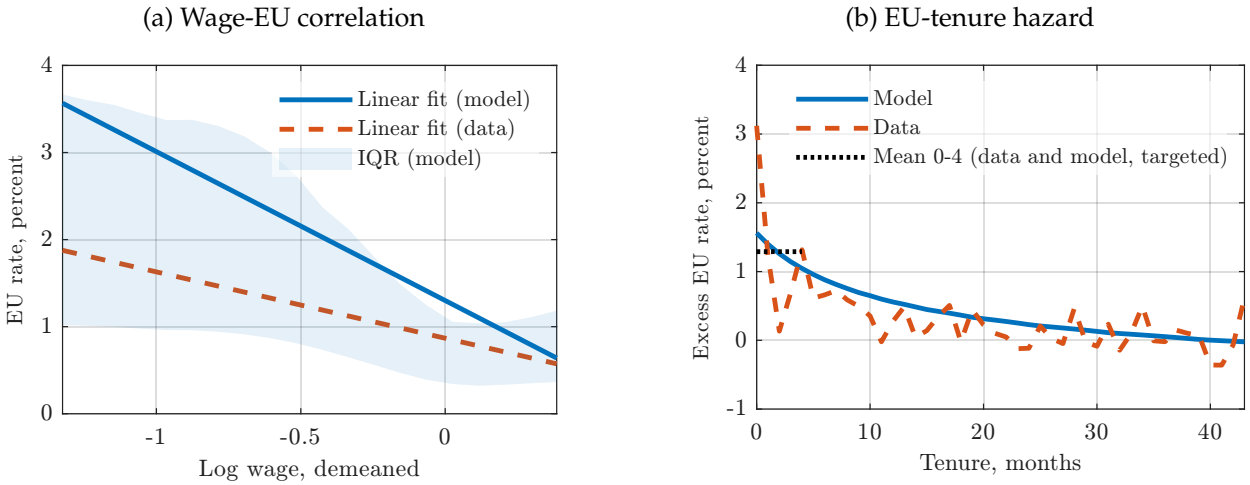

Notes: This figure plots EU rates by wage and tenure in the model and data. Panel (a) plots EU rates against log wages. The lines show linear best fits, and the shaded region denotes the interquartile range of EU rates at each wage level in the model's ergodic distribution. We demean log wages separately for the model and the data, in each case subtracting the employment-weighted mean log wage of the respective sample. Panel (b) plots the EU-tenure profiles, relative to their mean values at tenures 40–44. The horizontal dashed lines show the mean at tenure 0–4.

$r = 0.0017$ to generate a 2% annual interest rate, consistent with the return on liquid risk-free assets. This is supported by choosing government savings to clear the asset market at this interest rate given equilibrium household savings, satisfying $A_g = -A_h$.

We estimate a worker discount rate of $\rho = 0.013$, matching an average quarterly income-weighted marginal propensity to consume (MPC) of 0.2 out of a \$500 transfer (Auclert et al., 2020, 2025), where we detail its computation in Appendix C.6. This implies an annual discount rate of about 14%.[28] Despite the high discount rate, our model generates sensible consumption dynamics besides MPCs, as discussed below.[29] Following Kaplan et al. (2018), we set the borrowing limit $\underline{a}$ to correspond to average observed quarterly labor income. Finally, we choose the tax shifter τ to target an average personal income tax rate of 20%.

4.2 Evaluating the calibration and untargeted moments

Before exploring our model's ability to match the wage-job safety indifference curves elicited in our custom survey, we study the steady-state properties of the calibrated model and assess its

[28] Following Kaplan et al. (2018), we discipline the discount rate using MPCs, since single-asset models require relatively high impatience to generate empirically plausible MPCs. Nonetheless, our discount rate exceeds theirs: Kaplan et al. (2018) match an MPC of 16% with only a 9% discount rate (including mortality) in their single-asset model variant. This difference could reflect that wages are fixed during the life of a match in our model, which could reduce the amount of labor income risk—especially in safe jobs—relative to their model.

[29] Of course, the mechanism linking the discount rate to the MPC operates by pushing households to or close to the borrowing constraint. One may want to view the model's asset notion as liquid assets. To preview (see Figure 5 Panel (c) below), we are consistent with the empirical finding of about 30–40% share of constrained households (once wealthy hand-to-mouth households are accounted for, see Kaplan et al., 2014). We view the MPC as an informative moment capturing the impact of a broad notion of incomplete markets for the purposes of our paper.

Figure 5: Steady-state marginal distributions

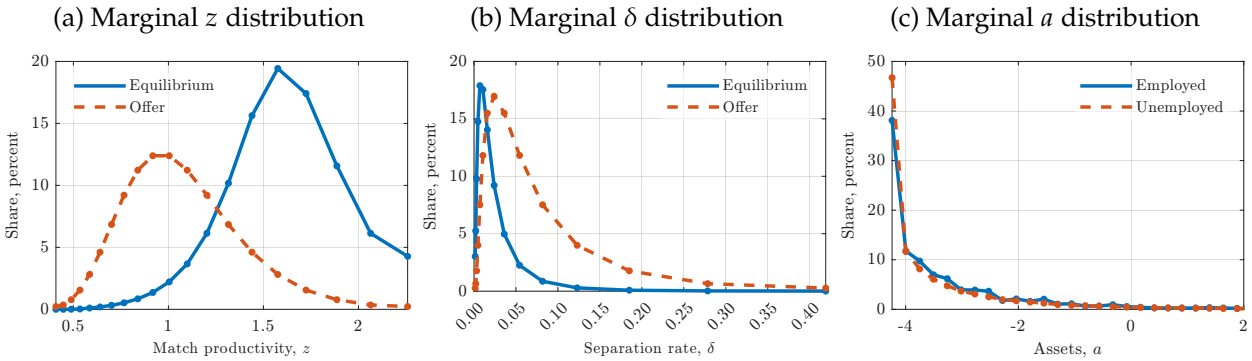

Notes: This figure plots the steady-state marginal worker distributions, obtained by integrating the joint distribution over each remaining dimension, alongside the corresponding offer distribution for each job characteristic. Each panel spans the calibrated node range of the relevant state variable. Asset distributions are normalized to sum to one separately for the employed and unemployed to aid legibility. Circles indicate the grid points on which the model is solved.

performance along other untargeted moments.

Worker distributions. Worker and firm decisions jointly determine the equilibrium distribution of workers across jobs. Figure 5 plots the resulting marginal distributions for productivity, separation risk, and assets. We reiterate that under our wage protocol and calibration, wages and productivity move one for one in logs.

Panel (a) shows the equilibrium distribution of workers over match productivity and the offer distribution for comparison. The equilibrium is shifted to the right through selectivity at the UE margin as well as EE moves toward more productive jobs. Panel (b) plots the marginal distribution of workers over their job's separation rate, again with the offer distribution. The average worker holds a substantially safer job than the average job offer. For instance, 68% of workers, but only 26% of job offers, lie in jobs with monthly EU rates of 1.3% or less. Two forces generate this gap. First, safer jobs survive longer by definition. Second, workers endogenously escape higher-risk jobs more quickly via EE moves, and reject high-risk offers, as they search for safety together with higher wages.

We also note that the joint equilibrium wage-job safety distribution is reasonable. Figure 4 Panel (a) also shows the interquartile range of EU risk at each wage level in the ergodic worker distribution. The model generates substantial dispersion in separation risk, echoing the empirical heterogeneity observed in Figure 1 Panel (b), and if anything generating a steeper negative slope than that in the data.

Finally, Figure 5 Panel (c) plots the marginal asset distributions for employed and unemployed workers. As expected, the high discount rate keeps workers close to the borrowing constraint, and the level of liquid wealth is too low relative to the data.[30] In the model, 35% and 47% of

[30]Kaplan et al. (2018) report that liquid wealth—which they argue is the appropriate notion to use in a single-asset model—is 26% of annual GDP in the US. In contrast, our model generates liquid wealth that is −18% of GDP.

employed and unemployed workers are borrowing constrained, respectively. These shares are consistent with empirical estimates of hand-to-mouth households, which range from 30% to 40% once wealthy hand-to-mouth households are accounted for (Kaplan et al., 2014). Furthermore, and perhaps more importantly, our model generates a sensible consumption drop on job loss, which we turn to next.

Labor market trajectories following job loss (untargeted). A large literature documents that workers who lose their job suffer large and persistent earnings losses (starting with Jacobson et al., 1993). This cost of job loss is important for the model to match, as it determines how much workers value job safety. We assess our model against this evidence by simulating a panel of workers, coarsening the simulated data to annual frequency, and applying a Callaway and Sant'Anna (2021) style estimator that compares the post-displacement path of workers who experience an EU event against a control group of workers who do not. Following the empirical literature, we require both groups to have at least three years of tenure at the event date, and we express results as differences relative to the pre-event years -2 to 0, where year 1 is the event date. This comparison is performed year by year, and the results are then averaged across displacement years. Details of the estimator and the underlying simulation are reported in Appendix C.7.

Figure 6 presents the results. Panel (a) plots the response of annual earnings, our direct comparison to the US-data benchmark of Davis and von Wachter (2011). Following their convention, we summarize the earnings loss by its 20-year present value at a 5% real discount rate, expressed in years of pre-displacement annual earnings: the model delivers 1.85 years, very close to the 1.71 they estimate from US administrative data, despite being untargeted.

Panels (b) and (c) implement the wage-employment-risk decomposition performed by Jarosch (2023) on German social-security data. Panel (b) splits the log-earnings response into a log-wage component and an employment component (defined as log earnings minus log wage). The loss is roughly evenly split between the two channels early on, but the employment effect recovers quickly while the wage scar remains persistent—the same qualitative pattern that Jarosch (2023) documents in Germany. Panel (c) plots the corresponding post-layoff EU-rate gap, computed as a linear probability model on the occurrence of at least one EU transition in a given year. Displaced workers face an elevated probability of subsequent job loss that decays slowly. We are reassured that the model generates the same qualitative patterns as in Jarosch (2023).

Consumption drop following job loss (untargeted). An important aspect of our mechanism is the consumption drop upon job loss. Especially under risk aversion, this response captures the immediate utility cost from unemployment and therefore governs workers' willingness to trade off wages against unemployment risk. Its severity reflects incomplete markets (as well as the generosity of unemployment insurance). Browning and Crossley (2001) estimate a 56% passthrough from income fall to consumption fall upon unemployment in Canadian data, with this passthrough measured as in Low et al. (2010) to facilitate comparison across studies. Despite being untargeted, our model generates a passthrough of 62%, close to the data, and below the 100% that would

Figure 6: Cost of job loss

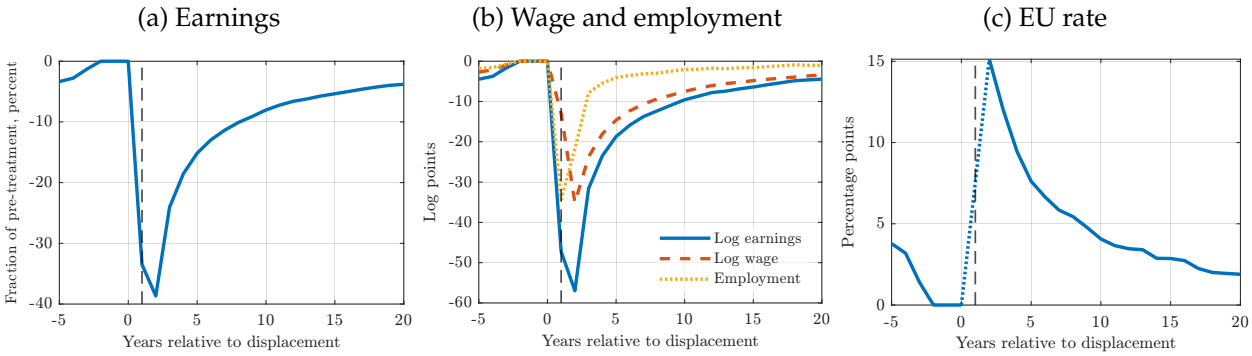

Notes: This figure plots the outcomes of treated workers, i.e., those who experience an EU separation at event time 1, relative to a control group of workers who do not yet receive treatment, pooled across reference years and relative to the pre-period window −2 to 0. See Appendix C.7 for the estimator and simulation details. Panel (a) plots the earnings level as a percent of the treated group's average pre-displacement earnings. Panel (b) plots the log-earnings, log-wage, and employment effect (log earnings − log wage) responses. Panel (c) plots the annual EU-rate response in pp. The mechanical +100 pp spike at event time 1 is omitted so that the post-displacement dynamics are visible.

be implied by a hand-to-mouth model.[31] This points to the high degree of self-insurance in the model, and gives us confidence that the cost of unemployment to workers in our model, which is a key factor shaping their wage-job safety indifference curves, is reasonable. We provide further details on this exercise in Appendix C.6.

On-the-job search efficiency (untargeted). We estimate OJS efficiency of $s_e = 1.31$. While this estimate contrasts with models featuring homogeneous separation risk, it is closer to the recent evidence from Faberman et al. (2022), who show that employed workers are 3.2 times more efficient at search than the unemployed.[32] For intuition, let A_e and A_u denote the average match acceptance rates of employed and unemployed workers, respectively. Then the job finding and switching rates are $ue = A_u\lambda(\theta)$ and $ee = s_e A_e \lambda(\theta)$. These imply an OJS efficiency of $s_e = (EE/UE)\times(A_u/A_e)$. Our calibration targets imply $EE/UE = 0.1$. We therefore estimate a high value of $s_e = 1.31$ because employed workers accept offers at only $A_e/A_u = 7.6\%$ of the rate that unemployed workers accept offers. This low relative job acceptance rate reflects heterogeneity in job safety and risk aversion, so employed workers—who occupy safe jobs in equilibrium—value safety and turn down many job offers they deem too risky.

EE moves with wage cuts (untargeted). In our model, workers climb up both the productivity and job safety ladders, but may trade off one dimension against the other when doing so moves

[31] We note that the passthrough exceeds our targeted MPC because of the persistence of the income shock and the associated wealth effects.

[32] Recent estimates in homogeneous job destruction models with exogenous search effort are Elsby and Gottfries (2022) ($s_e = 0.148$), Mercan and Schoefer (2020) ($s_e = 0.056$), and Carrillo-Tudela et al. (2025) ($s_e = 0.103$). In these models, one needs $s_e < 1$ to match the EE rate that is an order of magnitude lower than the UE rate, as workers accept all higher-paying jobs.

Figure 7: Indifference curves and the price of risk in the model and data

(a) Average within-worker IC

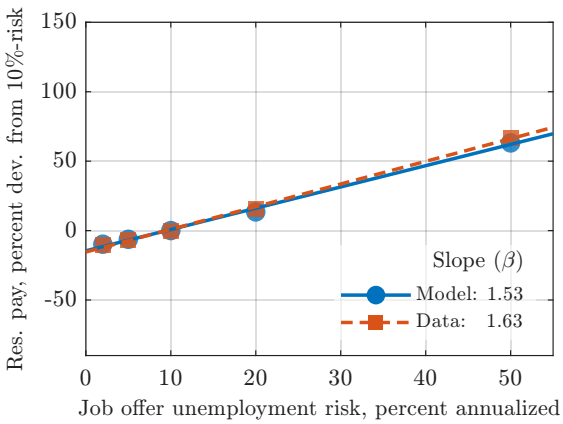

(b) ICs under model variants



Notes: This figure plots the average ICs in the model. Panel (a) plots the IC from the regression specification in Equation (1) with worker fixed effects, plotted as the solid lines, with the dots giving the equivalent non-parametric binned scatter both in the model and survey data. Each reservation wage is expressed relative to that for a job with 10% annual EU risk, so both curves pass through zero at $\widetilde{\delta}_j = 10\%$. For the model, the underlying monthly separation rates are annualized by continuous compounding. Panel (b) additionally plots the model ICs under risk neutrality, holding all other parameters fixed, and for hand-to-mouth workers with risk aversion at the baseline value, as described in Section 5.1.

them to a higher indifference curve. In equilibrium, 7% of EE transitions involve a wage cut, moving to a job with lower separation risk in return. Empirically, 34% of EE moves feature a wage cut (Tjaden and Wellschmied, 2014), which standard models rationalize by introducing a so-called *godfather shock* that may force workers to switch to a lower-paying job, or by features that backload wages (Cahuc et al., 2006). We provide a novel motive behind EE transitions with wage cuts: higher job safety accounts for about one-fifth of EE moves with wage cuts observed in the data.

4.3 The core mechanism: EE wage-safety indifference curves in the model and data

A key question for our mechanism is how workers trade off wages against unemployment risk when making job-to-job switches in reality. We elicit these preferences empirically by estimating wage-safety ICs from our custom survey in Section 2.2 and derive their theoretical counterparts in Section 3.2. We now turn to studying whether the ICs from our model, despite being untargeted, provide a good quantitative fit to the data. Figure 2 recaps the IC slopes for most of our model and survey-based variants.

IC slopes in the model vs. data. We replicate the empirical analysis of Section 2.2 using model simulated data. For brevity and with a slight abuse of notation, we express worker value functions in terms of wage w rather than productivity z, $W(a, w, \delta)$.[33] For a worker with state (a, w, δ), we

[33]Since $w(z) = \psi z$, the mapping from wage to productivity is trivial, which allows us to work with reservation wages rather than reservation productivities.

compute the reservation wage $w^r(a, w, \delta, \widetilde{\delta})$ that makes her indifferent between her current job and an offer with separation rate $\widetilde{\delta}$, defined by

$$W(a, w, \delta) = W(a, w^r(a, w, \delta, \widetilde{\delta}), \widetilde{\delta}). \tag{18}$$

For each (a, w, δ), this yields one reservation wage per $\widetilde{\delta}$. The resulting simulated data structure is therefore analogous to the indifference curve elicitation question in our survey. We restrict attention to $\widetilde{\delta}$ values spanning the range of hypothetical risks in the survey, and weight observations by the ergodic distribution of employed workers. Whenever we map model-based ICs to their survey counterparts below, we first annualize monthly separation rates by continuous compounding, $\delta^Y = 1 - e^{-12\delta}$, and then construct percentage wage changes (rather than using wage levels as in the comparative statics of Equation (8)).[34]

We first compute worker-specific indifference curve slopes, β_i, by estimating the linear specification in Equation (1) for each worker state (so that the subscript i in α_i and β_i is captured by the state variables). In the model, IC slopes depend on assets, wage, and separation risk, generating a wide range of slopes across employed workers. We compare their full distribution to the survey in Appendix C.8, where the model reproduces the mean and interquartile range, while the survey slopes are considerably more dispersed.

The model generates an *average* IC slope close to the data. Averaging β_i across workers in our survey yields a slope of 1.63, compared to 1.53 in the model. Hence, in both the model and the empirical data, employed workers require approximately a 1.5–1.6% wage increase to compensate for a 1 pp increase in job-loss risk across a 12-month horizon.

We further corroborate this result regarding the average IC slope in Figure 7 Panel (a), where we replicate the regression in Equation (1) in the model simulated data. Specifically, for a worker i considering a move to a job with (annualized) risk $\widetilde{\delta}_j$, we regress the reservation pay ratio on the risk in the new job, including worker fixed effects. The estimated coefficient β captures the average IC slope, which is equal to 1.63 in the empirical data, as reported in Table 1. Repeating the analysis in the model yields a slope of 1.53, and the figure shows that the implied ICs are visually indistinguishable across the range of unemployment risk changes in the sample. We additionally visualize ICs non-parametrically using binned scatter plots, shown by the dots, to investigate the fit of the model and data across the whole range of possible risk changes. We partition the risk changes into bins and compute average risk changes and reservation wage ratios within each bin after removing worker fixed effects.[35] The model matches the empirical data well also regarding the shape of the average IC, as the average IC is close to linear with respect to risk changes in both the model and the data.

[34] We follow the empirical IC construction by normalizing the ICs at the reservation wage for a separation risk of 10%. We find essentially identical results if we instead normalize by each individual's current job's separation risk.

[35] As we remove worker fixed effects, we study the slope rather than the level of these curves. We express each reservation wage relative to that for a job with 10% annual EU risk, so the curves pass through zero at that risk level.

IC slope heterogeneity. Next, we investigate how IC slopes vary with worker and job characteristics, comparing the model to the data. Empirically, Table 1 Column (3) shows that workers in safe jobs have an IC slope of 2.11 while those in risky jobs have a slope of 0.96. Doing the same in the model—defining safe jobs as those with below-median EU risk—we find that workers in safe jobs have an IC slope of 2.01 and those in risky jobs have a slope of 1.3. Thus the model similarly generates steeper ICs for workers in safe jobs.

The model explains this heterogeneity through differences in job values. As shown in Equations (6) and (8), employed workers with a higher value relative to unemployment, i.e., a higher $W(a, w, \delta) - U(a)$, have steeper ICs. Since safer jobs have higher value, i.e., $W_\delta(a, w, \delta) < 0$, workers in such jobs demand a larger wage premium for increases in risk. The model also qualitatively replicates the finding that workers in high-wage jobs have steeper ICs, again because high-wage jobs have higher value, i.e., $W_w(a, w, \delta) > 0$. Table 1 Column (2) shows that, in the data, workers in above-median wage jobs have an average IC slope of 1.91 while those in below-median wage jobs have a slope of only 1.33. The equivalent numbers in the model are 2.02 and 0.84.[36]

Finally, the model matches the empirical pattern that workers with higher subjective consumption drops upon job loss have steeper indifference curves. Empirically, Table 1 Column (4) shows that workers reporting that their consumption would stay the same or increase following job loss have an average IC slope of 0.98, while those reporting their consumption would fall by a little or a lot have slopes of 1.47 and 1.72, respectively. In the model, workers with below-median consumption declines upon unemployment have an average IC slope of 1.21, compared with 1.86 for those with above-median declines. While these magnitudes are somewhat larger in the model, the qualitative patterns are consistent with the data.[37]

The price of risk as a near-sufficient statistic for the IC slope: extension to the global perspective. We argue in Section 3.2 that PR in Equation (7) is a near-sufficient statistic for an employed worker's IC slope locally even when accounting for asset drift (which the PR expression suppresses). In Appendix C.2, we additionally demonstrate that this result also holds "globally." Specifically, we regress each worker's IC slope—constructed exactly analogously to the regression-based IC estimation across the wide range of job offers (rather than locally)—on her PR (adjusted in gross and percent terms, as in Equation (1)), all in the ergodic distribution of employed workers. The estimated regression coefficient is 0.99 with an R^2 of 0.995 as shown in Appendix Figure C.1 Panel (c).

[36] Interestingly, the model finds a greater sensitivity of IC slope to wage than in the data, due to understating the IC slope of low-wage workers by around 40 percent. This suggests that workers may value wage differences less in the data than in the model, perhaps due to unmodeled forms of household insurance or non-wage amenities.

[37] We measure the consumption drop in the model as the instantaneous decline upon job loss, read directly off the steady-state consumption policy functions, and classify workers into those with above and below median declines. Exact comparison with the data is challenging since in the model essentially no workers have unchanged or increased consumption following job loss.

5 Inspecting the mechanism: incomplete markets and the wage-safety tradeoff

We now investigate how incomplete markets and the caution in job mobility they induce have quantitatively important effects on workers' career choices and aggregate labor market outcomes.

5.1 Micro effects on the w-δ tradeoff: indifference curves for (in-)complete markets

As our model indifference curves match the data well, we now turn to the role of incomplete markets in determining their slope and shape. To this end, we perform counterfactual experiments where we make workers first risk neutral, and then risk averse but hand-to-mouth. The former case captures perfect insurance (while still applying discount rate ρ) against income shocks while the latter case is the limit of no insurance.[38] In these experiments, we hold all the parameters at their baseline calibrated values and operate in partial equilibrium, i.e., we hold "prices" (r, θ) and taxes (τ, Υ) fixed, and change only the workers' utility function or ability to insure.

