Risk Premia, Firm Insurance, and Endogenous Labor Income Risk*

Maarten Meeuwis
Washington University in St. Louis

Dimitris Papanikolaou
Kellogg School of Management and NBER

Lawrence D. W. Schmidt
University of California, San Diego and NBER

July 2026

Abstract

We study how aggregate financial conditions shape firm insurance and, through it, labor income risk. In a directed search model with dynamic wage contracts and two-sided limited commitment, firm insurance against idiosyncratic shocks erodes when risk premia rise. Using U.S. administrative data, we document new evidence supporting the model: pass-through of firm shocks to earnings rises in bad times, especially for lower-paid workers near the separation margin. The model reproduces many untargeted time-series and cross-sectional features of earnings risk and implies substantial welfare costs of idiosyncratic risk, high private discount rates on human capital, and large gains from recession-contingent transfers.

*We thank Marios Angeletos, Susanto Basu, Corina Boar, Brett Green, Virgiliu Midrigan, Giuseppe Moscarini, Jonathan Parker, and Martin Souchier for useful comments. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: P-7503840 (P-7518691), Disclosure Review Board (DRB) approval number: CBDRB-FY25-0440).

Earnings data reveal that workers face large and heterogeneous risks that vary systematically over the business cycle.[1] Firms partially insure workers against these risks through implicit contracts, yet this insurance is incomplete and varies across workers.[2] These stylized facts raise three questions. First, what are the economic drivers of these empirical patterns? Second, what is the welfare cost of this idiosyncratic risk, and how does it vary across workers? Third, how do labor market interventions that mitigate this risk affect welfare? None of these can be answered from a description of the earnings process alone. Earnings losses do not translate one-for-one into welfare losses, since part of the observed variation reflects workers' own choices. Answering these questions requires a structural model that recognizes that these earnings and employment dynamics are equilibrium outcomes of workers' and firms' decisions.

To this end, we develop a model of endogenous labor income risk. The key ingredients are labor market frictions, endogenous wage dynamics through long-term contracts, and time-varying risk premia—a stand-in for financial market shocks. The interaction of search frictions and risk premia generates counter-cyclical unemployment as in Hall (2017). Firms partially insure risk-averse workers through dynamic wage contracts, but neither firms nor workers can commit to continue a match that yields less than its outside option. When risk premia rise, the value of employment relationships falls, limited commitment constraints tighten, and the scope for firm insurance shrinks as some jobs are endogenously destroyed. Workers therefore lose insurance precisely when labor markets are slack and job prospects are weak.

In our model, aggregate and idiosyncratic shocks are partially passed through to worker earnings along both intensive (earnings of job stayers) and extensive margins (probability of separation and reemployment earnings of job losers). Within continuing matches, optimal contracts trade off an *insurance motive*, which smooths wages against idiosyncratic shocks, and a *retention motive*, which backloads compensation to deter on-the-job poaching. For higher-wage workers who remain employed, pass-through operates mainly through wage adjustment within continuing matches, where the retention motive is strongest. For lower-wage workers, by contrast, pass-through operates mainly through endogenous job destruction: because they are close to the separation threshold and their gains from employment are especially backloaded, higher risk premia make these matches fragile and job loss costly.

Time variation in risk premia is the key driver of aggregate fluctuations in our model. Increases in risk premia lower the surplus value of employment, leading to fewer vacancy postings, higher

[1] Individual income growth rates are highly volatile and feature fat tails, countercyclical left-tail risk, and large differences in volatility across income groups (Guvenen, Ozkan, and Song, 2014; Guvenen, Karahan, Ozkan, and Song, 2021). Earnings losses from job displacement are large and roughly twice as large in recessions (Davis and Von Wachter, 2011). Unemployment fluctuates sharply over the business cycle, far more than can be explained by productivity fluctuations alone (Shimer, 2005). Worker separations are strongly countercyclical.

[2] See, for instance, Guiso, Pistaferri, and Schivardi (2005); Friedrich, Laun, Meghir, and Pistaferri (2019); Chan, Xu, and Salgado (2019).

separation rates, and a slacker labor market. The model therefore generates large, persistent, and countercyclical scarring effects of job loss: workers who lose their jobs in bad times—whether through endogenous destruction or exogenous displacement—face weaker hiring prospects, lower surplus upon reemployment, and less insurance in newly formed matches, as do workers who enter the labor market in downturns. Even temporary spikes in risk premia can thus have persistent effects on earnings and employment.

A central, testable prediction of the model concerns the interaction between firm-specific shocks and aggregate risk premia. Because higher risk premia make low-surplus matches more fragile and shrink the scope for firms to smooth idiosyncratic shocks, the pass-through of firm-level shocks rises with risk premia, with the amplification concentrated among workers near the separation margin and operating primarily through job loss. We confirm this prediction using U.S. administrative data: the pass-through of firm-specific TFP shocks to worker earnings rises with aggregate risk premia, especially for lower-paid workers, and operates through separation from the initial employer rather than wage cuts among stayers. These results are not driven by standard business cycle channels. First, we rank workers based on their earnings relative to other workers in the same firm, which eliminates much of the concern that firm-level differences drive the observed patterns. Second, controlling for aggregate conditions (growth in GDP, or an NBER recession indicator) leaves our estimates unchanged.

We calibrate the model to jointly match asset pricing moments, labor market dynamics, and the heterogeneous pass-through of firm and aggregate shocks to worker earnings. A central object is the value of nonemployment, which is not directly observed but which we infer from the level and cyclicality of separation and job-finding rates across income groups, building on Meeuwis, Papanikolaou, Rothbaum, and Schmidt (2025). The model rationalizes the evidence we document by making this value less sensitive to productivity and less backloaded than the value of employment, so that separation rates—especially for low-income workers—respond strongly to risk premia, limiting the scope for firm insurance in bad times. This discipline allows our calibrated economy to reproduce realistic unemployment fluctuations, worker flows (overall and across income groups), and labor market tightness, both in unconditional moments and along the realized path of risk premium shocks.

As validation, we show that the model reproduces a broad set of features of earnings risk that we do not target in the calibration. Under the optimal contract, wages are stable when limited commitment constraints are slack but adjust sharply when they bind or when the match dissolves. This nonlinear mapping from productivity shocks to earnings generates non-Gaussian earnings dynamics even though the model's underlying shocks are Gaussian. The implied distribution of individual earnings growth matches the fat tails, excess kurtosis, and negative skewness documented by Guvenen et al. (2021), with model kurtosis 11.3 versus 14.9 in the data.

The model also quantitatively replicates many facts about how earnings risk varies across workers,

including declining volatility over most of the income distribution, higher volatility for movers than stayers, substantial dispersion in cumulative lifetime earnings growth and lifetime employment rates, and the rise in left-tail earnings risk in downturns. As in the data (Catherine, 2022; Schmidt, 2025), stock market declines are associated with more negatively skewed distributions of earnings growth, with the model generating a time series of downside risk similar to its empirical counterpart from Guvenen et al. (2014). In addition, it generates large and persistent earnings losses from job displacement that are amplified in recessions (Davis and Von Wachter, 2011), as well as persistent earnings shortfalls for cohorts entering the labor market during downturns (Kahn, 2010; Schwandt and von Wachter, 2019).

Armed with a model of endogenous labor income risk that can quantitatively replicate its key empirical properties, we now use it to quantify three objects: the welfare cost of idiosyncratic risk, the value of human capital, and the effects of labor market policies. Each depends on how much workers value nonemployment and on how employment and wages respond to shocks and policy—equilibrium outcomes of workers' and firms' choices that a model with an exogenous income process cannot capture.

The cost of job loss illustrates the distinction between earnings losses and welfare losses. A worker's earnings fall sharply after an exogenous separation, but the value of nonemployment offsets part of the loss, so her welfare loss is far smaller than her earnings loss; for a worker who separates endogenously, it is zero. This gap between earnings and welfare losses is largest for low-paid workers, who face the highest job-loss risk, and it resolves a long-standing puzzle—why firms let workers go rather than retain them at lower pay (Bewley, 1999)—because a match near the separation margin has little surplus left to preserve.