Figure 7 Panel (b) plots ICs for these two alternative models, alongside the baseline model's IC.[39] Each counterfactual moves the model away from the data (and the congruent baseline model) in a different direction. The orange dotted line plots the average indifference curve under risk neutrality, which reveals that ICs flatten on average under linear preferences, with the slope dropping from 1.53 to 1.14. Employed workers now demand a lower wage premium to move to riskier jobs as they care less about job safety.[40]

The purple dash-dotted line instead plots the average indifference curve when workers are risk averse but hand-to-mouth. The IC becomes steeper, with its slope rising from 1.53 to 2.65. This reflects that, absent self-insurance, workers become effectively more risk averse, and hence demand higher wage premia to move to riskier jobs. In addition, the average indifference curve with hand-to-mouth workers becomes visibly convex, as shown in the binned scatter plot, in contrast to the data and baseline model, which are approximately linear. This convexity follows directly from Section 3.2, where the HtM ICs are characterized by Equation (10). The concave utility function applied to the reservation wage, $u(w^r(w, \delta, \widetilde{\delta}))$, makes the IC convex by requiring ever larger wage increases to compensate for marginal increases in unemployment risk. Appendix A establishes these properties formally in an analytical HtM version of the model, where we show that risk aversion renders the indifference curves strictly convex and that their local slope equals the price of risk. In the baseline model this effect is dampened by self-insurance so that the average indifference curve is effectively linear, albeit steeper than under risk neutrality.

[38] The full insurance case corresponds to, for example, all workers forming an insurance union where they pool all income in order to maintain constant consumption across all states. The union is not allowed to invest in outside assets, so discounts the future at rate ρ, and makes job acceptance decisions on behalf of its members. This construction allows us to isolate the role of insurance from that of discounting.

[39] To isolate the role of preference changes, we run the average IC curve regressions using the worker distribution from the baseline model as weights, not the resulting ergodic distributions of the counterfactual models. However, the results are almost identical if we instead use the counterfactual distributions.

[40] The indifference curve is linear under risk neutrality, as in Pinheiro and Visschers (2015), which is true only for the monthly job separation rate $\widetilde{\delta}$. Figure 7 Panel (b) plots the IC for *annualized* EU rates, which introduces a mild convexity.

These counterfactual exercises demonstrate that risk aversion and imperfect insurance play an important role in shaping the wage-risk tradeoff in our model. This result also reveals that quantitatively, the existing literature, by almost exclusively assuming risk neutrality when considering risky job mobility, may substantially overstate the willingness of workers to switch jobs, and moreover miss important worker heterogeneity in risk taking capacity due to, e.g., asset positions. All this, in turn, opens the door to studying how policy interventions or structural changes impacting workers' willingness to take risks propagate into on-the-job search behavior. As a stepping stone toward these exercises, we now turn to the implications of self-insurance for aggregate outcomes.

5.2 Aggregate partial-equilibrium effects

In our calibration, workers appear, on average, well insured against unemployment risk when making EE switches: the baseline model's IC slope of 1.53 is closer to the RN model's 1.14 than it is to the HtM model's 2.65. To contextualize these slopes, we study how the extent of self-insurance shapes aggregate outcomes such as productivity, unemployment, or the EE rate. We maintain the partial equilibrium analysis and compare steady states, and we view this diagnostic exercise as a segue to the general equilibrium analysis, where we introduce concrete shifts in the environment (policy and aggregate shocks) while allowing prices to adjust.[41] Table 3 presents our results.

Eliminating the role of incomplete markets, by again moving to risk neutrality, has a significant effect on the overall EE rate, which rises by 12% relative to the baseline calibration, as workers in safe jobs make riskier EE moves. Moving to the hand-to-mouth economy causes EE moves to fall by 4.5%. By this metric, the baseline model is further from a world of risk neutrality than from one of hand-to-mouth behavior, a theme which mostly carries over to other variables in Table 3.

The table also reveals the main mechanisms whereby changing workers' risk tolerance affects the aggregate economy: when risk aversion falls, equivalently with complete markets, employed workers are willing to accept riskier job offers. As workers take riskier jobs, the overall EU rate increases, by about 13.3%.[42] This, in turn, results in a higher unemployment rate. At the same time, productivity rises because workers accept higher wage (i.e., productivity) jobs in exchange for the increase in risk. Changing risk tolerance also affects the equilibrium by changing workers' job acceptance behavior out of unemployment, as stressed in the previous literature (Herkenhoff, 2019). However, this effect is quantitatively much smaller than in the EE rate, as suggested by the smaller percentage change in the UE rate than in the EE rate.[43] This reflects our high estimated OJS efficiency s_e close to one, which makes unemployed workers' reservation wages close to the flow value of unemployment, b, and hence insensitive to changes in risk tolerance or indeed other parameters; moreover, the unemployed's acceptance decision can even be reduced to a one-

[41] Our results are similar in general equilibrium, but introduce additional channels that do not add to the discussion of the core mechanisms, which we study in Section 6.2 where we perform UI experiments in general equilibrium.

[42] Throughout, the EU rate reported in the tables and text is the *layoff* rate, i.e. involuntary separations into unemployment. We exclude the small volume of voluntary quits directly into unemployment generated by the discrete-choice smoother shocks (0.8% of all separations in the baseline). The textbook steady-state identity $u = \text{EU}/(\text{EU}+\text{UE})$ therefore holds exactly only once this quit flow is added back.

[43] This also applies to the general equilibrium effects of a UI change, which we formally decompose in Section 6.2.

Table 3: Partial equilibrium effects of varying market incompleteness/risk aversion

	Y/n	u rate	EE rate	EU rate	UE rate
Baseline calibration	1.571	0.049	0.025	0.013	0.254
Risk-neutral	1.574	0.055	0.028	0.015	0.256
% difference	0.19%	11.90%	11.98%	13.30%	0.77%
Hand-to-mouth	1.567	0.047	0.024	0.012	0.252
% difference	−0.23%	−4.65%	−4.49%	−5.29%	−0.51%

Notes: This table compares the equilibrium of the baseline model with incomplete markets (top row) to the counterfactuals where workers are risk-neutral (second row) or hand-to-mouth (fourth row). Rows three and five give the percent difference of each counterfactual from the calibration. These partial-equilibrium experiments illustrate the effects of worker decisions on aggregates, where the interest rate, labor market tightness, and taxes are fixed.

dimensional reservation-wage rule as, intuitively, for the unemployed, the price of risk is zero (see the unemployed's acceptance policies in Equations (5) and (9) and the associated discussion). As a result, we see a commensurate increase in the unemployment rate. Labor productivity increases by 0.19% in the risk-neutral/complete markets counterfactual (as workers prioritize wage gains over safety), but is lower by 0.23% with hand-to-mouth consumers. Next, we move to general equilibrium counterfactuals regarding policies and business cycles.

6 A new role of unemployment insurance in job-to-job mobility

In this section, we study what the risk-taking incentives of *employed* workers imply for UI policy. The key idea is that higher UI benefits provide insurance for job switching, limiting the downside risk of income loss in unemployment, and shape on-the-job search. As we have discussed in the introduction, the literature has largely focused on the role UI plays in productivity through the search behavior of the *unemployed*. With heterogeneity in job security, UI also determines *employed* workers' willingness to switch jobs, and the kinds of jobs they accept. The extent of market incompleteness and risk aversion further boosts this new channel's relevance. To this end, we start with a micro experiment that shows the effects on individuals' tradeoffs between wages and safety, and then move to aggregate general equilibrium effects.

6.1 Micro effects on the w-δ tradeoff: model versus data

In Section 5, by comparing risk neutrality to hand-to-mouth behavior, we show that more self-insurance induces more frequent and riskier EE moves in the model. We now show that more generous UI has a similar effect, also consistent with our survey results.

Specifically, we run a survey-based hypothetical micro experiment on how UI shapes ICs, though we do not quantitatively attempt to match this more demanding hypothetical scenario and

only note its qualitative consistency.[44] We ask respondents to report their EE reservation wages again under a hypothetical $100 increase in weekly UI benefits.[45] We then reestimate effects on the indifference curve slope within individuals in that experimental condition vs. in the baseline scenario. On average, the IC slope falls from 1.62 to 1.59 in the data.[46] In the model, we replicate this experiment by increasing b by $100 per week, and the average IC slope falls from 1.53 to 1.45.[47] In both the model and data, higher UI benefits reduce the wage premium workers demand to accept higher risk in EE moves. The model response is larger, but it is in the same direction and of comparable magnitude to the data. We do not attempt to reconcile the remaining gap, also in light of confidence intervals around our point estimate, but interpret this result as supportive of our overall mechanism and its interaction with UI specifically.

In the survey, we also elicit workers' subjective probability of making an EE move over the next year, both under the baseline and the hypothetical $100 increase in UI benefits. On average, respondents report an 18.9% probability at baseline, which rises to 22.2% with higher UI. In the model, the calibrated annual EE rate is 26%, which rises by 0.32 pp—holding the worker distribution fixed—in the higher UI counterfactual, qualitatively in the same direction but considerably smaller in magnitude. A key point is that in a model version with no EU risk heterogeneity, an increase in UI leaves the EE rate unchanged. This is because workers accept any EE offer with higher productivity, regardless of UI benefits, making the EE rate invariant to UI in this experiment.[48]

In summary, this counterfactual qualitatively confirms two findings in both the data and model: higher UI benefits make workers more willing to take riskier EE moves, and this higher risk tolerance increases the number of EE moves they are willing to make—and thus the EE rate. We next turn to studying UI in general equilibrium and for aggregate outcomes.

6.2 Aggregate general-equilibrium effects

We now ask what would happen if the level of UI benefits b generating a 40% average replacement rate in the baseline calibration were dramatically lowered. As a counterfactual, we select an 8% replacement rate. This low level is meant to approximate only basic social assistance payments, which, for our US context, would reflect, e.g., food stamps or housing benefits.[49] In doing so, we conduct a general equilibrium analysis. Namely, we adjust the tax shifter τ so that the program

[44] Appendix B.5 provides details.

[45] This corresponds to around a 20% increase in average monthly UI payments in both the data and model.

[46] This 1.62 is slightly different from the baseline estimate of 1.63 as we report the slope reestimated within the sample of respondents who answer the UI-based reelicitations.

[47] Appendix C.6 details the calculation of dollar units in the model that also applies to this experiment. We implement the change in partial equilibrium to isolate the effect of UI on worker choices. To mirror the survey—where the sample is fixed by construction—we compute the counterfactual average model slope using the baseline worker distribution.

[48] In general equilibrium, and allowing the worker distribution to adjust, the homogeneous-risk model always predicts a *decrease* in the EE rate following a benefit increase, while our baseline model can predict an *increase*. We revisit this difference in the general equilibrium analysis below.

[49] Setting $b = 0$ is not possible as it would imply zero or negative consumption for agents with negative asset holdings. The choice of 8% follows Koehne and Kuhn (2015) and captures average transfers to workers whose UI has expired. Thus, this scenario corresponds to the abolition of UI while keeping social assistance payments in place.

remains budget balanced. Market tightness θ also adjusts to respect the free entry condition.[50] As our results are insensitive to how the asset market clears, we present the baseline under the assumption that the government adjusts its savings, A_g, in order to clear the asset market at the original 2% interest rate.[51] We provide further details in Appendix C.9.

To isolate the roles of risk aversion and EU risk heterogeneity, we compare our baseline model with two auxiliary ones. The first auxiliary model is a risk-neutral version, which we have already discussed, maintaining heterogeneity in z and δ. The second one features risk-averse workers and incomplete asset markets, but assumes homogeneous EU risk. We recalibrate both these models to match appropriate subsets of our baseline calibration targets. Appendix C.10 details these recalibrations and reports the resulting parameter values. To further isolate the effect of unemployment insurance on job search decisions, we study the baseline model fixing either the UE or EE job acceptance margins at their baseline values.

Table 4 presents our results at the 8% replacement rate in percent deviations from each model's calibration with a 40% UI replacement rate. To complement this analysis, Figure 8 traces how the offer arrival rate λ and the EE, UE and EU rates respond to UI across the full range of benefit levels in all three model variants. We present the complete set of such comparative statics for the baseline model in Appendix Figure C.4.

Employed job search. This exercise reveals two findings about employed job search. First, abolishing UI lowers the EE rate in the baseline and risk-neutral models, but *raises* it in the homogeneous-δ model. This effect is at the heart of our paper: when UI dries up, workers exert more caution in their job mobility decisions. No such effect applies with homogeneous risk, where—as we formally show in our model—EE transitions occur along a one-dimensional wage ladder, with no role for unemployment values shaping on-the-job search behavior.

Second, this caution at the EE margin has broader consequences. Most immediately, EU rates change for compositional reasons. Specifically, since the missing EE moves would have gone into riskier jobs (which lower UI makes less attractive, in line with the flattening of the ICs we document in our survey under a hypothetical UI *increase* in Section 6.1), the EE responses lower EU rates in both models that feature EU-risk heterogeneity, whereas EU rates are constant by construction in the homogeneous-risk model. The "fix EE" row isolates the importance of the job ladder: when we fix the EE acceptance policy at its calibrated value, abolishing UI no longer reduces the EE rate (implying that offer arrival rates play a minor role), and the EU rate stops falling as well.[52] In other

[50] We find modest general equilibrium effects on tightness, so many outcomes resemble a partial equilibrium experiment with respect to θ. While the size of UI general equilibrium effects remains an active research area, this is consistent with the meta-analysis of Cohen and Ganong (2026), who find small effects. We note that our bargaining assumption keeps wage markdowns fixed, ruling out passthrough from UI benefits to wages. This may constitute an additional general equilibrium channel we abstract from, although evidence suggests it is muted (Jäger et al., 2020).

[51] Importantly, this is not equivalent to simply ignoring asset market clearing and fixing r. Adjusting A_g to clear the asset market impacts the government budget constraint and hence the necessary tax revenue to finance the UI changes. As a robustness check, we instead hold A_g fixed and let the interest rate adjust when changing benefits ±25% from steady state. The main results are unchanged across these two approaches.

[52] By contrast, the "fix UE" policy experiment reveals similar results, except that workers become even more cautious as

Table 4: Effects of UI reduction across model variants and acceptance margins

	Y/n	u rate	EE rate	EU rate	UE rate	λ
Baseline calibration	1.571	0.049	0.025	0.013	0.254	0.257
Percent deviation from relevant s.s.						
Baseline with lower UI	−1.29	−17.91	−12.81	−18.52	−0.62	−1.91
Baseline, fix EE policies	−0.01	−0.95	2.03	1.17	1.37	0.05
Baseline, fix UE policies	−1.28	−16.93	−13.47	−18.74	−1.33	−1.94
Risk-neutral	−0.42	−12.45	−6.81	−11.32	0.15	−0.75
Homog. risk	−0.45	−12.20	5.63	0.00	9.38	0.22

Notes: This table reports the effects of the UI reduction. Top row reports baseline moments. All remaining rows report percent differences from the calibrated steady state of the relevant model variant in a counterfactual in which gross UI benefits are reduced such that the replacement rate falls from 40% to 8%. The second row is the unconstrained GE response in the baseline model. The two "fix UE" and "fix EE" rows apply the same UI cut but additionally fix UE or EE acceptance policies to their baseline values, isolating the role of each acceptance margin in transmitting the UI change to aggregates. The bottom two rows repeat the counterfactual in re-calibrated risk-neutral and homogeneous-risk versions of the model. Appendix Table C.2 reports the recalibrated parameter values of these auxiliary models.

words, the reductions in the EE and EU rates in the full counterfactual come specifically from the response of the EE acceptance policy to lower UI. This is visible directly in Figure 8: Panels (b) to (d) overlay the offer arrival rate λ (thin dashed) on the same vertical scale as the EE, UE and EU responses, highlighting that λ moves only slightly—accounting for roughly a seventh of the EE response—with workers' job acceptance decisions accounting for the rest.

Incomplete markets amplify the caution in job mobility when UI dries up: under risk neutrality (our complete-markets proxy) the EE rate falls by 6.81%, whereas in the counterfactual in the baseline economy it falls by 12.81%—incomplete markets roughly *double* the EE response. This is primarily because the EE offer acceptance rate falls more in our model: the fall in the offer arrival rate λ is 1.16 pp larger in our model, explaining only a small share of the 6 pp larger fall in the EE rate. This is also visible in the larger fall in the EU rate in our model (−18.52%) than in the risk-neutral model (−11.32%), which is driven by workers rejecting more high-risk job offers when benefits are cut. This demonstrates how incomplete markets amplify the effect of risk on EE decisions. The homogeneous-risk model lacks both these channels. Since all higher-paying offers are accepted, the EE increase there is mechanical, driven by the unemployed becoming less picky and starting the job ladder at a lower z rung.

Unemployed job search. We now revisit the literature's more standard focus on the behavior of the unemployed and its downstream consequences. When UI is reduced, unemployed workers become less picky and accept less productive jobs, raising the UE rate—but only in the homogeneous-risk model (9.4%). In both the baseline and risk-neutral models, the acceptance rate out of unem-

this experiment amplifies the effect of lower UI because we prevent the unemployed from reoptimizing their acceptance rules, forcing them to endure the now-more-painful unemployment for longer.

Figure 8: General-equilibrium responses of labor-market outcomes to varying UI benefits

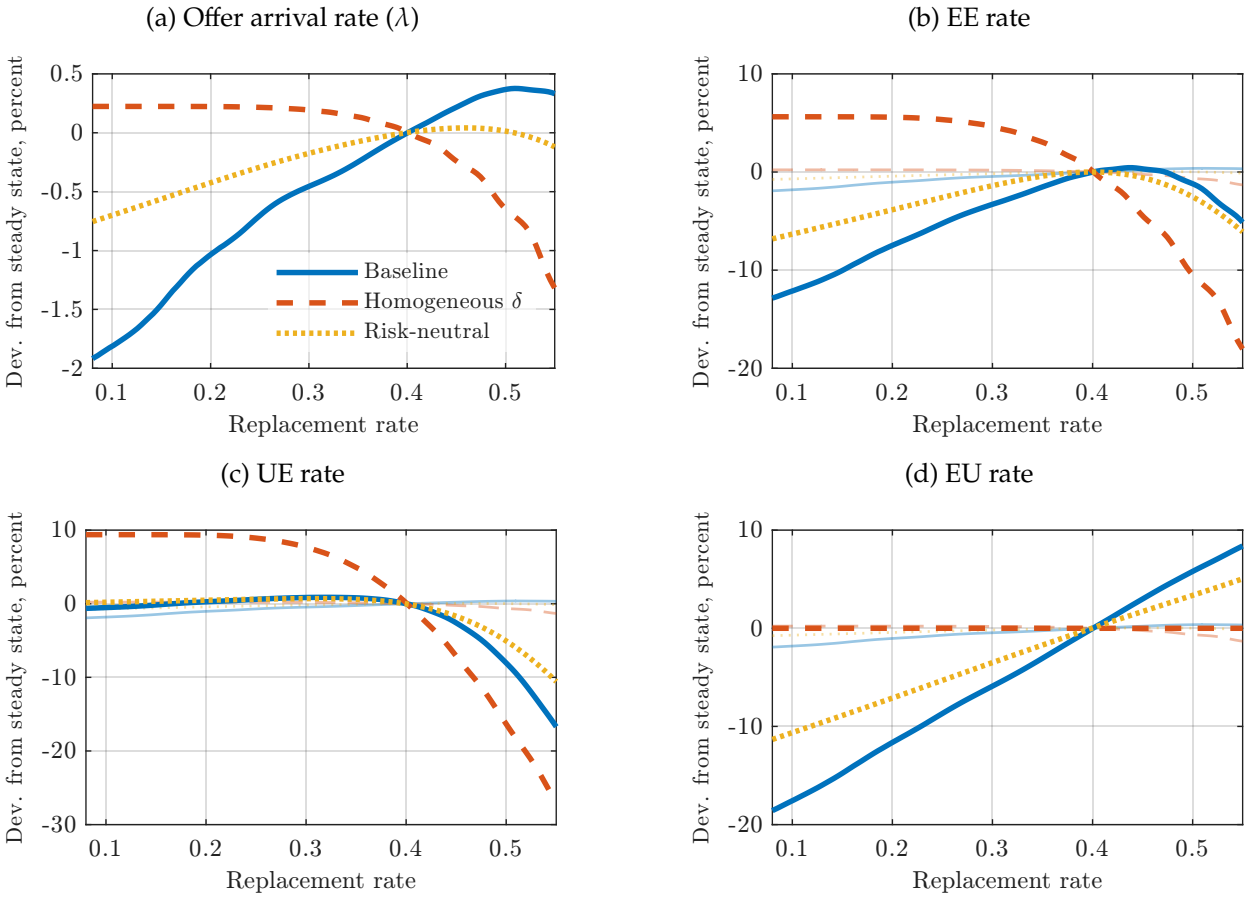

Notes: This figure plots the general-equilibrium responses of the offer arrival rate $\lambda(\theta)$, the EE, UE and EU rates to varying UI benefits b, displayed as the corresponding average replacement rate on the x-axis. Each panel shows the baseline alongside the homogeneous-risk and the risk-neutral variants. Panel (a) plots the offer arrival rate $\lambda(\theta)$ on its own scale. Panels (b) through (d) additionally overlay λ for each variant as faint dashed lines on the same vertical scale as the EE, UE and EU responses—to clarify that they are quantitatively irrelevant sources of the transition rate shifts, with acceptance probabilities dominating. Table 4 reports the corresponding values at the 8% replacement rate.

ployment is essentially flat for b below its calibrated value (Appendix Figure C.4 Panel (l)), so the UE rate barely moves (−0.6% in the baseline, +0.15% under risk neutrality).[53] Figure 8 Panel (c) shows this contrast: the UE rate rises monotonically as b falls in the homogeneous-risk model but is comparatively flat in the baseline and risk-neutral models.

[53]This insensitivity of the UE margin to b reflects a combination of the estimated offer distribution of productivity and our high estimated on-the-job search efficiency, above one in both the baseline ($s_e = 1.31$) and the risk-neutral model ($s_e = 1.27$). A high s_e pushes the reservation wage down toward the flow value of unemployment b, so that, given our estimated offer productivity distribution, almost every offer is accepted out of unemployment and a UI cut has essentially no acceptance "headroom" left to exploit. By contrast, in the homogeneous-risk model we estimate $s_e = 0.74 < 1$: the unemployed remain selective, and cutting UI makes them less picky, raising the UE rate. See Appendix Table C.2 for the recalibrated parameter values of these auxiliary models.

Labor demand. Although higher UI benefits are conventionally viewed as making hiring harder for firms, the opposite holds over the range of b we consider. Cutting benefits leads the job offer arrival rate λ to fall by 1.91% in our model, reflecting a perhaps counterintuitive drop in vacancies. This effect is absent in the homogeneous-δ model, where λ rises by 0.22%, and muted in the risk-neutral model, where λ falls by only 0.75%. In the baseline, when benefits rise, employed workers are more willing to take risk and so accept a larger fraction of outside offers (Appendix Figure C.4 Panel (k)). Vacancies meeting employed workers are therefore more likely to be filled, making it *easier* for firms to hire. This effect offsets the increased difficulty of hiring unemployed workers, who turn down more offers as UI rises, and the higher cost to firms of more frequent EE quits. In equilibrium, the additional ease of hiring employed workers dominates up to around a 50% replacement rate, so that λ traces out a hump shape in b (Figure 8 Panel (a)). The "fix EE" row of Table 4 confirms that the EE margin is key for this result: with the EE acceptance policy fixed at baseline, abolishing UI *raises* rather than lowers λ, reversing the sign of the labor-demand response. We conclude this discussion of labor demand effects by reiterating that our wage bargaining protocol deliberately precludes potential direct wage effects of UI, and overall, that the labor demand margin appears quantitatively secondary compared to worker decisions.[54]

Unemployment. Changes in the UE and EU rates translate directly into the unemployment rate. Unemployment falls by more in our baseline model (−17.9%) than in either auxiliary model (risk-neutral −12.5%, homogeneous-δ −12.2%). Relative to the homogeneous-δ model, unemployment falls by more because employed workers shy away from riskier offers and the average EU rate drops, more than offsetting the larger UE rate rise in the homogeneous-δ model. Relative to the risk-neutral model, unemployment falls by more because the EU rate falls more strongly (−18.5% vs. −11.3%): under incomplete markets, workers reject risky offers more frequently with lower UI.