We perform two quantitative experiments. First, we quantify the welfare cost of uninsurable idiosyncratic risk. Firm insurance is incomplete, leaving workers with large and heterogeneous exposures that have first-order implications for welfare. Even with modest aversion to idiosyncratic risk, a worker would require a permanent consumption increase of more than 30 percent to be compensated for bearing it. Although substantial, this cost is smaller than in models that treat the income process as exogenous and assign a low value to nonemployment (Krebs, 2007; De Nardi, Fella, and Paz-Pardo, 2020; Constantinides, 2024; Guvenen, Ozkan, and Madera, 2024): our model infers a higher value of nonemployment, disciplined by observed labor market behavior, which lowers the welfare cost of risk—especially for lower-paid workers for whom separation risk is highest. Idiosyncratic risk also depresses the value of human capital: the implied private discount rate rises from about 7 to 9 percent per year for the median worker, and from 13 to 16 percent for low-wage incumbents, who face the most volatile earnings paths. Workers with higher expected earnings growth—low-experience workers, a rough proxy for the young—face even higher discount rates.

Second, we evaluate different labor market policies, internalizing how policy changes endogenously

affect workers' labor supply decisions, firms' labor demand, and optimal contracts. We find that recession-contingent unemployment subsidies raise ex-ante welfare; these welfare gains do not reflect higher mean consumption—mean consumption is in fact lower than without policy—but optimal insurance provision. The gains are driven by reallocating consumption across time (from good to bad aggregate states) and across workers (toward those near the separation margin, where marginal utility is highest). These policies, however, dampen incentives to work, lowering average employment and output. Employment subsidies operate on a different margin: they support employment and output but provide less insurance, because they direct resources toward more productive workers who are more likely to remain employed. The optimal policy balances insurance against employment support—the welfare-maximizing policy heavily emphasizes insurance, while sustaining higher employment shifts this mix toward employment subsidies.

Our paper connects several strands of literature that have largely developed in parallel. One strand studies how discount rate shocks and financial conditions affect employment and labor market allocations (Hall, 2017; Kehoe, Midrigan, and Pastorino, 2019; Kehoe, Lopez, Midrigan, and Pastorino, 2023; Meeuwis et al., 2025); the other studies long-term contracts and firm insurance in frictional labor markets (Rudanko, 2011; Balke and Lamadon, 2022; Acabbi, Alati, and Mazzone, 2026). Relative to prior work, our framework emphasizes how the extent of firm insurance varies with aggregate financial conditions and how this interaction shapes both earnings risk and labor market allocations. Meeuwis et al. (2025) show that, in a model with exogenous wage contracts and complete markets, risk premium shocks disproportionately reduce the earnings and employment of lower-paid workers, and that this pattern is distinct from the generic fact that low-wage workers fare worse in downturns. In addition to endogenizing wage dynamics with incomplete risk sharing and risk averse households, we further ask whether aggregate financial conditions alter the extent to which firms insure workers against firm-specific shocks. Endogenous wage contracts are essential for this question: they allow us to interpret variation in pass-through as variation in firm-provided insurance, to trace its implications for worker welfare, and to evaluate labor market policies in a setting where wages respond to policy changes.

Malgieri and Citino (2025) also study how aggregate financial frictions affect workers' earnings. A key difference, however, is that risk premium shocks in our model affect both firms and workers, whereas their financial friction affects firms only. Because their model also features firing costs, tighter financial frictions can lead firms to borrow from workers—in effect, workers insure firms by remaining in low-surplus matches. That channel is absent here. As a result, their mechanism operates through the wage dynamics of continuing workers, whereas in our setting quantitatively important effects also operate through the extensive margin. Relatedly, Souchier (2023) and Nattinger and Rothbaum (2025) study how long-term contracts shape workers' earnings exposure to productivity shocks and firm-level volatility shocks, respectively. More broadly, Michelacci and Quadrini (2009)

show how financial constraints affect the wage profiles firms offer workers, while Kehoe et al. (2019) shows how debt constraints reduce the value of employment relationships and depress job creation.