Productivity. When b is dramatically reduced, productivity—output per worker, and hence measured over accepted and active matches—drops by about three times as much in the baseline counterfactual as in the homogeneous-risk or risk-neutral cases. Both market incompleteness and EU-risk heterogeneity are quantitatively important for this effect. Two channels link UI to productivity. The first is standard: with lower benefits, unemployed workers become less selective and accept lower-productivity jobs out of unemployment. Second, lower benefits make employed workers less willing to take EU risk in pursuit of more productive jobs. As discussed, the first channel is muted in our model, and it is the second that is the main source of the productivity effects of UI: fixing the EE acceptance policy at its pre-reform level removes nearly all of the productivity response (the baseline fall of 1.29% shrinks to 0.01%), while fixing the UE margin barely matters. Hence UI ameliorates some of the excessive weight that risk-averse, incompletely

[54] In the seminal contribution of Acemoglu and Shimer (1999), firms offer insurance to risk-averse workers, which leads to excessive job creation with lower wages. A positive level of UI, provided it is not too generous, *increases* output net of vacancy costs as it *discourages* job creation, which in turn provides higher wages. In our model, initial improvements in UI generosity also increase output net of recruitment costs (Appendix Figure C.4 Panel (f)) but they do so by *encouraging* job creation—the rising side of the λ hump in Figure 8 Panel (a), for which the OJS margin is key.

insured employed workers place on safety. This channel—encouraging risk taking by already *employed* workers in their job mobility decisions—is new to the longer-standing literature on the productivity effects of UI, which has focused exclusively on the selectivity of the *unemployed* (see, e.g., Acemoglu and Shimer, 1999, 2000; Marimon and Zilibotti, 1999).[55]

7 The countercyclical "climb to safety"

In our final application, we study the cyclical implications of our mechanism. Intuitively, in slack labor markets such as in recessions, lower labor market tightness increases unemployment durations and hence the price of risk: the value of unemployment U falls relative to W, widening $W - U$. Our mechanism therefore implies a countercyclical "climb to safety" among employed workers, in particular those at the borrowing constraint.

A simple model recession. To surgically shift labor market tightness, we hit our baseline model with an unanticipated and transitory shock to the vacancy posting cost that quadruples κ at impact and mean-reverts at a monthly rate of 0.025 (i.e., a half-life of roughly 28 months). Through the free-entry condition, the higher cost lowers market tightness θ and reduces the offer arrival rate $\lambda(\theta)$, generating a recession of Great-Recession magnitude. We solve the transition in a semi-partial equilibrium, where θ clears free entry along the transition path. We keep the tax shifter τ and the interest rate r at their steady-state values.

Aggregate responses. Figure 9 presents the impulse responses to the κ shock, which is engineered to have the unemployment rate roughly double by the second year (Panel (a)). The offer arrival rate λ (reflecting labor market tightness) falls by 71% on impact (Panel (b)). The UE rate follows suit, reflecting the insensitivity of the unemployed workers' reservation wages to the shock. The EE rate initially falls similarly, but then recovers more quickly.

The result that EE transitions are, if anything, less procyclical than the offer arrival rate may seem surprising: a rising price of risk should make employed workers more selective, so that—beyond the fall in λ itself—the EE transition probability conditional on an offer should fall as well. Yet the EE response is, along the transition, weakly less negative than the λ drop, implying that this selectivity margin plays a minor role in the aggregate.[56] Indeed, Panel (c) confirms that the price of risk PR = $(W - U)/(\widetilde{w}_{i,10\%}\, u'(c))$ does rise by 52% on impact. Similarly, the within-worker average IC slope increases by 32%.

[55] We explore in depth how OJS and risk heterogeneity affect the logic of productive efficiency in Appendix A, in the spirit of Acemoglu and Shimer (1999). See also Footnote 54.

[56] On impact, 99% of the fall in the EE rate is due to the collapse in λ and only 1% is due to rising selectivity; along the transition, the EE acceptance rate recovers faster than the arrival rate. Broadly, the limited role of the acceptance margin for EE fluctuations echoes the result in Mercan and Schoefer (2020), that a stable average EE acceptance rate paired with a cyclical offer arrival rate can match US business cycle fluctuations in quit and EE rates.

Figure 9: Impulse responses to a transitory vacancy-posting-cost shock

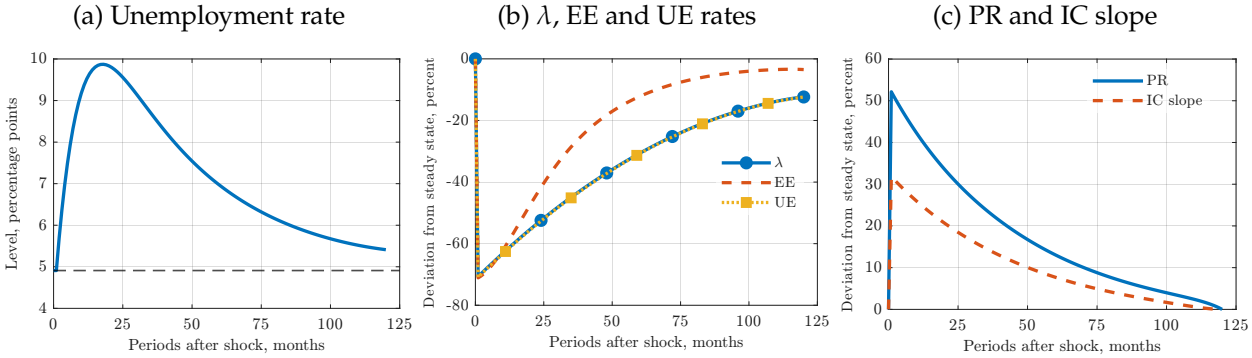

Notes: This figure plots the model impulse responses to an unanticipated and transitory positive shock to the vacancy posting cost. We solve the transition path in semi-partial equilibrium, where λ satisfies free entry, while the tax shifter τ and interest rate r are held at steady state. Panel (a) plots the unemployment rate. Panel (b) plots the offer arrival rate λ, EE rate, and UE rate. Panel (c) plots the price of risk PR = $(W - U)/(\widetilde{w}_{i,10\%}\, u'(c))$ and the regression-based average IC slope.

Climb to safety. Rather than in the number of EE moves, the higher price of risk instead manifests itself in the *composition* of EE moves. Figure 10 shows this shift directly. Specifically, for every EE switch, we record the percentage point change in (annualized) unemployment risk from the current to the new job, and the associated percent wage change. For reference, Panel (a) plots a heat map of these changes (i.e., fraction of EE moves that land in each cell) in *steady state*—without the κ shock. Almost all EE moves are to better-paying jobs (about 93%), which can be separated into riskier and safer jobs (54% and 39%, respectively).

To illustrate the impact of the recession on this composition, Panel (b) plots the *change* in these fractions, cell by cell, comparing the switches immediately after the recession to those before. The distribution of EE transitions shifts to the southwest: moves that trade pay for safety become more common, while moves that take on unemployment risk in exchange for wage gains become rarer. The average move becomes safer and the share of moves that trade lower pay for greater safety nearly doubles, from 7% to 14%. Hence, the climb to safety becomes apparent in the *composition* of mobility—but not in the aggregate EE rate. Moreover, in Appendix C.11, we also pinpoint the role of specific groups in the climb to safety and EE dynamics in recessions: safe-job, borrowing-constrained, low-asset workers, who have the most to lose and the least room to self-insure.

Countercyclical unemployment insurance. A natural policy response to this repricing of risk is countercyclical UI generosity, which stabilizes the risk-taking capacity of job switchers. To study this scenario, we repeat the recession experiment while temporarily raising the benefit b by 50 percent, mean-reverting at the same rate as the vacancy-cost shock, with the interest rate and income-tax shifter held at their steady-state values.[57] Because our wage rule insulates $\lambda(\theta)$ from

[57] Financing the UI extension through taxes would partially offset the insurance it provides, by increasing taxes on the employed during the downturn. Our approach isolates the pure insurance channel.

Figure 10: The composition of EE moves in steady state and its change during recessions

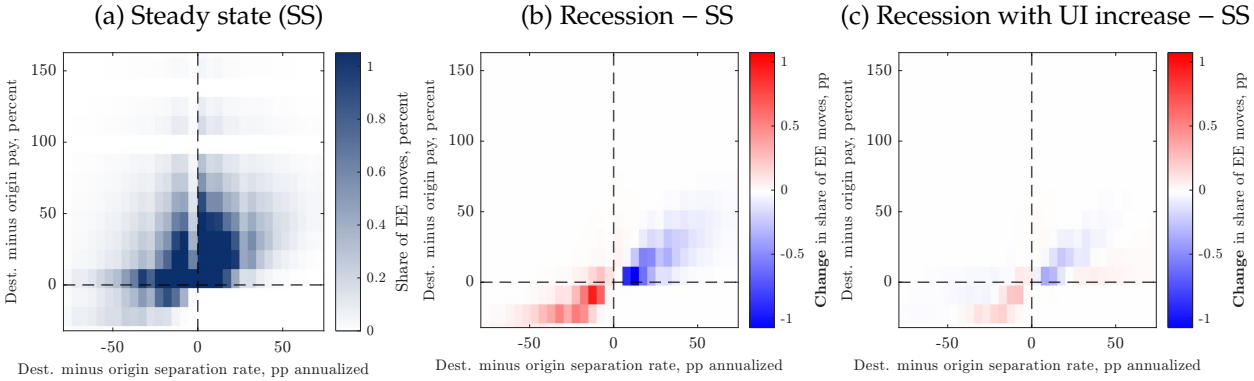

Notes: This figure plots the change in job characteristics across EE moves: the destination job's unemployment risk minus the origin's and the percent change in pay $w(z)$. Panel (a) is the steady-state density of EE moves across these cells. Panels (b) and (c) plot the *change* in this density relative to steady state at the moment immediately after the shock hits for the recession and the recession accompanied by a temporary +50% UI, respectively.

the direct effect of b, the extension leaves the decline in λ_t roughly unchanged but dampens the repricing of risk of job switching relative to the κ-shock only exercise. Figure 10 Panel (c) replicates Panel (b) (the baseline recession) but now compares changes from the steady state to the recession onset with the more generous UI regime in the latter scenario.

Panel (c) reveals that the temporary +50% UI expansion almost entirely erases the recession-induced climb to safety.[58]

Of course, the UI extension has other effects in a recession. Higher reservation values reduce job acceptance rates and so directly slow job finding of the unemployed, which in turn changes the composition of jobs from which the EE moves stem. Additionally, the UE—as well as EE—selectivity responses feed through to the offer arrival rate via the free entry condition. In fact, we find a moderate effect of these margins: upon impact λ falls by an additional 2.1% under the +50% expansion. These effects and the UI-induced voluntary EU quits lead unemployment to peak higher: about 11.7%against 9.9% in the κ-induced recession alone. This implies a new variant of the conventional tradeoff in countercyclical UI policy: stabilizing the price of risk—and thereby EE transitions—comes at the cost of a larger rise in unemployment.

8 Conclusion

We explore an intuitive and basic idea: workers approach job mobility decisions with caution, weighing wage gains against job safety. When workers are risk averse and markets are incomplete,

[58] In unreported results, the PR spike more than halves, from +52% (Figure 9 Panel (c)) to +20%, and the rise in the IC slope falls similarly, from +32% to +18%. Moreover, we find that a more modest +20% UI extension cuts the price-of-risk spike to +38% and the IC-slope rise to +26%. Finally, rather than the average change in δ from realized EE moves, another metric is simply the average δ of the destination jobs of EE moves. This metric falls from 24.9% to 22.2% on impact in the recession—reflecting a qualitatively similar but smaller effect as it does not reflect the compositional effects of the origin jobs' riskiness. The +50% expansion undoes the drop in the destination job's δ entirely (25.4%).

their caution becomes excessive: they too frequently stay put in safe but lower-productivity jobs and reject risky high-wage jobs.[59] The resulting misallocation curbs aggregate productivity and creates a new role for unemployment insurance. We formalize this channel in a general-equilibrium search model, and show that three features are jointly necessary for it to operate: on-the-job search, heterogeneity in separation risk across jobs, and incomplete markets. With notable exceptions in other emerging work (Larkin, 2026; Caratelli, 2025; Pellegrini, 2026), existing work combines any two of these features but not all three, and so misses both the qualitative mechanism and its quantitative bite. Ours is the first paper to combine all three and conduct a dedicated analysis of how these features play out in equilibrium.

Moreover, we provide new direct evidence on this tradeoff. In a custom survey of US workers fielded as part of NORC's AmeriSpeak panel, we elicit the wage-job safety indifference curves at the heart of this mechanism. On average, a 1 pp increase in annual layoff risk requires a 1.63% higher wage to leave respondents indifferent, with substantial dispersion across workers related to wage, safety, and the consumption drop upon unemployment. Although untargeted, the average indifference curve slope is reproduced very closely by our model, so that our key mechanism is empirically disciplined. The IC slopes provide a useful bridge between the model and data and capture the caution in job mobility that is the focus of the paper (summarized in Figure 2).

Overall, taking the riskiness of job-to-job mobility seriously implies that misallocation from incomplete markets operates not only through the unemployed worker's job-finding decision but also, and to a comparable degree, through the employed worker's reluctance to climb a risky job ladder. The model also uncovers new roles for standard labor-market policies. Higher unemployment benefits do more than insure the unemployed: they subsidize the acceptance of risky offers by the employed and, by raising employed workers' willingness to switch, may *encourage* rather than depress job creation. The two-dimensional risky job ladder is thus an important tool for both the positive analysis of labor reallocation and the design of social insurance—and offers a new perspective on the ongoing debate over which factors and policies may boost labor market dynamism, including "active labor market policies for the employed" (Schoefer, 2025).

Our results suggest several avenues for future work. First, with rising wealth inequality, a growing share of low-wealth households could slow the job ladder, linking our mechanism to the secular decline in US labor-market fluidity in the last few decades (see, e.g., Baksy et al., 2026)—in a sense the reverse direction of the mechanism explored by Larkin (2026).

Second, labor market policies such as imperfectly portable benefits, severance payments, and tenure-based employment protections all operate on the wage-safety margin of job switching and hence link labor market institutions with fluidity (Schoefer, 2025). Our framework offers a natural laboratory for evaluating these policies and their aggregate effects.

Third, we conduct our analysis within a model calibrated to the US context and field the survey among US workers. The US ranks among the most fluid labor markets, with short unemployment durations and high EE transition rates (see, e.g., Elsby et al., 2013; Schoefer, 2025). For that reason,

[59] See Footnote 1 regarding the productive efficiency concept (rather than broader welfare notions) we use.

our results likely underestimate the relevance of the mechanism. We view a full cross-country extension as an important object of future research, linking low UE rates with low EE rates and risk avoidance, along with downstream consequences of such depressed dynamism.[60] Indeed, in our own preliminary "European calibration," with higher unemployment rates stemming from lower UE rates, the price of risk for workers rises by about 50%—triggering substantial risk avoidance and "climb to safety" in less fluid labor markets.

References

Acemoglu, D. and R. Shimer (1999): "Efficient Unemployment Insurance," *Journal of Political Economy*, 107, 893–928.

——— (2000): "Productivity Gains from Unemployment Insurance," *European Economic Review*, 44, 1195–1224.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 89, 45–86.

Algan, Y., A. Chéron, J.-O. Hairault, and F. Langot (2003): "Wealth Effect on Labor Market Transitions," *Review of Economic Dynamics*, 6, 156–178.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," *NBER Working Paper No. 26647*.

——— (2025): "Fiscal and Monetary Policy with Heterogeneous Agents," *Annual Review of Economics*, 17, 539–562.

Baksy, A., D. Caratelli, and N. Engbom (2026): "The Long-Term Decline of the U.S. Job Ladder," *NBER Working Paper No. 34981*.

Benabou, R. (2002): "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70, 481–517.

Birinci, S., F. Karahan, Y. Mercan, and K. See (2026): "Labor Market Shocks and Monetary Policy," *Journal of Monetary Economics*, 159, 103928.

Bloemen, H. G. and E. G. F. Stancanelli (2001): "Individual Wealth, Reservation Wages, and Transitions into Employment," *Journal of Labor Economics*, 19, 400–439.

Braxton, J. C., K. Herkenhoff, and G. M. Phillips (2024): "Can the Unemployed Borrow? Implications for Public Insurance," *Journal of Political Economy*, 132, 3025–3076.

Browning, M. and T. F. Crossley (2001): "Unemployment Insurance Benefit Levels and Consumption Changes," *Journal of Public Economics*, 80, 1–23.

Cahuc, P., F. Postel-Vinay, and J.-M. Robin (2006): "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 74, 323–364.

Callaway, B. and P. H. C. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225, 200–230.

Caratelli, D. (2025): "Labor Market Recoveries across the Wealth Distribution," *Office of Financial Research Staff Discussion Paper No. 24-01*.

Carrillo-Tudela, C., A. Clymo, and M. Coles (2025): "Equilibrium Job Turnover and the Business Cycle," *CEPR Discussion Paper No. 16733*.

[60] For instance, Schoefer (2025) reports that in a representative survey, the most important reason to reject high-wage offers among German employed workers is loss of job security.

Chaumont, G. and S. Shi (2022): "Wealth Accumulation, On-the-Job Search and Inequality," *Journal of Monetary Economics*, 128, 51–71.

Chetty, R. (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 173–234.

Clymo, A., P. Denderski, and L. Harvey (2022): "Wealth, Quits and Layoffs," *Working Paper*.

Cohen, J. P. and P. Ganong (2026): "Disemployment Effects of Unemployment Insurance: A Meta-Analysis," *American Economic Review: Insights*, 8, 1–18.

Davis, S. J. and J. Haltiwanger (1999): "Gross Job Flows," *Handbook of Labor Economics*, 3, 2711–2805.

Davis, S. J. and T. von Wachter (2011): "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, Fall, 1–72.

den Haan, W. J., P. Rendahl, and M. Riegler (2018): "Unemployment (Fears) and Deflationary Spirals," *Journal of the European Economic Association*, 16, 1281–1349.

Dennis, J. M. (2019): "Technical Overview of the AmeriSpeak Panel: NORC's Probability-Based Household Panel," *Technical Report, NORC at the University of Chicago*.

Drenik, A., S. Jäger, P. Plotkin, and B. Schoefer (2023): "Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data," *The Review of Economics and Statistics*, 105, 206–216.

Eeckhout, J. and A. Sepahsalari (2024): "The Effect of Wealth on Worker Productivity," *The Review of Economic Studies*, 91, 1584–1633.

Elsby, M. W. L. and A. Gottfries (2022): "Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations," *The Review of Economic Studies*, 89, 1370–1419.

Elsby, M. W. L., B. Hobijn, and A. Şahin (2013): "Unemployment Dynamics in the OECD," *The Review of Economics and Statistics*, 95, 530–548.

Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2022): "Job Search Behavior among the Employed and Non-Employed," *Econometrica*, 90, 1743–1779.

Fujita, S., G. Moscarini, and F. Postel-Vinay (2024): "Measuring Employer-to-Employer Reallocation," *American Economic Journal: Macroeconomics*, 16, 1–51.

Griffy, B. S. (2021): "Search and the Sources of Life-Cycle Inequality," *International Economic Review*, 62, 1321–1362.

Hagedorn, M. and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98, 1692–1706.

Haltiwanger, J., H. Hyatt, and E. McEntarfer (2018): "Who Moves Up the Job Ladder?" *Journal of Labor Economics*, 36, S301–S336.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 104, 2075–2126.

——— (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

Hendren, N. (2017): "Knowledge of Future Job Loss and Implications for Unemployment Insurance," *American Economic Review*, 107, 1778–1823.

Herkenhoff, K., G. M. Phillips, and E. Cohen-Cole (2024): "How Credit Constraints Impact Job Finding Rates, Sorting, and Aggregate Output," *The Review of Economic Studies*, 91, 2832–2877.

Herkenhoff, K. F. (2019): "The Impact of Consumer Credit Access on Unemployment," *The Review of Economic Studies*, 86, 2605–2642.

Hornstein, A., P. Krusell, and G. L. Violante (2011): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, 101, 2873–2898.

Huang, J. and X. Qiu (2026): "Precautionary Mismatch," *Working Paper*.

Hubmer, J. (2018): "The Job Ladder and Its Implications for Earnings Risk," *Review of Economic Dynamics*, 29, 172–194.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83, 685–709.

Jäger, S., B. Schoefer, S. Young, and J. Zweimüller (2020): "Wages and the Value of Nonemployment," *The Quarterly Journal of Economics*, 135, 1905–1963.

Jarosch, G. (2023): "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, 91, 903–942.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Kaplan, G., G. L. Violante, and J. Weidner (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, Spring, 77–138.

Koehne, S. and M. Kuhn (2015): "Should Unemployment Insurance Be Asset Tested?" *Review of Economic Dynamics*, 18, 575–592.

Krolikowski, P. (2017): "Job Ladders and Earnings of Displaced Workers," *American Economic Journal: Macroeconomics*, 9, 1–31.

Krusell, P., T. Mukoyama, and A. Şahin (2010): "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations," *The Review of Economic Studies*, 77, 1477–1507.

Larkin, K. (2026): "Job Risk, Separation Shocks and Household Asset Allocation," *Working Paper*.

Lentz, R. and T. Tranæs (2005): "Job Search and Savings: Wealth Effects and Duration Dependence," *Journal of Labor Economics*, 23, 467–489.

Lise, J. (2013): "On-the-Job Search and Precautionary Savings," *The Review of Economic Studies*, 80, 1086–1113.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432–1467.

Marimon, R. and F. Zilibotti (1999): "Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits," *The Economic Journal*, 109, 266–291.

Mercan, Y. and B. Schoefer (2020): "Jobs and Matches: Quits, Replacement Hiring, and Vacancy Chains," *American Economic Review: Insights*, 2, 101–124.

Mui, P. and B. Schoefer (2025): "Reservation Raises: The Aggregate Labour Supply Curve at the Extensive Margin," *The Review of Economic Studies*, 92, 442–475.

Ozkan, S., J. Song, and F. Karahan (2023): "Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital," *Journal of Political Economy Macroeconomics*, 1, 506–550.

Pellegrini, E. (2026): "Wealth Inequality and Labor Mobility: The Job Trap," *Working Paper, Boston College*.

Pinheiro, R. and L. Visschers (2015): "Unemployment Risk and Wage Differentials," *Journal of Economic Theory*, 157, 397–424.

Ravn, M. O. and V. Sterk (2017): "Job Uncertainty and Deep Recessions," *Journal of Monetary Economics*, 90, 125–141.

——— (2021): "Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach," *Journal of the European Economic Association*, 19, 1162–1202.

Schoefer, B. (2025): "Eurosclerosis at 40: Labor Market Institutions, Dynamism, and European Competitiveness," *NBER Working Paper No. 33975*.

Schoefer, B. and O. Ziv (2024): "Productivity, Place, and Plants," *The Review of Economics and Statistics*, 106, 1167–1186.

Tjaden, V. and F. Wellschmied (2014): "Quantifying the Contribution of Search to Wage Inequality," *Amer-*

ican Economic Journal: Macroeconomics, 6, 134–161.

Topel, R. H. and M. P. Ward (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107, 439–479.

Online Appendix:
Cautious Careers: Job Mobility under Incomplete Markets

Alex Clymo, Piotr Denderski, Yusuf Mercan, and Benjamin Schoefer

A Analytical (hand-to-mouth) model and risk-aversion-induced productive inefficiencies **3**
A.1 Model setup . . . 3
A.2 Job acceptance decisions . . . 4
A.3 Risk aversion, policy, and productive efficiency . . . 6

B Empirical appendix **13**
B.1 Details on the publicly available datasets . . . 13
B.2 Details on our custom survey . . . 16
B.3 Further estimation results . . . 20
B.4 Robustness of baseline IC estimate . . . 21
B.5 The UI generosity treatment . . . 23

C Quantitative model appendix **26**
C.1 Numerical implementation: state space and grids . . . 26
C.2 The price of risk as a near-sufficient statistic for IC slopes . . . 26
C.3 Smoothing the discrete acceptance decisions . . . 28
C.4 Wage determination . . . 30
C.5 Tenure distributions . . . 31
C.6 Consumption dynamics in the model . . . 31
C.7 Cost of job loss: estimator and simulation . . . 32
C.8 The distribution of indifference-curve slopes in the model and the data . . . 33
C.9 UI in general equilibrium . . . 33
C.10 Calibration of the auxiliary models . . . 34
C.11 Cross-sectional anatomy of the cyclical climb to safety . . . 34

References: Online Appendix **38**

Tables

B.1 Descriptive statistics, preferred sample . . . 20
B.2 The tradeoff between pay and job safety across demographic groups . . . 22
B.3 The tradeoff between pay and job safety across specifications . . . 24
B.4 UI generosity treatment: subjective EE probability and IC slopes . . . 25
C.1 ±25% UI changes under different asset market clearing assumptions . . . 34
C.2 Calibrated parameters: baseline and auxiliary models . . . 36
C.3 Who reroutes toward safety . . . 37

Figures

A.1 Indifference curves: planner vs. market 9
B.1 Companion to Figure 1 Panel (b): raw, cell-level plot of job-specific EU rates vs. wages 15
B.2 Subjective job loss and quit probabilities 21
B.3 Distribution of individual IC slopes 21
B.4 Individual IC slopes across elicitations 22
B.5 Sensitivity of the IC slope to the winsorization cutoff 23
C.1 The price of risk as a sufficient statistic for indifference-curve slopes 29
C.2 Quarterly MPCs to $500 windfall as a function of assets 32
C.3 Distribution of individual indifference-curve slopes: model versus survey 34
C.4 Effect of varying UI benefits in general equilibrium 35
C.5 Cross-cell heterogeneity at the impact of the recession 37

A Analytical (hand-to-mouth) model and risk-aversion-induced productive inefficiencies

This appendix dissects a hand-to-mouth version of the main-text model. First, this special case delivers analytical intuitions. Second, in this special case, we formalize the notion of productive (output) efficiency that serves as our benchmark throughout the paper (see Footnote 1 in the main text). Then, we discuss—feature by feature—which ingredients of our framework distort job mobility, and to what extent UI can correct these distortions. With this second goal, we aim to provide a *job-to-job* transition analog of the results of Acemoglu and Shimer (1999, 2000), who show that for risk-averse *unemployed* job seekers, a positive level of UI restores productive efficiency. By contrast, our focus is on the role of UI in a model with on-the-job search and separation-risk heterogeneity, and we dissect a novel source of inefficiency arising from the behavior of the *employed* along the job ladder.[1]

A.1 Model setup

We present an analytical, partial equilibrium version of the model in Section 3. More specifically, we take labor market tightness as given and abstract from fiscal considerations including taxation. Time is continuous and indexed by $t \in [0, \infty)$. Workers discount the future at rate ρ. Workers are hand-to-mouth consumers and obtain consumption utility $u(c)$ with $u'(c) > 0$ and $u''(c) \leq 0$. Thus, workers consume $c = w$ when employed in a job paying wage w and in unemployment they consume $c = b = b_{HP} + b_{UI}$, consisting of home production b_{HP} and unemployment insurance (UI) benefits b_{UI}. Unlike in the main text, to highlight the direct effect of UI, we abstract from how these benefits are funded, implicitly assuming that they are financed by lump-sum taxes elsewhere in the model. Compared to the main text, we introduce home production b_{HP} as it permits the complete removal of UI under risk aversion. Otherwise, such a reform could result in an undefined utility flow for the unemployed, e.g., under CRRA preferences. b_{HP} is therefore similar to the role of social assistance system (e.g., food stamps), which the main-text UI experiment keeps in place.

We assume the same wage setting protocol as in the main body of the paper, yielding $w(z) = \psi z + (1 - \psi)\chi$, with z denoting match productivity, χ denoting a strike payment and ψ governing worker's bargaining power. In other words, match productivity uniquely determines the wage level. Thus, for notational convenience, and foreshadowing the analysis in Appendix A.3, we replace productivity z with wage w in the offer distribution and value functions.

Offers arrive at exogenous rates λ_U for the unemployed and λ_E for the employed. We define $s_e = \lambda_E/\lambda_U$ as the relative search efficiency of the employed. Offers are drawn from the joint cumulative distribution function $F(w, \delta)$ with support $\mathcal{S} = \left[\underline{w}, \overline{w}\right] \times \left[\underline{\delta}, \overline{\delta}\right] \subset \mathbb{R}^2_+$. Wage w is fixed throughout the match. Matches are destroyed at rate δ, where the worker becomes unemployed.

Value functions. The value of being employed in a job with characteristics (w, δ) is denoted by $W(w, \delta)$, and U denotes the value of unemployment. Workers accept job offers which (weakly) increase their value relative to their current state. We define offer acceptance sets as $(\widetilde{w}, \widetilde{\delta})$ combinations that a worker accepts for two cases: A_U for the unemployment-to-employment (UE) margin, and $A(w, \delta)$ for the employer-to-employer (EE) margin, analogously to the main text:

$$A_U = \left\{(\widetilde{w}, \widetilde{\delta}) \in \mathcal{S} : W(\widetilde{w}, \widetilde{\delta}) \geq U\right\} \quad \text{and} \quad A(w, \delta) = \left\{(\widetilde{w}, \widetilde{\delta}) \in \mathcal{S} : W(\widetilde{w}, \widetilde{\delta}) \geq W(w, \delta)\right\}. \tag{A.1}$$

[1] The technical details of the intermediate steps and proofs in this section will be relegated to a supplemental online appendix. Since these largely employ standard methods, we emphasize economic intuitions throughout this appendix.

Using the definition of these sets, the value of unemployment solves the following equation:

$$\rho U = u(b) + \lambda_U \int_{A_U} \Big(W(\widetilde{w}, \widetilde{\delta}) - U \Big) \, dF(\widetilde{w}, \widetilde{\delta}). \tag{A.2}$$

Similarly, the value of employment in a job with wage w and separation rate δ is given by:

$$\rho W(w, \delta) = u(w) + \lambda_E \int_{A(w,\delta)} \Big(W(\widetilde{w}, \widetilde{\delta}) - W(w, \delta) \Big) \, dF(\widetilde{w}, \widetilde{\delta}) + \delta \big[U - W(w, \delta) \big]. \tag{A.3}$$

The following defines the decentralized, i.e., market, allocation of the model.

Definition 1 (Decentralized allocation)**.** *The decentralized allocation consists of the following objects: value functions $\{U, W(w, \delta)\}$ and the distribution of workers across jobs $g_E(w, \delta)$ such that U and $W(w, \delta)$ solve Equations (A.2)–(A.3) and the sets A_U and $A(w, \delta)$ induce the worker distribution $g_E(w, \delta)$ given job arrival and destruction rates.*

A.2 Job acceptance decisions

In this section we characterize workers' optimal choices in the market allocation. To begin with, we show that a one-dimensional reservation wage is the full description of the search policy for *unemployed* workers. That is, in the hand-to-mouth case, unemployed workers—no matter their risk aversion or the UI level—do not exert any selectivity along *job security*. Hence, selectivity implications regarding job security exclusively emanate from EE transitions, with no role for the UE margin. By contrast, unemployment risk and risk aversion matter separately and interactively for EE transitions.

UE transitions. Evaluating Equation (A.3) at reservation jobs for unemployment reveals, since the $\delta \left(W(w, \delta) - U \right)$ term drops out, that the value is independent of δ:

$$\rho W \left(w, \delta \right) = u(w) + \lambda_E \int_{A_U} \Big(W(\widetilde{w}, \widetilde{\delta}) - U \Big) \, dF(\widetilde{w}, \widetilde{\delta}) \;\; \forall \; (w, \delta) \text{ s.t. } W(w, \delta) = U. \tag{A.4}$$

Moreover, using Equations (A.2) and (A.4), we can pin down a standard one-dimensional reservation wage w_r out of unemployment (rather than a two-dimensional set), solving the equation:

$$u(w_r) = u(b) + (\lambda_U - \lambda_E) \int_{A_U} \Big(W(\widetilde{w}, \widetilde{\delta}) - U \Big) \, dF(\widetilde{w}, \widetilde{\delta}), \tag{A.5}$$

where unemployed workers accept any job with $w \geq w_r$. Hence, the distribution of offers accepted out of unemployment is truncated below by a reservation wage that is independent of the individual riskiness of the jobs taken. *This perfect independence is not exact in the case with assets, i.e., in the main text model.*[2]

Since $u(c)$ is strictly increasing, Equation (A.5) immediately implies the following result:

Corollary 1 (Properties of the reservation wage w_r)**.** *The search efficiency in unemployment relative to that in employment determines the sign of $w_r - b$: $sgn(\lambda_U - \lambda_E) = sgn(w_r - b)$ provided a positive measure of offers satisfies $W > U$. Furthermore, the reservation wage is independent of the job offer's specific risk $\widetilde{\delta}$.*

[2]This generalizes the result in Pinheiro and Visschers (2015) based on risk neutrality and the point in Footnote 15 in the (also risk-neutral) model of Jarosch (2023).

This is a standard McCall result, extended to risk aversion and separation risk heterogeneity: when jobs arrive faster in unemployment, $\lambda_U > \lambda_E$, the unemployed demand a strictly higher reservation wage than the flow income in unemployment, $w_r > b$, and vice versa. For $\lambda_U = \lambda_E$, $w_r = b$.

Despite the reservation wage being independent of the riskiness of any individual offer, risk aversion and job riskiness both still shape the *level* of this reservation wage. The utility function appears in Equation (A.5) directly and by influencing the value functions $W(\widetilde{w}, \widetilde{\delta})$ and U. We show in the next section how risk aversion lowers w_r. Finally, when job search is equally effective in unemployment and employment, $\lambda_E = \lambda_U$, Equation (A.5) reduces to $u(w_r) = u(b) \Rightarrow w_r = b$, i.e., the reservation wage is equal to the flow value of unemployment. This is true regardless of the level of risk aversion or risk heterogeneity.

EE transitions. The gains from moving from job (w, δ) to $(\widetilde{w}, \widetilde{\delta})$ can be seen using Equation (A.3). When considering a new job, the worker assesses both the utility gain from higher wages, $u(\widetilde{w})$ vs. $u(w)$, and the separation risk difference, $\widetilde{\delta}$ vs. δ. For all reservation jobs $(\widetilde{w}, \widetilde{\delta})$ out of any job (w, δ), we obtain the following expression for *indifference curves* (IC) between wages and separation risk:[3]

$$u(\widetilde{w}) - u(w) = \left(\widetilde{\delta} - \delta\right) \cdot \Big[W(w, \delta) - U\Big] \quad \forall\ (\widetilde{w}, \widetilde{\delta}) \text{ s.t. } W(\widetilde{w}, \widetilde{\delta}) = W(w, \delta). \tag{A.6}$$

Equation (A.6) defines ICs in $(\widetilde{w}, \widetilde{\delta})$ space. These ICs fully characterize the EE moves. Since $W(w, \delta)$ is increasing in w and decreasing in δ, the left-hand side defines the minimum wage premium a job needs to offer to induce the worker to switch for a given change in risk. Defining $w_e(\widetilde{\delta}, w, \delta)$ as the value of $\widetilde{w}$ that solves Equation (A.6) gives the reservation wage for EE moves: workers in a job with (w, δ) accept any offer with risk $\widetilde{\delta}$ as long as it offers a wage of at least $w_e(\widetilde{\delta}, w, \delta)$.

To characterize the local behavior of the IC, we differentiate the ICs with respect to $\widetilde{\delta}$:

$$\frac{\partial w_e(\widetilde{\delta}, w, \delta)}{\partial \widetilde{\delta}} = \frac{W(w, \delta) - U}{u'(w_e(\widetilde{\delta}, w, \delta))} \quad \text{and} \quad \frac{\partial^2 w_e}{\partial \widetilde{\delta}^2} = -\frac{u''(w_e)}{u'(w_e)} \left(\frac{\partial w_e}{\partial \widetilde{\delta}}\right)^2. \tag{A.7}$$

The ICs implied by Equation (A.6) have three important features. First, the EE ICs are shaped by risk aversion, both directly by the utility functions on the left-hand side, and indirectly through its effects on $W(w, \delta) - U$. As shown by Equation (A.7), for a given job (w, δ), ICs are linear in $(\widetilde{w}, \widetilde{\delta})$ space under risk neutrality since $u'(c) = 1$. Under risk aversion ($u''(c) < 0$) they are strictly convex.

Second, the right-hand side of Equation (A.6) features the difference in the separation rates, $\widetilde{\delta} - \delta$. Under homogeneous δ, this term is always zero, and hence workers simply climb up the job ladder whenever the new job offers a higher wage: $w_e(\delta, w, \delta) = w$, *irrespective of risk aversion and unemployment value*. If $\widetilde{\delta} \neq \delta$, then this term is non-zero and risk shapes job mobility.

Third, the difference in separation rates in Equation (A.6) is multiplied by the value gap, $W(w, \delta) - U$. The higher this gap is—e.g., in high-wage jobs—the larger the additional unemployment risk looms. For instance, for the reservation job out of unemployment, this price is zero as $W(w, \delta) = U$, confirming the earlier result that the reservation wage for UE transitions is invariant to job safety. As shown by Equation (A.7), the ICs are upward-sloping for all jobs that are strictly

[3] To see this, evaluate Equation (A.3) for both (w, δ) and $(\widetilde{w}, \widetilde{\delta})$ and note that $W(w, \delta) = W(\widetilde{w}, \widetilde{\delta})$ along an IC. Further, when two jobs offer the same utility their acceptance sets are identical, i.e., $A(\widetilde{w}, \widetilde{\delta}) = A(w, \delta)$. Our Equation (A.6) therefore extends prior work by Pinheiro and Visschers (2015) to a setting with risk aversion. Jarosch (2023) does not feature an analog of Equation (A.6) but does present indifference curves in his Figure 1. Neither paper features risk aversion and the resulting inefficiencies from incomplete markets.

better than unemployment, i.e., $W(w,\delta) > U$. We can summarize these ideas succinctly by defining the marginal rate of substitution between wages and risk. This is the slope of the indifference curve at *current job* characteristics ($\widetilde{\delta} = \delta$), showing the wage increase a worker would require to compensate her for a marginal increase in risk: $MRS(w,\delta) \equiv \partial w_e(\delta, w, \delta)/\partial\widetilde{\delta} = (W(w,\delta) - U)/u'(w)$. Differentiating the MRS and using the monotonicity of $W(w,\delta)$ leads to the following result:

Corollary 2 (Monotonicity of $MRS(w,\delta)$)**.** *The marginal rate of substitution between wage and risk is strictly increasing in the wage and strictly decreasing in the risk for jobs preferred to unemployment.*

The MRS is equal to the price of risk defined in Equation (7) in the main text, adapted to a hand-to-mouth model. Hence, in the hand-to-mouth model the IC slope (evaluated locally around the current job) is exactly equal to the price of risk: $\text{PR}(w,\delta) = MRS(w,\delta)$.

A.3 Risk aversion, policy, and productive efficiency

To isolate the novel effect of risk aversion on—and through—job mobility, we bypass the standard source of inefficiency that is due to bargaining between workers and firms.[4] Thus, for the remainder of this section, we assume the following:

Assumption 1. *Workers have full bargaining power so that wages capture full match output, $w(z) = z$.*

This is equivalent to setting $\psi = 1$ and $\chi = 0$ in the wage setting rule, Equation (16).

Planner's problem: productive efficiency. In the spirit of Acemoglu and Shimer (1999), we introduce a planner's problem that focuses on productive efficiency. To this end, we consider a planner who can dictate all workers' job acceptance decisions and freely allocate consumption across workers. We formally set up this problem using the method of Lucas and Moll (2014) (derivations will be made available as a supplemental online appendix) and show that the timeless solution to the social planner's problem is captured by the following value functions U^{sp} and $W^{sp}(w,\delta)$:[5]

$$\rho U^{sp} = b_{HP} + \lambda_U \int_{A_U^{sp}} \left(W^{sp}(\widetilde{w},\widetilde{\delta}) - U^{sp}\right) dF(\widetilde{w},\widetilde{\delta}) \tag{A.8}$$

$$\rho W^{sp}(w,\delta) = \underbrace{w}_{=z} + \lambda_E \int_{A^{sp}(w,\delta)} \left(W^{sp}(\widetilde{w},\widetilde{\delta}) - W^{sp}(w,\delta)\right) dF(\widetilde{w},\widetilde{\delta}) + \delta\left[U^{sp} - W^{sp}(w,\delta)\right]. \tag{A.9}$$

The planner maximizes the net present value of output, including home production. For notational convenience and consistency with the market allocation, we write the value functions in terms of wages since $w = z$. We define the planner's allocation and efficiency as follows.

Definition 2 (Planner's allocation)**.** *The planner's allocation consists of value functions $\{U^{sp}, W^{sp}(w,\delta)\}$ and the distribution of workers across jobs $g_E^{sp}(w,\delta)$ such that U^{sp} and $W^{sp}(w,\delta)$ solve Equations (A.8)–(A.9) and the job acceptance sets A_U^{sp} and $A^{sp}(w,\delta)$ induce the distribution of workers described by $g_E^{sp}(w,\delta)$ given job offer arrival and destruction rates.*

[4] For example, in Jarosch (2023), i.e., a risk neutral model (allowing for returns to experience, which we abstract from for the question at hand), the decentralized allocation is socially efficient only when workers have full bargaining power in a sequential auction framework for wage renegotiation.

[5] We adopt the notational convention that all social planner-related objects are defined analogously to their decentralized equilibrium counterparts and we distinguish them by the use of the superscript sp.

Definition 3 (Productive efficiency). *Let $\mathcal{E} = \{U, W(w,\delta), g_E(w,\delta)\}$ be the decentralized equilibrium and let $\mathcal{E}^{sp} = \{U^{sp}, W^{sp}(w,\delta), g_E^{sp}(w,\delta)\}$ be the planner's allocation. $\mathcal{E}$ achieves productive efficiency (i.e., $\mathcal{E}$ is efficient) if, and only if, $g_E(w,\delta) = g_E^{sp}(w,\delta)\ \forall w, \delta$.*

Inspection of Equations (A.2)–(A.9) leads to the following result: the planner's and worker's problems coincide when there is no UI and workers are risk-neutral.

Lemma 1. *Let Assumption 1 hold. If $b_{UI} = 0$ and $u(c) = c$, then the decentralized allocation achieves productive efficiency.*

This follows immediately by direct comparison of the decentralized and planner allocations. Thus, our result extends a version in Acemoglu and Shimer (1999), that without risk aversion, there is no productive efficiency motivation for UI benefits. Intuitively, while our model departs structurally on various dimensions, the essential incremental extension is to include heterogeneity in EU risk across jobs and on-the-job search.

In the rest of the section, we explore the productive-efficiency motivation for UI benefits that emerges in the case with risk aversion, departing from the baseline case of the market allocation that achieves productive efficiency under risk-neutral preferences and no UI benefits. Furthermore, we show that risk aversion affects both UE and EE transitions. Thus, a natural question is the feasibility of restoring productive efficiency through appropriate policy choices, e.g., UI benefits. We proceed in two steps, first studying the model without, and then with heterogeneous job-loss risk.

A.3.1 No heterogeneity in job separation risk: productive efficiency at the EE margin, potential inefficiency at the UE margin, and the role of UI

We begin with an auxiliary model that abstracts from differences in job separation risk: $\underline{\delta} = \overline{\delta} = \delta$. This assumption immediately yields a first important result: there is no wage-job safety tradeoff, and EE transitions are independent of risk aversion. Workers move from one job to another, provided the new job offers a higher wage. This is true for both the planner and market allocations, meaning that EE decisions are always productively efficient.[6] This assumption renders the EE decisions independent of b and of the unemployment value more broadly: all jobs have the same risk, so unemployment does not factor into the cost-benefit calculation of EE transitions.

Second, we move to UE decisions. Since EE decisions are efficient under homogeneous unemployment risk, whether the decentralized allocation attains productive efficiency solely depends on the UE margin. This invites the question of how on-the-job search modifies the insights on the efficiency properties of UE transitions in the broad logic of Acemoglu and Shimer (1999). The reservation wage with no δ heterogeneity satisfies the following expression:

$$u(w_r) = u(b) + (\lambda_U - \lambda_E) \int_{w_r}^{\overline{w}} \frac{1 - F(\widetilde{w})}{\rho + \delta + \lambda_E(\widetilde{w})} u'(\widetilde{w}) d\widetilde{w}, \tag{A.10}$$

where $\lambda_E(\widetilde{w}) = \lambda_E \int_{\widetilde{w}}^{\overline{w}} f(x)\, dx = \lambda_E(1 - F(\widetilde{w}))$ is the EE outflow rate from a job paying wage $\widetilde{w}$. Since UE flows are pinned down by w_r, we can characterize the inefficiency brought about by risk aversion and analyze how policy can correct it using only this condition.

Proposition 1 (EE efficiency, UE inefficiency). *Let assumption 1 hold and further assume $\underline{\delta} = \overline{\delta} = \delta$, $u'' < 0$, $\lambda_E \leq \lambda_U$, and $w_r^{sp} > \underline{w}$. Then:*

[6] This result does not require workers to hold full bargaining power, i.e., $w = z$, but would emerge whenever the ranking of wages and productivity coincide (i.e., whenever worker bargaining power is positive and no imperfectly correlated second dimension of job quality exists, such as separation risk).

1. *Job acceptance sets by employed workers in the decentralized allocation and the planner's allocation coincide irrespective of b.*

2. *If $\lambda_E < \lambda_U$, then the reservation wage in the decentralized equilibrium when $b_{UI} = 0$ is strictly below that of the social planner, $w_r < w_r^{sp}$. There exists a strictly positive level of UI, b_{UI}^*, that restores productive efficiency in the decentralized equilibrium.*

3. *The level of UI to restore productive efficiency is strictly decreasing in λ_E.*

4. *If $\lambda_E = \lambda_U$, then the market allocation is efficient when $b_{UI} = 0$, regardless of risk aversion.*

Thus, when offers arrive faster in unemployment, risk-averse unemployed workers in the absence of UI accept some jobs the planner would refuse, inefficiently raising total employment: risk aversion makes them value the option of continued search—surplus that accrues at high-consumption, low-marginal-utility states—less than the planner does, so $w_r < w_r^{sp}$.[7] Providing UI benefits can undo this inefficiency. Specifically, holding all other parameters equal, the higher is OJS efficiency, the lower is the need for positive UI benefits to restore productive efficiency (Point 3). Intuitively, as λ_E rises toward λ_U, the reservation wage in the market economy shrinks toward b. Unemployed workers know they can continue to search for new jobs once employed, becoming less selective at the UE margin and reducing under-selectivity of the market relative to the planner. In the limit of $\lambda_E = \lambda_U$ (Point 4), $w_r = b = b_{UI} + b_{HP}$ and $w_r^{sp} = b_{HP}$ exactly, so equalizing the market and planner reservation wages requires UI to be zero. In this case, positive UI would break efficiency by causing the market to reject jobs the planner would take, regardless of the degree of risk aversion.

These results echo those by Acemoglu and Shimer (1999), who exclusively study the UE margin. In their general equilibrium directed search model there is no OJS. The random search, partial equilibrium counterpart to their analysis is $\lambda_E = 0$—a special case of this proposition. As we show, the presence of OJS mitigates or even nullifies the gap between the decentralized equilibrium w_r and the social planner's reservation wage w_r^{sp} out of unemployment, a result that is, to our knowledge, new.

A.3.2 Productive efficiency when jobs are heterogeneous in unemployment risk

We now show that heterogeneity in job separation risk introduces a new role for UI in promoting productive efficiency, even when $\lambda_E = \lambda_U$, where in the case of homogeneous separation risk no such inefficiency exists either along the UE or the EE margins. We begin by stating a result implied by the convexity of wage-job safety ICs under risk aversion and the linearity of the planner's ICs.