Our model's implications for earnings risk also relate to a broader literature that explains similar facts through different, often complementary mechanisms. Hubmer (2018) shows that a job ladder model can account for non-Gaussian earnings growth through frictional job mobility and employment transitions. Our model adds additional richness but abstracts away from the job ladder: nonlinear wage contracting generates non-Gaussian earnings dynamics within matches, while cyclical variation in earnings risk arises mainly through the extensive margin. Huckfeldt (2022) and Acabbi et al. (2026) explain the scarring effects of recessions through selective hiring and the collapse of human capital ladders, providing mechanisms that are complementary to our discount-rate-based channel (see also Moscarini and Postel-Vinay, 2018, for a survey on cyclical job ladders). Ai and Bhandari (2021) study optimal dynamic contracts between workers subject to idiosyncratic shocks and diversified owners that can generate a time-varying price of aggregate risk. In their framework, countercyclical tail risk in idiosyncratic productivity is exogenous and gives rise to endogenous variation in risk premia (see also McKay, 2017, for a study of implications of exogenous cyclical variation in tail risk for aggregate consumption). To our knowledge, our paper is also the first to quantitatively replicate the cyclical pattern documented in Guvenen et al. (2014) *endogenously*, where countercyclical tail risk in earnings growth arises in periods of elevated risk premia.

Our empirical analysis also relates to the literature on firm insurance. Guiso et al. (2005) show that firms partially insulate workers from shocks to firm performance, while Friedrich et al. (2019) and Chan et al. (2019) document heterogeneity in the pass-through of firm shocks to worker earnings across the earnings distribution. We show that this pass-through also varies systematically over time with aggregate risk premia, with the amplification concentrated among lower-paid workers and closely tied to the extensive margin.

Finally, we relate to the literature on optimal policies over the business cycle. A large literature studies how unemployment insurance and related labor market transfers should vary over the business cycle, weighing the insurance value of transfers against search incentives and macroeconomic stabilization (Mitman and Rabinovich, 2015; Jung and Kuester, 2015; Kroft and Notowidigdo, 2016; Landais, Michaillat, and Saez, 2018b,a; McKay and Reis, 2021; Kekre, 2023). Our model highlights a distinct rationale for recession-contingent transfers: when risk premia rise, employment relationships lose value, limited commitment constraints tighten, and firm-provided insurance erodes, so public transfers substitute for private insurance precisely in the states—and for the workers—where marginal utility is highest.

Taking stock, the paper makes two main contributions. First, a single mechanism—a time-varying risk premium acting through optimal dynamic contracts—provides a unified explanation for a broad set of facts about how earnings risk and mobility vary over time and across workers, and

for the level and cyclical variation in workers' exposures to firm-level shocks. Second, because this earnings risk is an equilibrium outcome of workers' and firms' choices rather than an exogenous process, the model can recover objects that earnings data alone cannot: the welfare cost of risk, the value of human capital, and the effects of policies to which employment and wages respond. The remainder of the paper is organized as follows. Section 1 presents the model, and Section 2 discusses its key mechanisms. Section 3 presents the empirical evidence on heterogeneous and time-varying pass-through of shocks to earnings and describes the calibration. Section 4 evaluates the model's fit. Section 5 studies how state-dependent exposure shapes earnings risk, valuation, and policy.

1 Model

The model features a frictional labor market (Diamond, 1982; Mortensen, 1982; Pissarides, 1985) in which firms partially insure workers through dynamic contracts subject to limited commitment and time-varying risk premia. Workers are risk averse, heterogeneous, and subject to persistent productivity shocks. Firms search for workers through a directed search process (Montgomery, 1991; Moen, 1997). Firms cannot commit to continue matches that ex post yield negative value, so when surplus falls they either cut wages to restore profitability or terminate the match. As in Kehoe et al. (2019, 2023), worker productivity grows faster in employment than in nonemployment. The model also features on-the-job search (Postel-Vinay and Robin, 2002), which generates a retention motive in wage setting.

We model risk premium shocks as shifts in agents' effective risk aversion toward aggregate risk. These shocks induce fluctuations in the effective discount rates that workers and firms apply to risky future cash flows, in the spirit of Lettau and Wachter (2007) and Meeuwis et al. (2025). A rise in risk premia lowers the value of risky future payoffs. Because hiring and retention decisions depend on the present value of uncertain future surplus, fluctuations in risk premia affect labor market allocations directly. They also change the surplus available within the match and, through that channel, the amount of insurance firms provide workers.

1.1 Environment

Time is discrete. There is a unit measure of ex ante identical workers and a large number of ex ante identical firms. Workers differ in productivity and can be employed, unemployed and searching, or out of the labor force. Firms produce output with labor and post vacancies targeted to workers of specific productivity types.

Timing

Each period consists of five stages:

1. A fraction ζ of workers die and are replaced by new nonemployed workers, and aggregate and worker-level productivity shocks are realized.

2. A fraction s of existing matches is destroyed for exogenous reasons, and the remaining worker–firm pairs decide whether to continue or terminate their match.

3. Firms randomize the continuation values promised to surviving workers over a two-point lottery. This randomization ensures concavity of the firm's value function.

4. Search and matching take place. Employed workers receive an opportunity to search on the job with probability χ, unemployed workers search for new jobs, firms post vacancies, and new matches are formed. Following Balke and Lamadon (2022), the current employer cannot match a worker's outside offer ex post, so a successful outside offer ends the existing match.

5. Continuing and newly formed matches produce output and pay wages. Nonemployed workers then choose whether to enter the unemployment pool and search next period or to remain out of the labor force, receiving the corresponding flow payoff.

Production

The output produced by employed worker i at time t is

$$y_t(\Omega_{i,t}) = A_t h_{i,t} z_{i,t} \lambda_{i,t}. \tag{1}$$

Here, A_t is aggregate productivity, while $h_{i,t}$, $z_{i,t}$, and $\lambda_{i,t}$ are components of worker productivity. We summarize the worker's productivity state by $\Omega_{i,t} = (h_{i,t}, z_{i,t}, \lambda_{i,t})$.

Aggregate productivity follows a random walk with drift:

$$\Delta \log A_{t+1} = \mu_A + \sigma_A \, \varepsilon_{A,t+1}, \quad \varepsilon_{A,t+1} \sim N(0,1). \tag{2}$$

The first component of worker productivity, h, is permanent and evolves according to

$$\Delta \log h_{i,t+1} = g_{i,t} + \sigma_h \varepsilon_{h,i,t+1}, \qquad \varepsilon_{h,i,t+1} \sim N(0,1). \tag{3}$$

We interpret h as general human capital that grows with experience. Its growth rate depends on the worker's employment status, with $g_{i,t} \in \{g_E, g_O\}$. As in Ljungqvist and Sargent (1998), human capital grows faster in employment than out of employment, so $g_E > g_O$.

The second component, z, follows a mean-reverting process:

$$\log z_{i,t+1} = \psi_z \log z_{i,t} + (1 - \psi_z) \log \bar{z} + \sigma_z \varepsilon_{z,i,t+1}, \qquad \varepsilon_{z,i,t+1} \sim N(0,1). \tag{4}$$

Unlike h, which captures general human capital, z captures a component of worker productivity that matters primarily in employment. We model this distinction by assuming that the flow values

of unemployment and nonparticipation scale perfectly with h but only imperfectly with z. As a result, shocks to h do not affect labor market allocations, whereas shocks to z do.

The third component of worker productivity, $\lambda \in \{\bar{\lambda}_L, \bar{\lambda}_H\}$, captures experience accumulated while employed. While the worker remains employed, λ evolves according to

$$T_\lambda = \begin{bmatrix} 1-f & f \\ 0 & 1 \end{bmatrix}. \tag{5}$$

An inexperienced worker becomes experienced with probability f each period, and an experienced worker remains experienced as long as employment continues. Unlike h, which captures general human capital that accumulates gradually over time, λ is lost when the worker passes through nonemployment and is reset to $\bar{\lambda}_L$. By contrast, a worker who moves directly from one job to another without a nonemployment spell retains her λ type. As with z, λ affects productivity in employment but not the flow value of nonemployment.