Proposition 2. *Let Assumption 1 hold, suppose $u'' < 0$, and suppose there is heterogeneity in job separation risk. Then, the job acceptance sets in the decentralized allocation $A(w, \delta)$ and in the planner's allocation $A^{sp}(w, \delta)$ in general do not coincide, unless the set of jobs is sufficiently small. Thus, the EE margin is not productive-efficient.*

The first part follows directly from Lemma 1, where we establish that the planner's ICs coincide with those under risk-neutral preferences—hence, are linear—and no UI. Under risk aversion, the ICs are convex. Thus, fully restoring productive efficiency in the presence of heterogeneity in

[7] Assuming $w_r^{sp} > \underline{w}$ ensures that the planner rejects a positive mass of job offers out of unemployment. The assumption $\lambda_E \leq \lambda_U$ signs this distortion in the natural direction. By Corollary 1, it keeps reservation wages weakly above the flow value of unemployment, so the market is under-selective and a *positive* UI level is the right corrective instrument. If instead $\lambda_E > \lambda_U$, reservation wages fall below the flow value of unemployment and correcting an over-selective market would require a *negative* b_{UI}, i.e., a tax on the unemployed.

Figure A.1: Indifference curves: planner vs. market

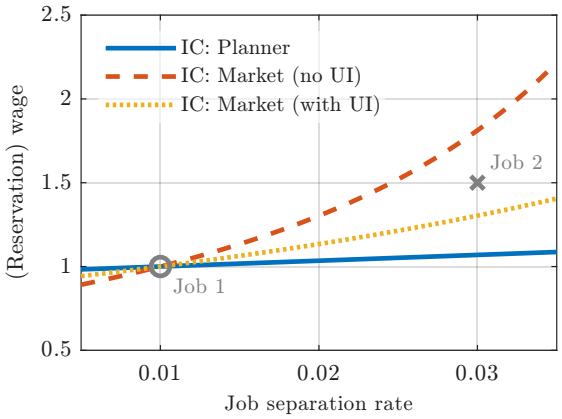

Notes: This figure plots the ICs in the two-job model. Value is increasing in the north-west direction. There are two jobs in the model denoted by the two points on the plots. The ICs use Equation (A.6) to find the wage required to exactly compensate for a change in risk relative to job 1. Parameter values are $w_1 = 1, w_2 = 1.5, \delta_1 = 0.01, \delta_2 = 0.03, b_{HP} = 0.1, \rho = 0.05, \sigma = 1.75, \lambda_U = \lambda_E = 0.2$ and $\Psi_1 = 0.5$. The blue line is the planner's IC. The red dashed line plots the market IC with $b_{UI} = 0$ and the orange dotted line plots it for $b_{UI} = 0.1$.

separation risk is not always possible, as we discuss below. To illustrate this analytically (and give a proof by example supporting Proposition 2), we use a two-job special case of our model.

The basic logic is intuitive. Generally, an EE inefficiency emerges when the ICs of the social planner and the market diverge out of viable origin jobs—and if there are jobs in the offer distribution that fall in between these two ICs. To fix ideas, Figure A.1 plots such a case, using two jobs: a safe but low-wage origin job "1" and a candidate destination job "2," which is riskier but pays more. The social planner is more risk-tolerant, indicated by a flatter (and linear) IC, while the market requires a larger compensation to take on the risky job. Without UI, the market rejects, but the social planner accepts job 2. These divergent choices reflect productive inefficiency. However, by shifting the price of risk by shaping the relative values of employment and unemployment, UI can—at least for a given set of origin and destination jobs—restore productive efficiency. But more generally, given the rich heterogeneity and many possible job offers, a single labor market policy instrument cannot align all such choices (but still can improve efficiency, as we describe below) even in this simpler case of no self-insurance. Formally, acceptance regions $A(w, \delta)$ and $A^{sp}(w, \delta)$ do not coincide for all (w, δ) origin jobs, in the viable support of the offer distribution for both destination and origin jobs. We now formally describe this two-job setup in which UI can restore efficiency, and build intuitions.

Special case: EE distortions and UI in a two-job model. We now present the model that underlies Figure A.1, and elaborate on the economics. Suppose that there are two jobs, (w_1, δ_1) and (w_2, δ_2). Job 1 is offered with probability Ψ_1, job 2 is offered with $\Psi_2 = 1 - \Psi_1$. Suppose $\lambda_E \leq \lambda_U$. Let $W_1 = W(w_1, \delta_1)$ and $W_2 = W(w_2, \delta_2)$. Let $u_U = u(b)$, $u_1 = u(w_1)$ and $u_2 = u(w_2)$ for the market allocation. Similarly, we have, with a slight abuse of notation (economizing on the sp superscript), $u_U = b_{HP}$, $u_1 = w_1$, and $u_2 = w_2$ for the planner. With the appropriate choice of these utilities

(u_U, u_1, u_2), the value functions for either the planner or the market solve:

$$\begin{aligned}\rho U &= u_U + \lambda_U \Psi_1 \max\{W_1 - U, 0\} + \lambda_U \Psi_2 \max\{W_2 - U, 0\} \\ \rho W_1 &= u_1 + \delta_1 (U - W_1) + \lambda_E \Psi_2 \max\{W_2 - W_1, 0\} \\ \rho W_2 &= u_2 + \delta_2 (U - W_2) + \lambda_E \Psi_1 \max\{W_1 - W_2, 0\}.\end{aligned} \tag{A.11}$$

With only three consumption levels, we can use a positive affine transformation of the utility function to make the comparison of the market and planner straightforward. We assume CRRA utility, $u(c) = c^{1-\sigma}/(1-\sigma)$ and discuss changes in risk aversion, σ.

Lemma 2. *Suppose agents have CRRA utility with $\sigma \geq 0$ and Assumption 1 holds. Consider the two-job model with $b < w_1 < w_2$. Then the worker's* choices *in the market equilibrium are identical to the choices in a modified problem where they have three utility levels $u_U = b$, $u_1 = w_1$, and $u_2 = \widetilde{w}_2$, where $w_1 < \widetilde{w}_2 \leq w_2$. For given b, w_1, w_2, the level of risk aversion σ defines a unique $\widetilde{w}_2$, with higher risk aversion implying a lower $\widetilde{w}_2$. For $\sigma = 0$ we have $\widetilde{w}_2 = w_2$, and as $\sigma \to \infty$, $\widetilde{w}_2 \to w_1$.*

Using this lemma and (generic) Equation (A.11), we see that the problems of the planner and market are identical except that the planner values the consumption stream of the high-wage job at w_2, while risk aversion causes the market to value it at the lower value of $\widetilde{w}_2$ where $w_1 < \widetilde{w}_2 \leq w_2$.[8]

Next, consider the following thought experiment. For the planner, assume that both jobs are preferred to unemployment. Assume that job 2 has higher pay and weakly higher risk: $w_2 > w_1$ and $\delta_2 \geq \delta_1$. Suppose the planner ranks the risky job above the safe job: $W_2^{sp} > W_1^{sp} > U^{sp}$. Thus, in the planner's allocation unemployed workers should accept either job, and workers should move from the safe job to the risky but high-paying job.

Our main result from this section is contained in the following proposition:

Proposition 3. *Let Assumption 1 hold. Consider the two-job model where the planner prefers the risky job to the safe job, and both jobs to unemployment, $W_2^{sp} > W_1^{sp} > U^{sp}$.*

1. *Suppose that both jobs are equally risky: $\delta_2 = \delta_1$. Then the market allocation without UI ($b_{UI} = 0$) is productively efficient regardless of the level of risk aversion.*

2. *Suppose that job 2 is strictly riskier than job 1: $\delta_2 > \delta_1$.*

   (a) *If workers are sufficiently risk averse, then in the market allocation with $b_{UI} = 0$, workers prefer the safe to the risky job. Workers move from the risky to the safe job, take both jobs from unemployment, and the market allocation does not attain efficiency.*

   (b) *For any level of risk aversion, there exists a positive level of UI benefits $b_{UI} > 0$ that makes the workers prefer the risky job. For w_1 sufficiently large, the lowest such level of UI does not induce workers to prefer unemployment over the safe job, thus restoring the planner's allocation.*

This proposition encompasses the new insights from our framework. Point 1 shows that in the absence of risk heterogeneity, there is no motivation for UI even with risk aversion, as the market allocation is productively efficient under $b_{UI} = 0$. This follows from Proposition 1. Since the UE

[8] Fixing $u_U = b$ and $u_1 = w_1$ while letting only $u_2 = \widetilde{w}_2$ vary with σ is a normalization as choices are invariant to positive affine transformations of utility, whose two free parameters pin down the first two utility levels, loading all curvature into the third. Second, this does not presume that the worker accepts job 1 out of unemployment; in principle, the option-value under $\lambda_U > \lambda_E$ could make her reject it. However, risk aversion never flips the ranking $u(b) < u(w_1)$, and with $b_{UI} = 0$ it (weakly) *lowers* the reservation wage out of unemployment relative to the planner's (the under-selectivity of Proposition 1). Given our maintained assumption that the planner prefers both jobs to unemployment, the market therefore accepts job 1 out of unemployment for every $\sigma \geq 0$, so the b-versus-w_1 margin does not bind.

margin is undistorted, by construction, in this parameter configuration, and without heterogeneity in EU risk, the EE margin is always undistorted.[9]

In contrast, Point 2(a) states that when there is heterogeneity in EU risk, incomplete markets and risk aversion can distort the EE margin. If workers are sufficiently risk averse, they will prefer a safe but low productivity job, even though the social planner prefers the risky high productivity job. The UE margin remains undistorted, but since workers make EE moves in opposite directions in the market and planner solutions, their allocation differs, with the share of safe but lower-productivity jobs higher in the market allocation.[10] This leads to lower productivity, EU rates, and unemployment in the market allocation than in the productively efficient allocation.

Finally, Point 2(b) shows that UI can play a role in correcting this distortion at the EE margin. Intuitively, by raising the value of unemployment, higher UI makes employed workers less afraid of unemployment and more willing to take risks, pushing them to prefer the high-wage high-risk job, as desired by the planner.Importantly, this motivation for positive UI exists *because of, not despite, on-the-job search*, and holds even when $\lambda_E = \lambda_U$—exactly when the inefficiency at the UE margin disappears and the productive inefficiency it reflects emanates purely from the EE margin.

These results are illustrated in Figure A.1, for a numerical example where $\delta_2 > \delta_1$ and $\lambda_E = \lambda_U$, with the precise parameter values reported in the figure note. We plot ICs from job 1 in the $w - \delta$ space, with jobs 1 and 2 illustrated by the points marked with a circle and cross, respectively.[11] The planner's IC (blue line) is linear and relatively flat, reflecting the planner's limited need for compensation against risk. Since job 2 lies above the planner's IC, job 2 is preferred to job 1 by the planner. The market's IC (red dashed line) is convex and steeper than the planner's, reflecting the worker's larger distaste for risk under incomplete markets. This means that job 2 lies below the market's IC, so job 1 is preferred to job 2 by the market for this level of risk aversion. The orange dotted line shows the role of UI in correcting this distortion, and plots the market's IC for a sufficiently positive level of UI. Higher UI insulates against unemployment risk, reducing the price of risk encoded in the IC. This flattens the market IC, bringing job 2 above the IC and aligning the market's choice with the planner's.

Impossibility of fully recovering productive efficiency in general when $\underline{\delta} < \bar{\delta}$. This two-job example illustrates how UI can push risk-averse workers' EE decisions back toward those of the planner. The broader logic also holds in models with more than two jobs, with the obvious but crucial caveat that it is no longer possible to *fully* replicate the planner's solution in general using just a single UI instrument. This result follows immediately, since indifference curves under risk aversion are strictly convex in the market allocation while they are linear in the planner's allocation.[12] In our rich setting with two-dimensional match heterogeneity over a continuum of jobs, a single instrument generally fails to fully restore efficiency. However, as we show in our quantitative work, a positive level of UI is still helpful for raising aggregate productivity. Moreover, in Figure C.4 Panel (f), we show that output net of vacancy posting costs actually increases in UI benefits over a range of low benefit levels in our baseline model. This directly shows how, over this

[9] When $\delta_2 = \delta_1$, a restriction implied by the assumed planner ranking of jobs and unemployment is that $w_2 > w_1 > b$. Trivially in this case, workers also prefer to leave job 1 for job 2.

[10] When $\delta_2 > \delta_1$, for the ranking $W_2^{sp} > W_1^{sp} > U^{sp}$ to hold it must even more so be the case (relative to $\delta_2 = \delta_1$) that $w_2 > w_1 > b$. Thus, in the decentralized allocation workers are still willing to accept both jobs relative to unemployment when $b_{UI} = 0$ as they increase their income and only risk is for it to fall back to b. Thus, the UE margin is undistorted when $b_{UI} = 0$.

[11] The indifference curves are computed using the formula in Equation (A.6).

[12] For sufficiently spread-out offer distributions it is not possible for $A(w, \delta)$ and $A^{sp}(w, \delta)$ to coincide for all (w, δ), since these sets are generated from indifference curves with different shapes.

range, raising UI creates more resources on net, raising the productive frontier at the same time as providing more insurance.

B Empirical appendix

B.1 Details on the publicly available datasets

In this section we elaborate on data construction, measurement and sample selection in the Current Population Survey (CPS) and the Survey of Consumer Expectations (SCE). These data underlie Figure 1 and various moments that our quantitative model targets in Section 4.1.

B.1.1 Aggregate moments used in the calibration

The unemployment rate (5%) and the aggregate monthly EU rate (1.3%) are economy-wide 1996–2007 averages, CPS-based series published by the BLS, and the aggregate monthly EE rate (2.6%) is the average of the series constructed by Fujita et al. (2024), who adjust for missing responses induced by CPS design changes over time.[13,14] We do not separately target the UE rate; it is implied by the steady-state flow balance given the unemployment rate and the EU rate. We choose the 1996–2007 window to average all three moments because the labor market was broadly in balance in this period and because the EE rate series is only available starting from December 1995.[15] To be clear, the only calibration target constructed from our own CPS micro-data sample (described below) is the EU-tenure gradient of Figure 1 Panel (a): we target a 1.3 pp difference between the mean EU rate at tenures 0–4 and at tenures 40–44 months in the specification with worker controls.

B.1.2 Current Population Survey

The CPS has a rotating-panel design, in which households are surveyed for four consecutive months, then they rotate out for eight months and are surveyed for another four months, after which they permanently leave the sample. This structure allows us to match at most three-fourths of the sample in one month to the next.

Sample selection. We restrict our sample to prime-age household heads aged 25 to 55. We keep observations from 1994, when employer changes become identifiable in the CPS, through May 2026. (The results are nearly identical if we instead use the shorter time period restricted to stable macroeconomic conditions that provide our aggregate calibration targets, see Footnote 15 above.) We drop workers who report being self-employed, in the armed forces or an unpaid family worker. We further exclude individuals enrolled in school. We only keep workers who reside in the 50 states and Washington, D.C.

Worker flows. We use monthly CPS data from IPUMS, categorizing each individual in each month t as employed (E), unemployed (U) or out of the labor force (N). We track individuals across months using CPSIDP—the IPUMS person identifier built from rotation group, household and

[13] We note that in particular the EE rate may have trended down subsequently, although this shift is up for debate (Fujita et al., 2024).

[14] Specifically, we retrieve these series from FRED. For the EE rate we use `FMPSA3MA`. We calculate the unemployment rate by dividing the seasonally adjusted unemployment stock (`UNEMPLOY`) by the seasonally adjusted labor force (`CLF160V`). We calculate the EU rate by dividing the seasonally adjusted EU flows (`LNS17400000` forwarded by one month) by the seasonally adjusted employment stock (`CE160V`).

[15] The time window is inconsequential for our quantitative analysis. If we instead use the 1996–2026 averages, we obtain 5.5% for the unemployment rate, 1.3% for the EU rate, and 2.4% for the EE rate, values quite close to our calibration targets. Conversely, if we restrict the CPS and SCE samples (described below, which end in 2026) to the 1996–2007 period, we find essentially identical results for our CPS and SCE analyses in Section 2 and Figure 1—in particular, including the calibration target for the EU-job tenure gradient.

line number, race, sex, and age—and compute monthly individual-level transitions by comparing labor force status in adjacent months. For example, $\mathbb{I}(EU_{it}) = 1$ if a worker is employed in t and unemployed in $t+1$, and 0 otherwise. We construct analogous indicators for all other flows among E, U and N. For employer-to-employer (EE) transitions, we use the IPUMS variable EMPSAME, which asks an employed worker whether they are with the same employer as in the previous month; we record an EE transition when a worker is employed in two adjacent months and reports an employer change. This question exists only starting in 1994, when dependent interviewing was introduced, so EE transitions cannot be measured earlier (Fallick and Fleischman, 2004).

Because these indicators compare point-in-time monthly statuses rather than continuous work histories, they are subject to time-aggregation bias. A worker who separates and is reemployed within the same month is recorded as a job stayer (or, if the employer differs, as an EE mover), and a worker who transitions through unemployment between two interview dates but reports employment at both is not recorded as an EU separation.

Job tenure. We complement the core monthly data with the Job Tenure & Occupational Mobility Supplement, which measures job tenure as the length of time worked at the current job. For jobs that begin in-sample—via UE, NE, or EE transitions—we set tenure to 0. We then increment tenure by one month for each month the worker remains in the same job, i.e., until an EU, EN, or EE separation. Tenure thus comes from two sources, and we weight accordingly: workers drawn from the supplements receive the supplement sampling weights, while workers whose jobs begin in-sample receive the basic monthly weights. Our tenure-based analysis in Panel (a) uses only the subset of workers for whom tenure can be measured—whether observed in the supplements or constructed from monthly transitions—whereas Panels (b) and (c) use the larger CPS sample built from the monthly files.

Worker and job controls. For job controls we include industry and occupation dummies. We use the 1990 Census Bureau industrial classification system (IND1990) and the Census Bureau 2010 occupation codes (OCC2010), as harmonized over time by IPUMS. We aggregate these industries and occupations to broad categories, resulting in 13 industries and 25 occupations. As for worker-side observables, we assign individuals to 5 education categories: less than high school, high school graduate, some college/vocational school/associate degree, college graduate, and advanced degree. We use three broad race categories: white, black and other. We define potential experience as age minus 25 (the youngest age in our sample, see above) and include its quartic in our controls. We further include a gender dummy and marital status dummies, and control linearly for the number of own children in the household, both the total and those under age five.

Wages. We calculate hourly wages using information on weekly earnings and hours worked in the outgoing rotation groups, reported up to twice per worker. First, for workers who are paid hourly, we take their reported hourly wages. Second, for salaried workers, we divide their reported weekly earnings by usual hours worked or, if they report their hours vary, by hours worked last week at their main job. We deflate wages by the CPI to calculate real wages (although we include time fixed effects in our analysis). Finally, we winsorize real wages at the bottom and top 0.3 percent. The wage control in our regressions is the log of each worker's average real wage across their (up to) two earnings observations.

Analysis samples and weighting, by figure panel. Figure 1 Panels (a)–(c) draw on the same CPS extract, sample restrictions, and years, but each panel uses the subsample and weights its object

Figure B.1: Companion to Figure 1 Panel (b): raw, cell-level plot of job-specific EU rates vs. wages

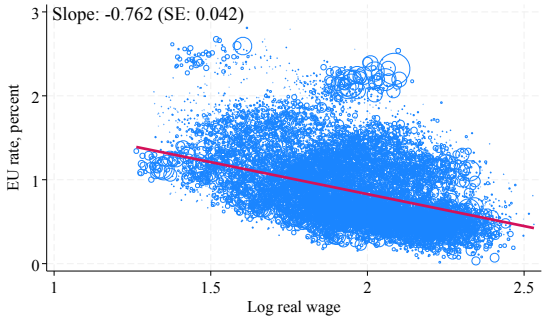

Notes: This figure plots the raw, cell-level monthly job-specific (industry-occupation-state) EU rates against log real wages from the regressions described in Footnote 5, fixing time effects and worker observables at baseline values—complementing Figure 1 Panel (b). Marker size reflects the survey-weighted employment share of each cell. The line is the OLS fit of predicted EU on predicted log wage, run at the cell level and weighted by these shares, reporting a robust standard error. Source: CPS.

requires. Panel (a) uses employed workers with measurable job tenure of 0–44 months, weighted as described in the job tenure paragraph above. Panel (b) uses all employed workers with non-missing controls, weighted by the basic monthly weights. Panel (c) restricts the sample to workers observed to make at least one EE transition while in the sample, so that the switcher and stayer distributions are computed over the same respondents; the underlying EU-prediction regression is weighted by the basic monthly weights, while the histograms are unweighted.

B.1.3 Survey of Consumer Expectations

The SCE is a nationally representative survey of a rotating panel of approximately 1,300 household heads conducted by the Federal Reserve Bank of New York. Respondents participate in the panel for up to 12 months, with a roughly equal number rotating in and out of the panel each month. We use the core monthly files in our analysis.

Sample selection. We restrict our sample to employed individuals aged 25 to 55, mirroring the CPS sample. We use observations from June 2014, when the survey began asking about employer changes, through August 2025. We exclude the self-employed and drop respondents whose labor market status is ambiguous because they report multiple statuses. Since our analysis requires within-worker changes across consecutive months, we drop a respondent's observations after the first skipped survey month.

Worker flows. We determine respondents' employment status in each month from the survey's employment status question, classifying respondents as employed, unemployed, or out of the labor force, and we calculate monthly transition dummies using the panel dimension of the survey, as in the CPS. We identify an EE transition when a respondent is employed in two consecutive months and reports working for a different employer than in the previous month.

Subjective transition probabilities and risk attitude. For the probability of a job loss, we use the answer to the question "What do you think is the percent chance that you will lose your main/current job during the next 12 months?" (`Q13new`). For the probability of quitting we use

the answer to "What do you think is the percent chance that you will leave your main/current job voluntarily during the next 12 months?" (`Q14new`). For the risk attitude measure we use the answer to "More generally, how would you rate your willingness to take risks in daily activities?" (`QRA2`) asked over a scale ranging from 1 to 7.

Analysis samples and weighting, by figure panel. Figure 1 Panel (d) is based on one-month within-worker first differences of the subjective job-loss probability, so it uses respondents employed in consecutive months; as in the CPS analysis, we restrict the sample to individuals who switch employers at least once during their panel participation, so that the switcher and stayer distributions are computed over the same respondents. The histograms are unweighted. Panels (e) and (f) use all employed respondent-months with non-missing risk attitude and outcome responses, weighted by the survey-provided respondent weights.

B.2 Details on our custom survey

In this section, we elaborate on our custom survey. We use three waves of the NORC AmeriSpeak Omnibus survey, fielded in May and June 2025, each containing our tailor-made module on reservation pay and subjective unemployment risk.

B.2.1 Questions used in the paper

Questions `FRKC1A`–`FRKC14` are asked only of respondents in paid employment (`EMPLOY1` $\in \{1, 2, 3\}$). Placeholders of the form `[INSERT ...]` are piped responses from earlier items. The reservation wages used in the headline estimates are taken from question `FRKC4`, elicited before respondents receive any information about UI benefits; the post-UI-information block `FRKC10` serves as a robustness check and as the reference elicitation for the UI generosity treatment. The current job reservation wage concept w_i^* from question `FRKC3` enters only the anchor robustness check. The UI generosity effects come from questions `FRKC11` and `FRKC12`.

EMPLOY1—Employment. Which of the following describes your current employment status?

(1) Working remotely — as a paid employee

(2) Working in-person — as a paid employee

(3) Working both in-person and remotely (a hybrid model) — as a paid employee

(4) Working — self-employed/own a business

(5) Not working — retired

(6) Not working — disabled

(7) Not working — other

FRKC1A—Pay frequency. How does your job pay you? Pick the option that fits best. (1) hourly, (2) weekly, (3) biweekly (every other week), (4) monthly, (5) annual.

FRKC1B—Current pay. What is your current [INSERT FRKC1A] pay from your job? Please provide the number in US dollars. Please report gross [INSERT FRKC1A] pay before taxes.

$ [NUMBOX; range 0–1,000,000,000]

FRKC2—Unemployment risk (current job). We now ask you about the unemployment risk in your current job. What do you think is the probability that this job leaves you unemployed at some point in the next 12 months? (For instance, you get dismissed, you quit, your firm closes and there are layoffs. . .). The numbers below are between 0 and 100 and are reported in percent.

- **0 percent** means you are sure you will not become unemployed at some point in the next 12 months.
- **10 percent** means that you think there is a 1 in 10 chance of you becoming unemployed at some point in the next 12 months.
- **50 percent** means it is equally likely that you will become unemployed or not become unemployed at some point in the next 12 months.
- **100 percent** means you are sure you will become unemployed at some point in the next 12 months.

Pick a probability between 0 and 100 percent: [SLIDER, 0–100] %

FRKC3—Reservation pay at current-job risk. You described your current job as follows:

- Your [INSERT FRKC1A] pay is $ [INSERT FRKC1B].
- The risk of your current job ending and you becoming unemployed at some point in the next 12 months is [INSERT FRKC2]%.

Now imagine a job offer at another employer. If you do not accept this offer, you can still stay at your current job. Suppose the new job is comparable to your current job. In particular, the unemployment risk is the same as in your current job, i.e., there is a [INSERT FRKC2]% probability that the new job ends and you become unemployed at some point in the next 12 months. However, the new job may have a different [INSERT FRKC1A] pay.

Below, please indicate the lowest [INSERT FRKC1A] pay level (in $) this new job needs to offer for you to accept it:

$ [NUMBOX; range 0–1,000,000,000]

FRKC4—Reservation pays across risk. You stated that you would accept a new job that has:

- an unemployment risk of [INSERT FRKC2]% at some point in the next 12 months,
- at least a(n) [INSERT FRKC1A] pay of $ [INSERT FRKC3], and is otherwise comparable to your current job.

But what if the new job had a different level of unemployment risk? By unemployment risk, we again mean the probability that the new job leads to unemployment at some point in the next 12 months. The table below shows different levels of unemployment risk in this new job for you to consider: 2%, 5%, 10%, 20%, and 50%. For each risk level, please enter the lowest [INSERT FRKC1A] pay level at which you would be willing to switch to this new job.

Unemployment risk level	Lowest [INSERT FRKC1A] pay at which I would accept this job
2% unemployment risk	$ [NUMBOX]
5% unemployment risk	$ [NUMBOX]
10% unemployment risk	$ [NUMBOX]
20% unemployment risk	$ [NUMBOX]
50% unemployment risk	$ [NUMBOX]

FRKC6—Consumption drop upon job loss. Suppose you lost your job and became unemployed.

Thinking through all the ways that your income and needs may be affected, what would happen to the amount of money you spend? Think about all the goods and services you may buy (e.g., restaurant meals, groceries, vacations, movies, new gadgets. . .).

After you lose your job, you would spend:

- Much less than before
- A bit less than before
- About the same as before
- A bit more than before
- Much more than before
- Don't know

FRKC10—Reservation pays after UI information. Typically, workers eligible for unemployment insurance can expect those benefits to be about 40% of their pre-tax pay. For instance, someone making $1,000 a week in their previous job can expect about $400 a week in unemployment insurance benefits. (This is a rough guideline.)

We now want to find out how the unemployment insurance system may affect your decision about which jobs to take. We are again interested in the lowest [INSERT FRKC1A] pay level at which you would be willing to accept a new job that is comparable to your current job but has a different unemployment risk level.

For your reference, the table below repeats your previous answers for various unemployment risk levels in the new job.

Below, you can revise your earlier answers considering the information about unemployment benefits the survey provided above. *Fill this column with [INSERT FRKC1A] pay levels. If you do not want to change your answers, you can simply type the same numbers you gave earlier.*

Unemployment risk level	Your previous answer	After learning about UI benefit levels
2% unemployment risk	$ [INSERT FRKC4A]	$ [NUMBOX]
5% unemployment risk	$ [INSERT FRKC4B]	$ [NUMBOX]
10% unemployment risk	$ [INSERT FRKC4C]	$ [NUMBOX]
20% unemployment risk	$ [INSERT FRKC4D]	$ [NUMBOX]
50% unemployment risk	$ [INSERT FRKC4E]	$ [NUMBOX]

FRKC11—Reservation pays under a $100/week UI supplement. Imagine the following scenario: in addition to the normal unemployment insurance benefits, the government now increases the amount by $100 per week. For instance, the worker that would have usually received $400 a week in unemployment insurance benefits now will receive $500, and a worker that would have received $200 now will get $300, and so forth. For this question, assume that the additional $100 per week applies to you too. How would your answers change? For your reference, the table below reports your previous answers.

Unemployment risk level	Your previous answer	After the $100/week UI increase
2% unemployment risk	$ [INSERT FRKC10A]	$ [NUMBOX]
5% unemployment risk	$ [INSERT FRKC10B]	$ [NUMBOX]
10% unemployment risk	$ [INSERT FRKC10C]	$ [NUMBOX]
20% unemployment risk	$ [INSERT FRKC10D]	$ [NUMBOX]
50% unemployment risk	$ [INSERT FRKC10E]	$ [NUMBOX]

FRKC12—Quit probabilities under the two UI regimes. What is the probability you will quit your current job to move to another employer at some point in the next 12 months?

Below, we ask you to provide this probability under two scenarios. While answering the question, please assume you are equally likely to have the opportunity to change your employer in both scenarios.

- Scenario 1: Unemployment benefits are at their normal level (i.e., excluding the additional $100 per week).
- Scenario 2: Unemployment benefits are at their higher level (i.e., including the additional $100 per week).

Select a probability between 0 and 100 percent for the two scenarios:

Scenario 1 [SLIDER, 0–100] % Scenario 2 [SLIDER, 0–100] %

B.2.2 Data cleaning

First, we keep respondents who report being paid employees (working remotely, in person, or in a hybrid arrangement) and aged 18–64, dropping retirees, the disabled, the self-employed, and those not working for other reasons, which yields 1,499 workers across the three waves. Second, we keep respondents with a valid (positive, non-missing) answer to every one of the five reservation pay questions in the baseline (i.e., `FRKC4`), so that the panel is balanced (1,330). Third, we drop respondents who completed the module in three minutes or less (1,079). Fourth, we drop respondents whose five raw answers are not weakly monotonic in risk, in either direction (1,008, the preferred sample, of whom 827 work full-time). We impose no sign restriction on any respondent's own estimated IC slope.

For each respondent, we express the reservation pay as the percent deviation from the respondent's own answer at the 10 percent risk level. We winsorize this outcome symmetrically at the 2nd and 98th percentiles, calculated separately for each risk level on the cleaned sample of employed respondents aged 18–64, before we apply any further sample restrictions. We then reuse the same winsorized reservation ratios throughout (rather than re-winsorizing within each robustness subsample considered below).

Current pay is used for descriptive statistics and for the heterogeneity cuts in Table 1, but not for sample selection. We treat the survey missing codes, zeros, and implausible magnitudes as missing. Specifically, to determine implausible reported pay levels, we apply pay-frequency-specific bounds: the floor is the federal minimum wage of $7.25 per hour and the cap is 30× the federal minimum, each scaled to the respondent's reporting frequency (hourly, weekly = $40\,h$, biweekly = $80\,h$, monthly = $173\,h$, annual = $173\,h \times 12$). We annualize reservation pay and current pay with the same conversion factors.

Table B.1: Descriptive statistics, preferred sample

Variable	Mean/Share	SD	Variable	Mean/Share
Age (years)	43.2	12.4	*Education (5-cat)*	
Yearly gross pay (USD k)	81.7	52.6	Less than HS	0.035
Annual EU risk (%)	20.7	25.1	HS graduate or equivalent	0.124
Household size	3.1	1.5	Some college/assoc. degree	0.325
Female	0.455		Bachelor's degree	0.301
Full-time work	0.861		Postgrad./profession. degree	0.215
Pay frequency			*Marital status*	
Hourly	0.258		Married	0.530
Weekly	0.081		Widowed	0.018
Biweekly	0.543		Divorced	0.105
Monthly	0.094		Separated	0.026
Annual	0.024		Never married	0.321
Census region			Living with partner	0.000
Northeast	0.147		*Race / ethnicity*	
Midwest	0.298		White, non-Hispanic	0.620
South	0.317		Black, non-Hispanic	0.112
West	0.238		Other, non-Hispanic	0.004
			Hispanic	0.168
			2+, non-Hispanic	0.040
			Asian-Pacific Islander	0.057

Notes: This table presents descriptive statistics on the analysis sample of 1,008 respondents. The annual gross pay is computed for full-time workers. A small number of respondents have missing information for some variables (e.g., full-time status).

B.2.3 Descriptive statistics

Our baseline analysis sample consists of 1,008 respondents. We present the key descriptive statistics in Table B.1. Figure B.2 shows the distribution of elicited subjective probabilities of losing one's job and of quitting to move to another employer in the next 12 months. The latter distribution is displayed for two cases: baseline UI generosity, and the hypothetical scenario of higher UI benefits (by \$100/week, described in Appendix B.5).

B.3 Further estimation results

Individual-level IC slopes. We estimate the individual-level counterpart to Equation (1),

$$\widetilde{r}_{ij} = \alpha_i + \beta_i \widetilde{\delta}_j + \varepsilon_{ij}, \tag{B.1}$$

on the five reservation-pay answers of the respective elicitation. Figure B.3 plots the distribution of β_i for the baseline elicitation, estimated on full-time workers.

Stability across elicitations. Figure B.4 shows that respondents report stable ICs across three separate elicitations in the survey. Specifically, against the baseline β_i from our main and first IC question, it plots each respondent's two additional estimates (one following a reminder of the actual US UI replacement rate and one where we ask about a hypothetical \$100 increase). The

Figure B.2: Subjective job loss and quit probabilities

(a) Subjective job loss probability

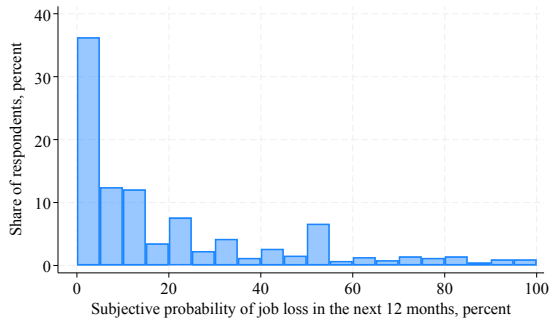

(b) Subjective quit probability

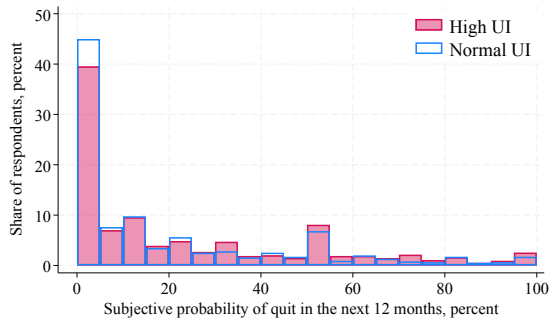

Notes: This figure plots subjective job loss and quit probabilities. Panel (a) plots the distribution of subjective probability of losing one's job in the next 12 months. Panel (b) plots the distribution of subjective probability of making an EE transition in the case of baseline (normal) UI and following a hypothetical \$100/week increase in (high) UI generosity.

Figure B.3: Distribution of individual IC slopes

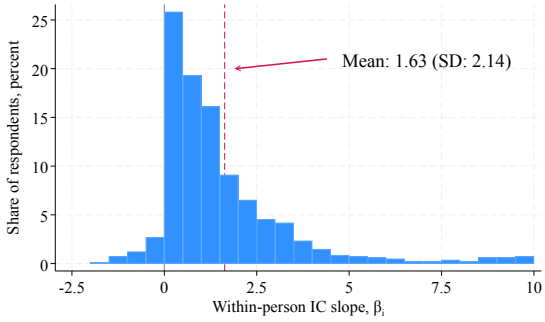

Notes: This figure plots the baseline within-person slope β_i distribution for full-time workers. The x-axis is trimmed to $[-2.5, 10]$. The text and figure report the untrimmed mean and standard deviation.

slopes are strongly correlated and cluster around the 45-degree line.

Demographic heterogeneity. Appendix Table B.2 reports estimates of Equation (1) augmented with interactions between the offered risk and demographic indicators,

$$\widetilde{r}_{ij} = \alpha_i + \beta\,\widetilde{\delta}_j + \gamma\left(\widetilde{\delta}_j \times D_i\right) + \varepsilon_{ij}, \tag{B.2}$$

where D_i indicates being female, young, white, married, or holding a bachelor's degree or higher. The worker fixed effect absorbs the level effect, so γ captures whether the IC slope differs by group.

B.4 Robustness of baseline IC estimate

Table B.3 reports the estimates behind Figure 3 Panel (b), where we vary one ingredient, for all workers (Panel A) and for full-time workers (Panel B).

Column (2) drops the monotonicity screen and thus retains "zigzag" response patterns. Column (3) replaces the balanced-panel requirement with a weaker one, requiring at least two valid

Figure B.4: Individual IC slopes across elicitations

(a) Post-UI information vs. baseline

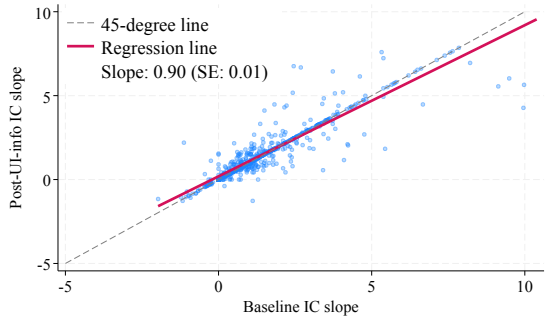

(b) UI increase vs. baseline

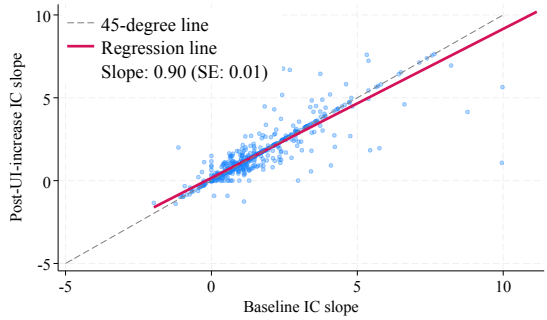

Notes: This figure plots the individual-level baseline IC slopes against the IC slopes for the two reelicitations, after a UI level reminder and the hypothetical UI increase. Each dot is a respondent with individual slopes estimated in both elicitations. The axes are trimmed to $[-5, 10]$. The slope and standard errors are indeed identical between the two panels to the second decimal.

Table B.2: The tradeoff between pay and job safety across demographic groups

	(1) $\widetilde{r}_{ij}$	(2) $\widetilde{r}_{ij}$	(3) $\widetilde{r}_{ij}$	(4) $\widetilde{r}_{ij}$	(5) $\widetilde{r}_{ij}$	(6) $\widetilde{r}_{ij}$
$\widetilde{\delta}_j$	1.775***	1.577***	1.423***	1.361***	1.613***	1.184***
	(0.103)	(0.095)	(0.127)	(0.096)	(0.120)	(0.197)
$\widetilde{\delta}_j \times$ Female	−0.342**					−0.287*
	(0.147)					(0.149)
$\widetilde{\delta}_j \times$ Young		0.130				0.287*
		(0.152)				(0.160)
$\widetilde{\delta}_j \times$ White			0.329**			0.280*
			(0.156)			(0.163)
$\widetilde{\delta}_j \times$ Married				0.486***		0.493***
				(0.145)		(0.158)
$\widetilde{\delta}_j \times$ College					0.034	−0.012
					(0.151)	(0.158)
N (job offers)	4135	4135	4135	4135	4135	4135
Within R^2	0.355	0.352	0.355	0.359	0.352	0.367
N (respondents)	827	827	827	827	827	827
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Full-time only	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the estimates of Equation (B.2) on full-time workers. The base coefficient is the IC slope of the omitted group: men, workers 41 and older, non-white workers, unmarried workers, and workers without a bachelor's degree. Standard errors clustered at the respondent level are reported in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

reservation-pay answers per worker, the minimum needed to identify a within-person slope; respondents without the 10-percent-risk answer still drop out because the anchored outcome is

Figure B.5: Sensitivity of the IC slope to the winsorization cutoff

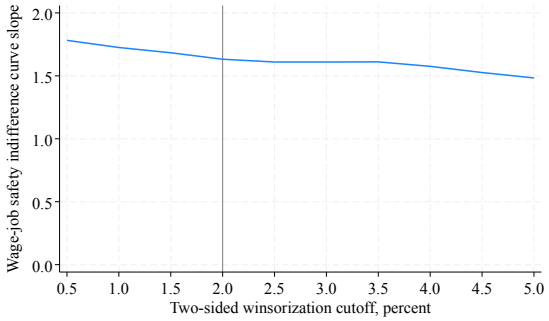

Notes: This figure plots estimates of Equation (1) for full-time workers against winsorization two-sided cutoffs. The vertical line marks the 2% baseline.

undefined without it. The weak monotonicity screen is only applied to the available reservation pay responses. Column (4) drops the three-minute completion screen and thus retains speeders.

Column (5) winsorizes the reservation ratio at the 1st and 99th instead of the 2nd and 98th percentiles. Figure B.5 varies the winsorization cutoff continuously. Column (6) trims instead of winsorizes. We still require all five risk levels to be answered, but because trimming sets the tail answers to missing rather than capping them, the resulting panel is mechanically unbalanced. We therefore require at least two remaining items per respondent to identify a within-person slope.

Column (7) re-estimates the baseline on the post-UI-information elicitation (question `FRKC10`) instead of the baseline elicitation (question `FRKC4`), applying the balance and monotonicity screens to that elicitation's own answers.

Column (8) keeps the baseline elicitation but anchors the reservation ratio to the separately elicited reservation wage at the current job's risk level, w_i^* (question `FRKC3`), instead of the own 10-percent-risk answer, so that no item is identically zero by construction; the specification requires a valid w_i^* and winsorizes all five items at the 2nd and 98th percentiles.

Column (9) replaces the sign-agnostic monotonicity condition with the requirement that the raw reservation pay at the 50 percent (highest) risk level strictly exceed that at the 2 percent (lowest) level. Column (10) instead replaces it with the requirement that the respondent's raw answers at every consecutive risk level are non-decreasing. Column (11) tightens this to strictly increasing.

B.5 The UI generosity treatment

We report results from the hypothetical \$100 per week UI increase in Table B.4. We generate a panel dataset at the respondent level for before and after (with and without) the hypothetical \$100/week UI benefit increase and estimate difference-in-differences (DiD) regressions. Columns (1) and (2) do so for the subjective quit probability under the normal and high UI scenarios s of question `FRKC12` and estimate the following regression:

$$P(EE)_{is} = \alpha_i + \tau^P T_s + \varepsilon_{is}, \tag{B.3}$$

where T_s indicates the high-UI scenario, so τ^P captures the within-person effect of UI generosity on the subjective EE probability. Columns (3) and (4) stack the post-UI-information and UI-increase

Table B.3: The tradeoff between pay and job safety across specifications

Panel A: All workers

	(1) Baseline	(2) Non-monotonic answers incl.	(3) Incl. partial ICs	(4) Speeders incl.	(5) 1% winsor	(6) 2% trim
$\widetilde{\delta}_j$	1.554***	1.550***	1.561***	1.462***	1.628***	1.438***
	(0.066)	(0.065)	(0.066)	(0.060)	(0.079)	(0.057)
N (job offers)	5040	5395	5074	6075	5040	4994
Within R^2	0.340	0.323	0.339	0.311	0.279	0.369
N (respondents)	1008	1079	1017	1215	1008	1006
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Full-time only	No	No	No	No	No	No

	(7) Post-UI info	(8) w_i^* anchor	(9) Endpoints increasing	(10) Weakly increasing	(11) Strictly increasing
$\widetilde{\delta}_j$	1.584***	2.027***	1.992***	1.671***	2.437***
	(0.064)	(0.107)	(0.075)	(0.067)	(0.138)
N (job offers)	5055	5025	4280	4780	1540
Within R^2	0.360	0.249	0.427	0.375	0.485
N (respondents)	1011	1005	856	956	308
Individual FE	Yes	Yes	Yes	Yes	Yes
Full-time only	No	No	No	No	No

Panel B: Full-time workers

	(1) Baseline	(2) Non-monotonic answers incl.	(3) Incl. partial ICs	(4) Speeders incl.	(5) 1% winsor	(6) 2% trim
$\widetilde{\delta}_j$	1.632***	1.624***	1.641***	1.556***	1.725***	1.490***
	(0.074)	(0.074)	(0.075)	(0.069)	(0.091)	(0.064)
N (job offers)	4135	4410	4161	4930	4135	4097
Within R^2	0.352	0.333	0.350	0.327	0.285	0.385
N (respondents)	827	882	834	986	827	826
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Full-time only	Yes	Yes	Yes	Yes	Yes	Yes

	(7) Post-UI info	(8) w_i^* anchor	(9) Endpoints increasing	(10) Weakly increasing	(11) Strictly increasing
$\widetilde{\delta}_j$	1.665***	2.145***	2.028***	1.743***	2.516***
	(0.073)	(0.124)	(0.083)	(0.076)	(0.150)
N (job offers)	4165	4125	3590	3940	1345
Within R^2	0.371	0.255	0.427	0.384	0.492
N (respondents)	833	825	718	788	269
Individual FE	Yes	Yes	Yes	Yes	Yes
Full-time only	Yes	Yes	Yes	Yes	Yes

Notes: This table reports estimates of Equation (1) for various samples. Column (1) is the preferred sample; every other column changes one ingredient relative to it, as described in the text. Standard errors clustered at the respondent level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

reservation-pay elicitations s at the offer level and estimate

$$\widetilde{r}_{ijcs} = \alpha_{ic} + \beta\,\widetilde{\delta}_j + \tau^{\beta}\left(\widetilde{\delta}_j \times T_s\right) + \varepsilon_{ijcs}, \tag{B.4}$$

where T_s again indicates the UI-increase elicitation and α_{ic} are worker-by-elicitation fixed effects. Because both elicitations follow the UI information, τ^{β} isolates the generosity effect on the IC slope, net of the information effect. Columns (2) and (4) keep respondents whose answers differ from their baseline ones.

Table B.4: UI generosity treatment: subjective EE probability and IC slopes

	(1) $P(EE)_{is}$	(2) $P(EE)_{is}$	(3) $\widetilde{r}_{ijc}$	(4) $\widetilde{r}_{ijc}$
UI generosity treatment	3.345***	7.378***		
	(0.575)	(1.234)		
$\widetilde{\delta}_j$			1.623***	1.789***
			(0.074)	(0.146)
$\widetilde{\delta}_j \times$ UI generosity treatment			−0.035*	−0.108*
			(0.020)	(0.063)
N (observations)	1460	662	7630	2460
Within R^2	0.044	0.098	0.367	0.355
N (respondents)	730	331	763	246
Fixed effects	Worker	Worker	Worker × elicitation	Worker × elicitation
Revised answer only	No	Yes	No	Yes
Full-time only	Yes	Yes	Yes	Yes

Notes: This table reports the UI generosity treatment effects. Columns (1) and (2) report estimates of Equation (B.3) on full-time workers in the preferred sample with quit-probability answers before and after the hypothetical UI generosity increase. Columns (3) and (4) report estimates of Equation (B.4). Columns (2) and (4) drop respondents who did not revise their respective answers. Standard errors clustered at the respondent level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Quantitative model appendix

C.1 Numerical implementation: state space and grids

We solve the model on a discretized state space. Employed workers are indexed by assets a, match productivity z, and unemployment risk δ. Unemployed workers are indexed by assets a alone. We discretize each dimension as follows:

- **Assets:** $n_a = 50$ points equally spaced on $[\underline{a}, \bar{a}]$, with the upper bound $\bar{a} = 7.76$ chosen so that at most 1% of the mass in the ergodic distribution lies at the upper bound.
- **Match productivity and unemployment risk:** $n_z = 20$ and $n_\delta = 15$ points, respectively, equally spaced in logs. The job-offer distributions over z and δ are independent lognormals. Each grid is truncated at ±3 standard deviations around the mean of its offer distribution.

We also discretize tenure to calculate tenure-related moments, as described in Appendix C.5.

C.2 The price of risk as a near-sufficient statistic for IC slopes

In Sections 3.2 and 4.3 we argue that the price of risk (PR) captures an individual's IC slope: how much extra pay a worker demands for higher unemployment risk, referring to an assessment that we now detail. Specifically, we show that PR predicts the model's IC slopes almost exactly. Regressing the latter on the former yields a coefficient of 0.99 ($R^2 = 0.995$).[16] Hence, PR is not just a local approximation but a near-sufficient statistic for the wage-safety tradeoff we elicit in the survey. Figure C.1 reports this main result, which we detail below.

The benchmark. Evaluating Equation (8) at the worker's own job ($\widetilde{\delta} = \delta$ and $\widetilde{z} = z$) and assuming that the worker has reached her job-specific optimal asset position, i.e., $\dot{a} = 0$, yields:

$$\left.\frac{\partial \widehat{w}^r(a,z,\delta,\widetilde{\delta})}{\partial \widetilde{\delta}}\right|_{\widetilde{\delta}=\delta,\ \dot{a}=0} = \frac{W(a,z,\delta) - U(a)}{u'\big(c^*(a,z,\delta)\big)} = \mathrm{PR}(a,z,\delta), \tag{C.1}$$

where $\widehat{w}^r(a,z,\delta,\widetilde{\delta})$ is the net reservation wage for a worker with characteristics (a,z,δ) facing a job offer with unemployment risk $\widetilde{\delta}$—which we characterize in Section 3.2. Two features break this exact equality in the baseline model when moving to the ICs that actually determine worker choices. First, workers away from their optimal asset position have $\dot{a} \neq 0$. Second, the IC slopes we estimate (both in the simulated and survey data) are not derivatives (i.e., they are not "local" slopes) but stem from the "global" comparison of far-away jobs that is encoded in Equation (6). (Intuitively, local HtM slopes are linear, global ones are curved absent risk neutrality.) As a case in point, the regression specification in Equation (1) fits a line across five risk levels spanning 2 to 50 percent per year, most of them far from the worker's own job. This second aspect gives the IC curvature and implies that the price of risk has no exact global counterpart. In our main result, Figure C.1, we find that neither departure matters quantitatively.

Variable construction. We now describe how we construct the variables plotted in Figure C.1. The underlying observations are each (a,z,δ) node—denoted by i—in the state space. For each node, we compute the precise reservation wage at each of the five risk levels considered in our

[16] We note a computational caveat and attenuation stemming from extrapolation outside of the wage grid below.

survey. PR follows from Equation (7). We form the reservation pay ratio by dividing reservation wages by the same anchor reservation wage at the job with 10% separation risk (using either the net reservation wage $\widehat{w}_{i,10\%}$ or gross reservation wage $w_{i,10\%}$ based on the application). To construct global IC slopes, we estimate worker-specific IC slopes, β_i, by running a model analog of the regression in Equation (1) worker by worker. To compare this slope to the PR, we normalize PR by the same anchor reservation wage, $\text{PR}_{i,10\%} = \text{PR}_i/\widehat{w}_{i,10\%}$. Then, we can interpret the price of risk as the value of the current job over unemployment, expressed in months of pay at the wage the worker would accept for a job with 10 percent annual(ized, see below) unemployment risk. (To preview, the variant in Panel (a) applies, and needs, no such normalization.) Each dot in the panels is one (a, z, δ) node, sampled in proportion to employment mass, and all best-fit lines are weighted by that mass. We construct these concepts at the model frequency (monthly) in Panels (a) and (b), and annualize them when moving to the survey analog in Panel (c), a transition we detail below.

The basic idea of the validation check. Our check is simple and intuitive: if PR predicts worker IC slopes accurately, we expect this relationship to trace out a 45 degree line in each panel, which we plot as a benchmark. Departures from this line are captured by the best-fit line featuring a different intercept and slope, and a second useful summary statistic of fit is the R^2.

Local IC slopes. Panel (a) assesses the approximation for local IC slopes in Equation (C.1). It gauges how the PR—which ignores asset terms—captures the actual local IC slope $\partial\widehat{w}_i^r/\partial\widetilde{\delta}$ in Equation (8), both expressed in levels of net pay per unit of the model's *monthly* separation hazard. The fitted slope is 0.98 with an intercept of 0.32 and an R^2 of 0.995, and the scatter plot is highly concentrated around the 45-degree line. The implication of this first check is that the asset-drift terms, $\dot{a}$, in Equation (8) are quantitatively negligible in our calibration.

Global IC slopes. The remaining two panels assess ICs globally, rather than locally. Panel (b) replaces the local derivative with the slope of the reservation-wage schedule across the five elicited risk levels, β_i^{net}, using the net-pay and monthly-hazard rates as in Panel (a) but plotting it against $\text{PR}_{i,10\%} = \text{PR}_i/\widehat{w}_{i,10\%}$. The tight fit remains, with a slope of 0.95, an intercept of 0.15, and an R^2 of 0.996. Hence, PR summarizes not only how a worker prices a marginal change in risk around her own job—"locally"— but also "globally": how she prices the whole range of hypothetical jobs, such as the wide range faced by survey respondents. The samples differ across the panels. Panel (a) uses every state, whereas Panels (b) and (c) include reservation wages that—which for workers in safe, high-paying jobs—might lie beyond the wage grid on which we solve the model. We therefore fit the line in Panels (b) and (c) on the states that remain on the grid, drawn in blue, and hold out the extrapolated states, drawn in gray, which carry 57.5 percent of employment.[17] We elaborate on this technical consideration at the end of this section.

From model units to survey units. Panel (c) presents Panel (b) in units consistent with our custom survey measurement (the empirical treatment built around Equation (1)), which relates percent changes in *gross* pay to percentage point changes in *annual* unemployment risk. This is also how we treat the model and survey data in their juxtaposition in the main text. We reconcile the PR concept with the survey data with two adjustments detailed below, obtaining $\widehat{\pi} \cdot \text{PR}^g_{i,10\%}$, where $\text{PR}^g_{i,10\%} = \frac{w_i^\Upsilon}{\tau(1-\Upsilon)} \cdot \text{PR}_{i,10\%}$.

[17] The fit remains tight once we include them. The R^2 falls only from 0.996 to 0.93 in Panel (b) and from 0.995 to 0.94 in Panel (c).

The first adjustment aligns time units. Our survey elicits the probability of job loss over twelve months, whereas the model's separation rate is a monthly hazard; we link the two with continuous compounding, i.e., through $\widetilde{\delta}^{Y} = 1 - e^{-12\widetilde{\delta}}$. This implies $d\widetilde{\delta}^{Y}/d\widetilde{\delta} = 12(1 - \widetilde{\delta}^{Y})$. Since the IC slope coefficient β_i is fitted across the five elicited risk levels rather than being a derivative at a point, we adjust it by the factor:

$$\widehat{\pi} \;=\; \frac{\mathrm{Cov}\big(\widetilde{\delta}_j,\, \widetilde{\delta}_j^{Y}\big)}{\mathrm{Var}\big(\widetilde{\delta}_j^{Y}\big)} \;=\; 0.118, \tag{C.2}$$

where $\widetilde{\delta}_j = -\frac{1}{12}\log(1 - \widetilde{\delta}_j^{Y})$ is the monthly hazard implied by annual risk level j. Econometrically, $\widehat{\pi}$ is the first-stage coefficient from projecting the monthly hazard on annual risk across the five risk levels. Equation (1) is then the reduced form, and $\beta_i/\widehat{\pi}$ is exactly the IV coefficient obtained by regressing the reservation pay ratio on the monthly hazard, using annual risk as the instrument.

Our second adjustment aligns the wage concepts. Equation (8) is expressed in net pay $\widehat{w} = \tau w^{1-\Upsilon}$, whereas our survey uses gross pay. Differentiating the tax function gives $d\widehat{w}/dw = \tau(1 - \Upsilon)w^{-\Upsilon}$. Then the gross counterpart of the price of risk, moreover normalized by the reservation wage at the 10% unemployment-risk job, is

$$\mathrm{PR}^{g}_{i,10\%} \;=\; \frac{w_i^{\Upsilon}}{\tau(1-\Upsilon)} \cdot \frac{\big[W(a,z,\delta) - U(a)\big]}{\widehat{w}_{i,10\%}\, u'(c_i)} = \frac{w_i^{\Upsilon}}{\tau(1-\Upsilon)} \cdot \mathrm{PR}_{i,10\%}. \tag{C.3}$$

Here the term $w_i^{\Upsilon}/[\tau(1-\Upsilon)]$ maps net into gross pay.

The price of risk predicts the IC slopes. Adjusting by these two factors, Panel (c) presents a slope of 0.99 with an intercept of 0.01 and an R^2 of 0.995. Hence, the IC slopes we estimate in the model, which average 1.53 vs. 1.63 in our survey, are predicted almost exactly by a single statistic.

The wage grid. In Panels (b) and (c), the reservation-wage schedule toward the highest elicited risk level of 50% per year requires, for workers in high-paying and safe jobs, a value beyond the range of the wage grid on which we solve the model. Such workers' IC slopes rely on extrapolation. These states are shown in gray and constitute around 57.5 percent of employment. To obtain a statistic of fit that sidesteps this technical discrepancy, for this exercise, we estimate the fitted line on the non-extrapolated states and report two R^2s, one excluding and one including those states. The R^2s remain high as they fall only from 0.996 to 0.93 and from 0.995 to 0.94, respectively. Hence, the theoretical benchmark holds wherever we solve the model on the grid without extrapolation. This extrapolation is not relevant to the rest of our quantitative analysis elsewhere in the paper.

C.3 Smoothing the discrete acceptance decisions

We describe a technical detail suppressed in the main text that amounts to smoothing the discrete choices; as we explain at the end of this section, this feature does not have quantitatively important implications for our model. The job acceptance indicators $\mathbf{1}^{u}$ and $\mathbf{1}^{e}$ in Equations (12)–(13) are binary, making $U(a)$ and $W(a,z,\delta)$ non-differentiable in parameters at the acceptance threshold. We smooth each decision by adding i.i.d. taste shocks $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ to the offered match value before the comparison, so that a worker accepts $(\widetilde{z}, \widetilde{\delta})$ when $W(a,\widetilde{z},\widetilde{\delta}) + \epsilon$ weakly exceeds their current value. The binary indicators are then replaced by smoothed acceptance probabilities

$$p^{u}(\widetilde{z},\widetilde{\delta}|a) = \Phi\!\left(\tfrac{W(a,\widetilde{z},\widetilde{\delta}) - U(a)}{\sigma_\epsilon}\right) \quad \text{and} \quad p^{e}(\widetilde{z},\widetilde{\delta}|a,z,\delta) = \Phi\!\left(\tfrac{W(a,\widetilde{z},\widetilde{\delta}) - \max\{W(a,z,\delta),U(a)\}}{\sigma_\epsilon}\right), \tag{C.4}$$

Figure C.1: The price of risk as a sufficient statistic for indifference-curve slopes

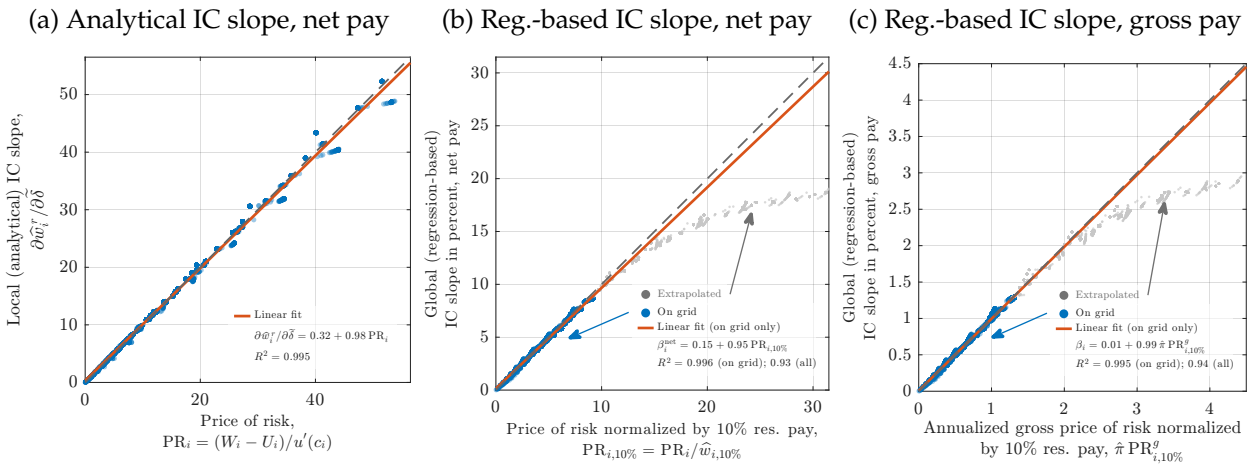

Notes: This figure plots each employed worker's IC slope against her price of risk (PR), across the ergodic distribution of employed workers. Each dot is one (a, z, δ) node sampled in proportion to that distribution, with darker markers arising from overlapping draws and thus indicating states with greater employment mass. The red line is the linear fit, whose equation and R^2 are reported in the legends. Panel (a) plots the local IC slope $\partial \widehat{w}_i^r / \partial \widetilde{\delta}$, the derivative of the net reservation wage evaluated at the worker's own separation rate, against the price of risk defined in Equation (7), both in levels of net pay per unit of the monthly separation hazard. Panel (b) changes both axes. On the y-axis, it replaces that local derivative with β_i^{net}, the respondent-specific slope estimated in a linear regression of her reservation wage on the five elicited risk levels. The underlying dependent variable retains the net-pay and monthly-hazard units of Panel (a) but divides it by the reservation (net) pay at 10 percent annual risk. Correspondingly, the price of risk on the x-axis is also normalized by $\widehat{w}_{i,10\%}$. Panel (c) restates Panel (b) in the units of Equation (1), plotting the estimated β_i against $\widehat{\pi}\,\mathrm{PR}^g_{i,10\%}$. It departs from Panel (b) in that $\mathrm{PR}^g_{i,10\%}$ is the gross-pay analog of the price of risk, defined in Equation (C.3), and $\widehat{\pi}$, defined in Equation (C.2), annualizes the monthly separation hazard. In Panels (b) and (c), computing the reservation-wage schedule at the highest elicited risk level requires, for the states shown in gray, a reservation wage beyond the wage grid on which the model is solved. The fitted line in those panels is estimated on the non-extrapolated grids, and the two reported R^2 are those of the regression excluding and including the extrapolated states. See Appendix C.2 for details.

where Φ is the standard normal CDF. Both terms collapse to $\mathbf{1}^u$ and $\mathbf{1}^e$ as $\sigma_\epsilon \to 0$. The expected-max identity, $\mathbb{E}_\epsilon[\max\{x+\epsilon, 0\}] = x\,\Phi(x/\sigma_\epsilon) + \sigma_\epsilon\,\phi(x/\sigma_\epsilon)$, then implies that each smoothed HJB acquires an additional non-negative option-value term of the form $\lambda(\theta)\,\sigma_\epsilon \sum_{\widetilde{z},\widetilde{\delta}} \pi(\widetilde{z}, \widetilde{\delta})\,\phi((W-V)/\sigma_\epsilon)$, where V is the relevant outside option. We implement these terms in the numerical solution so that the smoothed problem remains internally consistent at any $\sigma_\epsilon > 0$ and recovers the binary HJBs in Equations (3)–(4) as $\sigma_\epsilon \to 0$. We additionally allow employed workers the option to quit into unemployment from jobs that have a lower value than unemployment (due to i.i.d. taste shocks and the aggregate shifts in the business-cycle experiments of Section 7), and smooth this discrete quit choice identically to the job acceptance decisions. This quit option arrives at rate $\zeta_q = 1$, implying a frequency of once per month on average.

The smoothing parameter σ_ϵ is measured in value units, so we scale it to the resolution of the grid. We set σ_ϵ equal to c times the mean absolute difference in employment value W between adjacent nodes of the job-offer grid, averaged over the z and δ dimensions, i.e., the typical change in the value of a job when the offer moves by one grid step. With $c = 0.3$, the taste shock is roughly a third of a grid step, and σ_ϵ shrinks automatically as the grid is refined, so the smoothed model converges to the exact model in the limit. In the baseline calibration this gives $\sigma_\epsilon = 0.15$. This smoothing matters for our counterfactual exercises. With binary

acceptance on a discrete grid, a small change in a parameter, such as the UI benefit in Figure 8, moves entire cells of offers across the acceptance threshold at once, so aggregate transition rates respond in discrete jumps and the plotted comparative statics might become jagged. The smoothed probabilities instead shift acceptance mass continuously with the value gap, while leaving the calibrated equilibrium essentially unchanged. As evidence that its effect is mild, replacing the smoothed acceptance probabilities with the exact, un-smoothed optimal rule—holding fixed the calibrated value functions and stationary distribution—changes the aggregate job-finding rate by under half a percent and the EE rate by about three percent, and the taste shocks generate only a negligible flow of voluntary quits into unemployment, accounting for under one percent of all separations. The smoother therefore preserves the model's economic content and its fit to the calibration targets.

C.4 Wage determination

Wages are continuously renegotiated to overcome the issue of on-the-job search giving rise to a non-convex bargaining set (Shimer, 2006).[18] We also assume that the threat to terminate the match in the absence of a wage agreement is not credible whenever the surplus remains positive, as argued by Hall and Milgrom (2008). Hence, bargaining happens over the instantaneous flow surplus. To this end, we construct a two-stage non-cooperative bargaining game that yields a simple wage rule building on the insights by Gottfries (2026) and Elsby and Gottfries (2022).[19]

At each instant dt the worker prepares two offers w_{W1} and w_{W2}, for stage 1 and stage 2, respectively. The firm anticipates all possible stage-1 offers and commits to an acceptance rule and also commits to its counter-offer w_F in case it rejects w_{W1}. The worker's second-stage offer and the firm's counter-offer are take-it-or-leave-it, and with probability ψ it is w_{W2} that is selected.[20] In case of no agreement in stage 1, the worker receives a strike payment χ, whereas the firm gets nothing. If stage 2 ends in disagreement, no output is produced and the worker receives χ while the firm receives nothing. The parties then return to their outside options.[21]

With these disagreement payoffs, backward induction implies the following stage-2 offers. If the worker starts, she chooses w_{W2} to maximize her flow payoff, subject to the firm weakly preferring acceptance over disagreement, setting $w_{W2} = z$. If the firm starts, it chooses the smallest wage the worker would accept, which is $w_F = \chi$. Hence, the expected stage-2 payoffs are $\psi z + (1-\psi)\chi$ and $(1-\psi)(z-\chi)$ for the worker and firm, respectively.

Given these stage-2 continuation values, the firm accepts a stage-1 offer w_{W1} whenever $z - w_{W1} \geq (1-\psi)(z-\chi)$. For the worker, any $w_{W1} \geq \psi z + (1-\psi)\chi$ is at least as good as the expected stage-2 payoff. Therefore, the unique wage that both parties accept in stage 1 is $w = \psi z + (1-\psi)\chi$, which the worker proposes in equilibrium, resulting in Equation (16).

[18] We keep the wage structure deliberately simple to focus on how risk affects job search. How risk interacts with and shapes the wage structure is left for future research.

[19] We thank Axel Gottfries for conversations which were instrumental in designing this wage rule.

[20] Equivalently, each side submits two labeled envelopes: the worker lists w_{W1} for the first and w_{W2} for the second period, while the firm provides a yes/no for each w_{W1} in the first and lists w_F in the second period.

[21] The worker's payoff is evaluated in monetary terms, as if she were risk neutral. Since this game occurs and is fully resolved within an instant dt, the worker's risk aversion does not affect her decisions: she simply maximizes her expected payout. Intuitively, the worker's value function is $W(a, z, \delta)$, and in the game they bargain over an instantaneous, risky payment that raises assets from a to $a + w\,dt$. As $dt \to 0$ this payment becomes infinitesimal, so the worker's marginal value of wealth remains $W_a(a, z, \delta)$ regardless of the second-stage outcome. She therefore values the gamble at its expected payout and behaves as if risk neutral. If instead we treat the worker as risk averse even within the instant of the game, the result only strengthens. In stage 1 they now *strictly*, rather than weakly, prefer the conjectured wage, since agreement avoids the risk over which second-stage offer is selected.

C.5 Tenure distributions

The EU-tenure profiles in Figure 4 Panel (b) and the corresponding calibration target require the joint distribution of worker characteristics by job tenure. We denote by $g^e(a, z, \delta, l)$ the mass of employed workers with characteristics (a, z, δ) and job tenure $l \geq 0$, measured in months since the job began. Tenure increases while the worker remains in her job and is zero for any new hire, whether out of unemployment or through an EE move. We compute g^e with the same non-stochastic simulation procedure used for the ergodic distribution. Specifically, we start from the distribution of new hires $g^e(a, z, \delta, 0)$, we discretize tenure using 15 quadrature nodes over $[0, 50]$ months, and iterate forward using the portion of the transition matrix that captures within job movements in assets, with EU and EE separators removed from the mass of workers with a given tenure duration. The nodes are spaced so that they concentrate at short tenures, where the profile is steep. To match the monthly-frequency moments we construct from the CPS, we compute the EU rate by tenure, linearly interpolating between nodes and integrating $g^e(\cdot, l)$ over one-month tenure windows.

When we use tenure again elsewhere, namely to impose the three-year tenure restriction that defines the sample for the displacement analysis, we do so on the (stochastically) simulated panel described in Appendix C.7, using each worker's realized employment and transition history (rather than the quadrature nodes).

C.6 Consumption dynamics in the model

Marginal propensity to consume. We compute MPCs out of a discrete increase in income. To be consistent with the data we compute the MPC over a given horizon, e.g., total consumption increase over a quarter. Define $C^e(a, z, \delta, l)$ as the total expected consumption flow over the next l units of time of a worker who is currently employed with characteristics a, z, δ. Similarly define $C^u(a, l)$ for a currently unemployed worker with assets a. The quarterly MPC to a \$500 windfall for an employed worker is then given by $(C^e(a + 500/dlr, z, \delta, 3) - C^e(a, z, \delta, 3))/(500/dlr)$, where dlr is a conversion factor to convert dollars into model wage units, and similarly for an unemployed worker. To convert \$500 into model wage units we use the following approach. We take quarterly mean labor income in the data to be \$16,500. Quarterly mean labor income in the model is $3((1-u)E[w] + ub)$, giving us a conversion factor $dlr = 16{,}500/(3((1-u)E[w] + ub))$ such that x wage units in the model are equivalent to $dlr \times x$ expressed in dollars. This conversion factor is also used to compute the model-equivalent change in b for a \$100 UI increase.

In Figure C.2 we plot the MPC distribution. The quarterly MPCs to a \$500 windfall are plotted for the unemployed in Panel (a) and for the employed in Panels (b) and (c), holding the safest and riskiest jobs, respectively. The MPCs are decreasing in assets, highest for the unemployed, decreasing in wages, and slightly increasing in unemployment risk. The average (income-weighted) MPC is 19.3% for the employed, 46.7% for the unemployed, and 20.0% for the whole population. Hence, despite assets being low for most employed workers, they are far from being hand-to-mouth, as in that case their MPC would be 100%.

Consumption fall upon unemployment. Using Canadian data, Browning and Crossley (2001) find an average 14% decline in monthly consumption expenditures and a 25% income loss at unemployment as reported in Low et al. (2010). The mean UI replacement rate in their sample is 48%, so the fact that the income loss is only 25% represents substantial income insurance (for example from spousal income), which is not present in our model. For this reason—and to facilitate comparison with other countries where income losses may differ—Low et al. (2010) interpret these

Figure C.2: Quarterly MPCs to $500 windfall as a function of assets



Notes: This figure plots MPCs over one quarter to a $500 windfall. For the employed each line corresponds to a different productivity/wage. Safest and riskiest jobs correspond to the lowest and highest δ in the grid. High-wage jobs are at the highest z, median-wage at the median z, and low-wage at a low z near the minimum accepted z.

numbers as a 14/25 = 56% passthrough from income to consumption drop.[22] As Browning and Crossley (2001) measure monthly consumption, we can approximate this using our consumption function directly.[23] We measure the consumption fall upon unemployment for an employed worker as $dc(a,z,\delta) = (c^u(a) - c^e(a,z,\delta))/c^e(a,z,\delta)$. We measure the change in income as the UI level relative to post-tax wage: $dy(a,z,\delta) = (\tau b^{1-\Upsilon} - \tau w(z)^{1-\Upsilon})/(\tau w(z)^{1-\Upsilon})$. We compute the average of both of these numbers among laid-off workers, and find a passthrough of 62%.[24]

C.7 Cost of job loss: estimator and simulation

Here, we detail the event study plotted in Figure 6. We simulate a panel of workers from the calibrated steady state, apply a Callaway and Sant'Anna (2021) style group-time average treatment effect estimator adapted to our setting, and report the resulting profiles for earnings, log earnings, log wage, and the EU rate. This estimator has the advantage of being transparent and easy to implement, and we have verified that our earnings results are robust to using the original specification of Davis and von Wachter (2011).

Simulated panel. Starting from the ergodic distribution, we simulate $N = 50,000$ workers for $T = 50$ years with a time step Δt implying 48 nodes per year. Each worker is simulated forward using their policy and transition functions, and we then annualize the panel into a $[N \times T]$ array of annual earnings (the sum of wages over months employed), the average wage while employed, and three event counters $n^{EU}_{i,t}$, $n^{UE}_{i,t}$, $n^{EE}_{i,t}$ recording the number of separations into unemployment, unemployment-to-employment, and employer-to-employer transitions experienced by worker i during year t.

Sample and groups. At each reference year y, we restrict attention to workers who pass a three-year tenure rule: number of employed months = 12 in each of years $y-3$, $y-2$, $y-1$ and

[22] Other studies find similar numbers, as discussed in den Haan et al. (2018).

[23] In reality, their data compares people before they were unemployed to their response at an interview six months after becoming unemployed. The 14% consumption loss is the change in monthly consumption expenditures, i.e., $(C_u - C_e)/C_e = -14\%$, where C_e and C_u are total monthly consumption expenditures.

[24] The average consumption change is $\overline{dc} = \int dc(\cdot)\delta g^e(\cdot)dadzd\delta / \int \delta g^e(\cdot)dadzd\delta$, with $\overline{dy}$ defined analogously.

$n^{EU} = n^{UE} = n^{EE} = 0$ in each of those years.[25] In this sub-sample, the *treated* group is those workers who experience at least one EU event at y ($n^{EU}_{i,y} \geq 1$), and the *control* group is those who do not ($n^{EU}_{i,y} = 0$). The control group purposefully does not exclude workers who experience an EU event at some $t > y$ and so those are included in Figure 6 Panel (c). The left-hand side of the regression for Panel (c) is the EU separation indicator $\mathbf{1}\{n^{EU}_{i,t} \geq 1\}$.

Estimator. For each year y and outcome Y, the group-time difference at event time k is

$$\widehat{\text{ATT}}(y,k) = \overline{Y}_T(y,k) \; - \; \overline{Y}_C(y,k), \tag{C.5}$$

where $\overline{Y}_T(y,k)$ and $\overline{Y}_C(y,k)$ are the sample means of Y at calendar year $t = y + k - 1$ for the treatment and control groups, respectively. We then anchor each cohort to the three-year pre-period window $k \in \{-2,-1,0\}$ by subtracting its sample mean from each k. This avoids the need for worker fixed effects since the simulated panel is balanced by design. The pooled estimator averages the anchored per-cohort estimates with equal weight across a range of reference years, $\mathcal{Y}_{\text{bal}} = \{y : \max(2 - k_{\min}, T_{\text{ten}} + 1) \leq y \leq T - k_{\max} + 1\}$, ensuring that every y in the average contributes a fully observed event window. With the panel and event window used in Figure 6, this gives 24 reference years. The IRFs in Figure 6 plot $\widehat{\text{ATT}}(y,k)$ with k on the x-axis. Workers may appear in multiple cohorts as either treatment or control. For example, a worker who makes an EU transition at y_1 and re-builds three years of tenure by some $y_2 > y_1 + 3$ is in both cohorts.

C.8 The distribution of indifference-curve slopes in the model and the data

Figure C.3 compares the distribution of individual IC slopes β_i in the model and the survey, as a histogram in Panel (a) and as a cumulative distribution function in Panel (b). The model matches the average IC slope (1.53 in the model, 1.63 in the survey), and the interquartile ranges are reasonably close (0.87–2.17 in the model, 0.36–2.13 in the survey). The empirical distribution is, however, more dispersed (standard deviation of 2.14 versus 0.75 in the model), strongly right-skewed, and includes a small share of negative slopes (4.7% of respondents) that the model cannot generate. We attribute this additional dispersion to measurement error in the individual survey responses. The model measures each worker's IC slope exactly, whereas the survey elicits it with noise, which fattens the tails and produces economically implausible negative values.

C.9 UI in general equilibrium

Asset market clearing. Our results are robust to how the asset market clears. The main text presents the case where the government adjusts A_g to keep r fixed at a 2% annual rate. Table C.1 compares this assumption to the alternative in which the government holds A_g fixed and r adjusts to clear the market, under ±25% changes in the UI benefit level b relative to baseline.

The two market-clearing assumptions deliver similar outcomes for productivity ($\Delta Y/n$), the unemployment rate, and the worker flows (EE / EU / UE). The main difference is in how the financing of UI is reflected in the asset market: when A_g adjusts, r is held at 2% annual and A_g moves between −15.2% and 9.6%; when A_g is held fixed, the interest rate moves to clear the market—a 4.2 pp drop under the 25% UI cut and a 2.82 pp rise under the 25% UI increase. We view these interest-rate movements as implausibly large given the modest size of the UI change, reflecting that the model omits other sources of saving.

[25] This is the model analog of the restriction in Davis and von Wachter (2011) to more than three years of prior tenure with the same main employer.

Figure C.3: Distribution of individual indifference-curve slopes: model versus survey

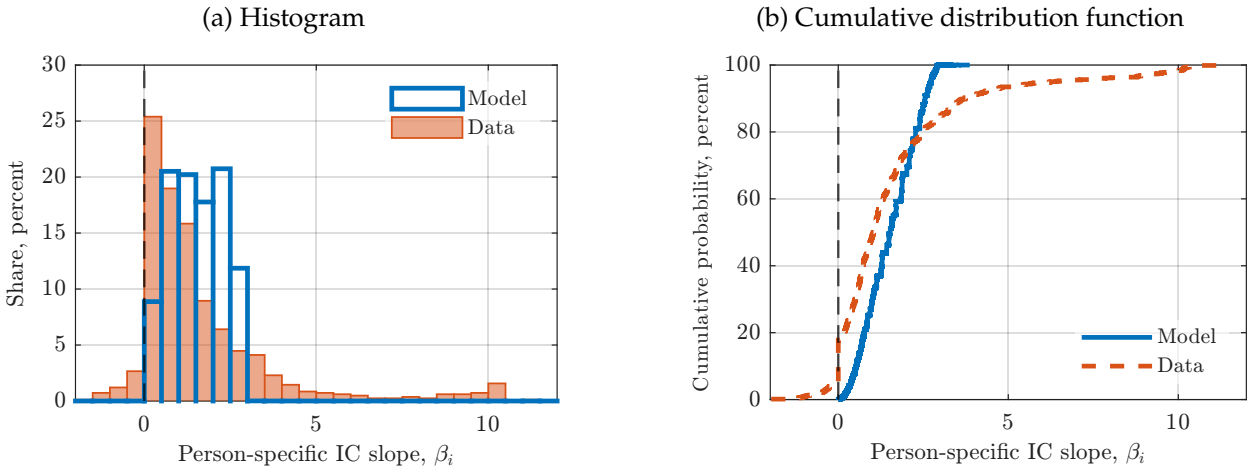

Notes: This figure plots the distribution of individual IC slopes β_i—the coefficient from Equation (1) estimated worker by worker—across employed workers in the model, weighted by the ergodic distribution, and across respondents in the survey. Panel (a) is a histogram over the full common support. Panel (b) is the cumulative distribution function.

Table C.1: ±25% UI changes under different asset market clearing assumptions

Approach	Δr	ΔA_g	Δ Avg. tax rate	$\Delta Y/n$	Δ u rate	Δ EE rate	Δ EU rate	Δ UE rate	$\Delta\lambda$
Adjust A_g, fix r (−25% UI)	—	−15.15%	−0.43pp	−0.34%	−0.34pp	−0.08pp	−0.08pp	0.22pp	−0.45%
Adjust A_g, fix r (+25% UI)	—	9.59%	0.81pp	0.53%	0.89pp	−0.03pp	0.07pp	−1.97pp	0.36%
Fix A_g, adjust r (−25% UI)	−4.20pp	—	0.36pp	−0.42%	−0.39pp	−0.11pp	−0.09pp	0.18pp	−0.60%
Fix A_g, adjust r (+25% UI)	2.82pp	—	0.28pp	0.55%	0.90pp	−0.01pp	0.08pp	−1.86pp	0.45%

Notes: This table presents results from changing UI by ±25% under two asset market clearing assumptions. The top panel holds r fixed at 2% annual and lets A_g adjust. The bottom panel holds A_g fixed and solves for the market-clearing r. Δr is the percentage-point change in the annualized interest rate. ΔA_g is the percent change in government savings. $\Delta Y/n$ is the percent change in aggregate labor productivity. The remaining columns (Avg. tax rate, u, EE, EU, UE) report percentage-point changes.

C.10 Calibration of the auxiliary models

The general-equilibrium UI counterfactuals in Section 6.2 (Table 4 and Figure 8) compare the baseline model against two recalibrated auxiliary models: a *risk-neutral* economy, in which workers have linear utility ($\sigma = 0$), eliminating precautionary asset accumulation and risk aversion; and a *homogeneous-δ* economy, in which all workers face the same unemployment risk. We re-estimate each auxiliary model to match the same empirical targets (where relevant) as the baseline. Table C.2 reports the resulting parameters. These recalibrated economies are distinct from the risk-neutral and hand-to-mouth counterfactuals in Section 5, which alter preferences while holding all remaining parameters at their baseline values.

C.11 Cross-sectional anatomy of the cyclical climb to safety

In Section 7 we study the aggregate properties of the "climb to safety." We now describe its cross-sectional features upon impact, analyzing which workers reroute their EE moves toward safety. Figure C.5 plots, across cells of the state space, the percentage change in the IC slope against

Figure C.4: Effect of varying UI benefits in general equilibrium

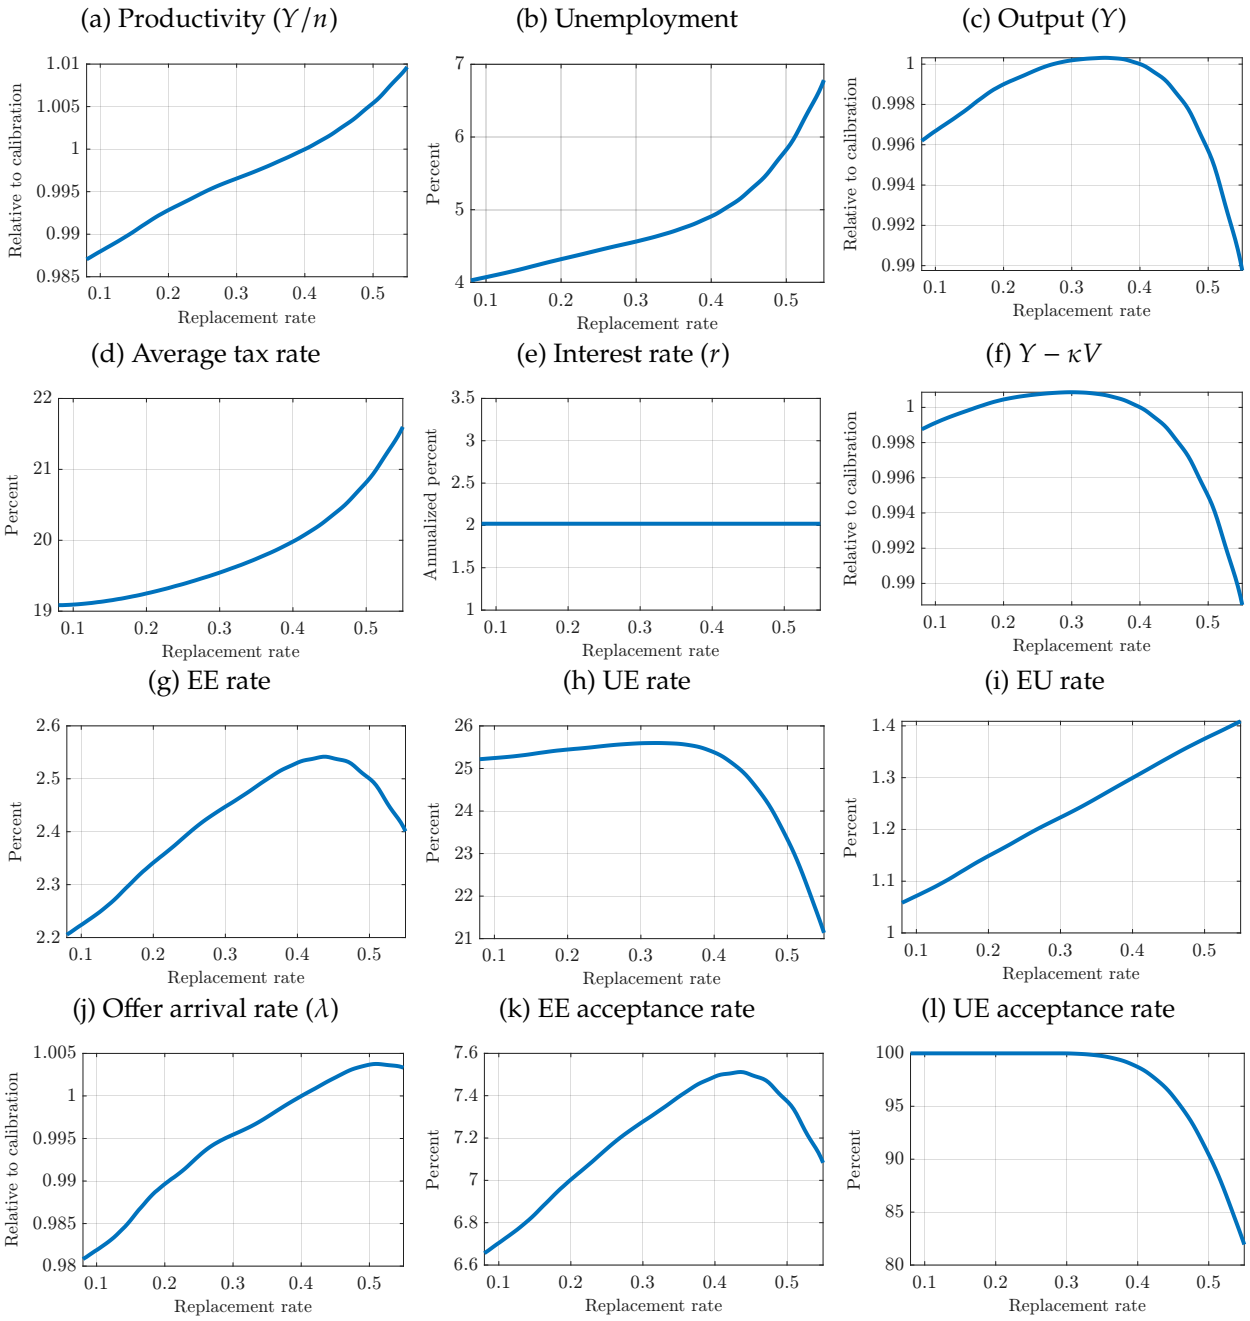

Notes: This figure plots comparative statics with respect to UI benefits in the baseline model, displayed as an average replacement rate. Each panel is plotted either in deviations from the calibrated steady state or in levels.

Table C.2: Calibrated parameters: baseline and auxiliary models

Parameter	Description	Value			Source/Target
		Baseline	Risk-neutral	Homog. δ	
Pre-determined:					
ρ_f	Firm discount rate	0.0043	0.0043	0.0043	5% annual
σ	Risk aversion	2	0	2	Standard
Υ	Tax progressivity	0.18	0.18	0.18	Heathcote et al. (2017)
η	Matching elasticity	0.5	0.5	0.5	Standard
ψ	Bargaining weight	0.9	0.9	0.9	Profit share of 0.1
χ	Strike payment	0	0	0	See text
$\rho_{z,\delta}$	Correlation EU risk and productivity	0	0	0	See text
Estimated:					
ρ	Discounting	0.0130	0.0130	0.0183	Average quarterly MPC
$\underline{a}$	Borrowing limit	−4.240	—	−4.027	1× quarterly income (Kaplan et al., 2018)
b	Unemployment benefit	0.566	0.561	0.538	Average replacement rate
s_e	Search efficiency on the job	1.31	1.27	0.74	Average EE rate
μ_δ	Unemployment risk mean parameter	−3.73	−3.80	−4.34	Average EU rate
σ_δ	Unemployment risk std parameter	0.954	1.014	—	EU tenure slope
μ_z	Mean of marginal density of z_0	−0.031	−0.031	−0.031	Normalize average z_0 draw
σ_z	Std. dev. of marginal density of z_0	0.285	0.274	0.271	Mean-Min ratio
κ	Vacancy posting cost	0.085	0.086	0.160	Unemployment rate
α	Match efficiency	0.257	0.251	0.289	Normalize labor market tightness θ
τ	Linear tax shifter	0.85	0.85	0.84	Personal income tax
A_g	Government saving	3.346	—	3.467	Asset market clearing → annual r
G	Wasteful gov. spending	0.252	0.245	0.239	Government budget balance (incl. debt)

Notes: This table reports the calibrated parameter values for the baseline model and the two auxiliary models used in the counterfactuals, each estimated to match the same empirical targets. A dash ("—") marks a parameter that does not apply to a given model. The risk-neutral model sets $\sigma = 0$ and makes workers hand-to-mouth, so the borrowing limit $\underline{a}$ and government savings A_g do not apply, and the discount rate ρ is fixed at its baseline value, since a quarterly-MPC target is not well defined under risk neutrality. The homogeneous-δ model has a single unemployment-risk type, so the risk-dispersion parameter σ_δ does not apply.

three outcomes. Panel (a) shows that the cells with the largest IC-slope rise see the largest fall in their EE rate (steady-state-weighted correlation −0.21). Panels (b) and (c) reveal the selectivity pattern behind the aggregate climb to safety: the accepted EE moves out of the more-affected cells tilt toward high-productivity destinations (correlation +0.25) and toward safer destinations (correlation −0.94). The cells whose IC slope steepens most are precisely those that move toward safety, as the IC slope predicts.

Table C.3 details which workers drive this reallocation, and in which direction. The workers who become *more* selective when the recession hits—those whose job-acceptance probability falls—make up roughly two-thirds of employment. They are disproportionately in safe, low-paying jobs (average annualized EU risk of 8.3%, against 24.8% for the less-selective group), are close to the borrowing constraint (half are exactly at it), and hold few assets. These are the workers with the most to lose from job loss and the lowest ability to self-insure, and they respond to the recession by clinging to safety. By contrast, the workers who become *less* selective are better-buffered households in riskier, higher-paying jobs, who accept a wider range of offers in order to leave their risky jobs. Both groups' responses to the recession exhibit the same price-of-risk force at the core of the paper, as both tilt the composition of EE moves toward safety.

Figure C.5: Cross-cell heterogeneity at the impact of the recession

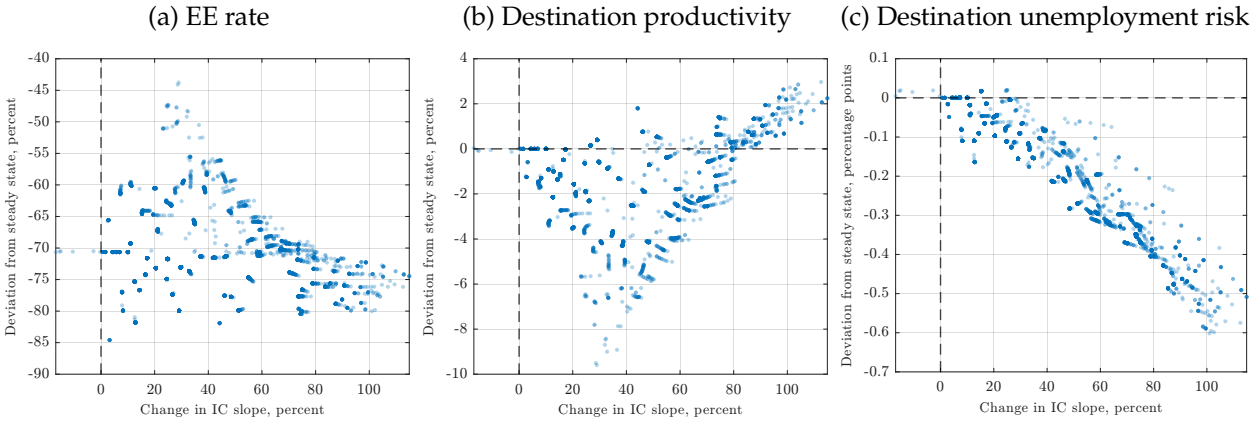

Notes: This figure plots per-cell changes on impact of the recession ($t = 1$ month) against the per-cell percentage change in the regression-based IC slope. Each dot represents a worker state (assets, wage, separation risk) drawn in proportion to the steady-state employment distribution, where darker markers arise from overlapping draws and thus indicate states with greater employment mass (sample size 5,000). Panels (a) and (b) plot the percentage changes in the cell's EE rate and in the destination productivity, respectively. Panel (c) plots the absolute change in the destination EU risk.

Table C.3: Who reroutes toward safety

	Share of employed	At borrowing constraint	Mean assets	Mean wage z	Mean risk δ (ann.)
Become more selective	69.1%	50.1%	−3.82	1.52	8.3%
Become less selective	30.5%	10.3%	−2.23	1.67	24.8%
All employed	100.0%	38.1%	−3.34	1.57	13.3%

Notes: This table reports characteristics of employed workers, split by whether their job-acceptance probability falls on impact of the recession ("become more selective") or rises ("become less selective"), weighted by steady-state employed mass. "At borrowing constraint" is the share of the group exactly at the borrowing limit $\underline{a}$; mean assets are in model units; mean wage z and mean risk δ (annualized) are averages over the workers' *current* jobs. The bottom row reports the population average.

References: Online Appendix

Acemoglu, D. and R. Shimer (1999): "Efficient Unemployment Insurance," *Journal of Political Economy*, 107, 893–928.

——— (2000): "Productivity Gains from Unemployment Insurance," *European Economic Review*, 44, 1195–1224.

Browning, M. and T. F. Crossley (2001): "Unemployment Insurance Benefit Levels and Consumption Changes," *Journal of Public Economics*, 80, 1–23.

Callaway, B. and P. H. C. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225, 200–230.

Davis, S. J. and T. von Wachter (2011): "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, Fall, 1–72.

den Haan, W. J., P. Rendahl, and M. Riegler (2018): "Unemployment (Fears) and Deflationary Spirals," *Journal of the European Economic Association*, 16, 1281–1349.

Elsby, M. W. L. and A. Gottfries (2022): "Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations," *The Review of Economic Studies*, 89, 1370–1419.

Fallick, B. and C. A. Fleischman (2004): "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," *Finance and Economics Discussion Series No. 2004-34, Federal Reserve Board*.

Federal Reserve Bank of New York (2013–2026): "Survey of Consumer Expectations [dataset]," Accessed June 7, 2026.

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, D. Backman, E. Breton, G. Cooper, J. A. Rivera Drew, S. Richards, D. Van Riper, and K. C. W. Williams (2025): "IPUMS CPS: Version 13.0 [dataset]," Minneapolis, MN: IPUMS, https://doi.org/10.18128/D030.V13.0. Underlying data provided by the U.S. Census Bureau and the Bureau of Labor Statistics.

Fujita, S., G. Moscarini, and F. Postel-Vinay (2024): "Measuring Employer-to-Employer Reallocation," *American Economic Journal: Macroeconomics*, 16, 1–51.

Gottfries, A. (2026): "Bargaining with Renegotiation in Models with On-the-Job Search," *Review of Economic Dynamics*, 61, 101343.

Hall, R. E. and P. R. Milgrom (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674.

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

Jarosch, G. (2023): "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, 91, 903–942.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432–1467.

Lucas, R. E. and B. Moll (2014): "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, 122, 1–51.

Pinheiro, R. and L. Visschers (2015): "Unemployment Risk and Wage Differentials," *Journal of Economic Theory*, 157, 397–424.

Shimer, R. (2006): "On-the-Job Search and Strategic Bargaining," *European Economic Review*, 50, 811–830.