Newborn workers enter the economy at time $t_0(i)$ without a job and with low employed experience, $\lambda_{i,t_0(i)} = \bar{\lambda}_L$. Their initial permanent human capital is $h_{i,t_0(i)} = \bar{h}$, and their initial z is drawn from $\log z_{i,t_0(i)} \sim N(\log \bar{z}, \sigma_{z0}^2)$.

Worker Preferences and Payoffs

Workers care about consumption. Their lifetime utility is defined recursively by

$$V_{i,t} = \left\{ c_{i,t}^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}\left[\Gamma_{t+1}\, V_{i,t+1}^{1-\gamma} \,|\, \mathcal{F}_{i,t}\right] \right\}^{\frac{1}{1-\gamma}}. \tag{6}$$

Here, $V_{i,t}$ is utility in consumption units, and $\mathcal{F}_{i,t}$ is worker i's information set at date t. This recursion follows Epstein and Zin (1989), with one modification: attitudes toward aggregate risk are governed by

$$\Gamma_{t+1} = \exp\left\{ -\frac{1}{2} x_t^2 - x_t\, \varepsilon_{A,t+1} \right\}, \tag{7}$$

where the effective degree of risk aversion toward aggregate shocks varies over time according to

$$\log x_{t+1} = \psi_x\, \log x_t + (1-\psi_x) \log \bar{x} - \sigma_x\, \varepsilon_{A,t+1}. \tag{8}$$

Equations (6)–(8) imply that workers evaluate continuation utility under a distorted conditional distribution for the aggregate shock $\varepsilon_{A,t+1}$. The Radon-Nikodym derivative in (7) captures this distortion. Appendix A.1 shows that this specification can be microfounded by ambiguity aversion over the aggregate shock ε_A, in the spirit of Hansen and Sargent (2001). The degree of ambiguity aversion varies over time through x_t and is perfectly negatively correlated with aggregate productivity shocks, in line with Campbell and Cochrane (1999); Kehoe et al. (2023), which simplifies the numerical solution of the model. An alternative interpretation is that agents hold distorted, pessimistic beliefs about future aggregate shocks, with (7) capturing the deviation from the objective distribution.

In sum, x represents the key aggregate shock in the model—time variation in risk premia—which we model as time-varying risk aversion towards aggregate shocks. Specifically, our formulation implies that the time variation in effective risk aversion applies only to aggregate shocks, not to idiosyncratic productivity shocks. Workers therefore exhibit source-dependent risk aversion, as in Skiadas (2013): aversion toward aggregate risk varies over time with x_t, while aversion toward idiosyncratic risk is constant at γ. The elasticity of intertemporal substitution is $1/\gamma$.

Workers are hand-to-mouth: they do not save or borrow. Employed workers therefore consume their wage, so $c_{i,t} = w_{i,t}$. Nonemployed workers receive and consume an outside-option payoff that depends on whether they search for a job. Workers who enter the unemployment pool and search consume

$$b_t(\Omega) = A_t\, h\, (\bar{b}_0 + \bar{b}_1 z). \tag{9}$$

As in Hall (2017) and Kehoe et al. (2023), the opportunity cost of employment scales one-for-one with aggregate productivity A and permanent human capital h, consistent with Chodorow-Reich and Karabarbounis (2016). Following Kehoe et al. (2019), we also allow it to depend on worker productivity z in order to match labor market flows across the earnings distribution.

Workers who do not search and remain out of the labor force consume instead

$$n_t(\Omega) = \bar{n}\, A_t\, h. \tag{10}$$

Unlike unemployment benefits, the payoff from nonparticipation does not depend on z. This captures the idea that unemployment benefits are imperfectly tied to prior earnings, whereas nonparticipation is not. It also helps keep labor-force participation from being too cyclical, consistent with the data (Mukoyama, Patterson, and Şahin, 2018).

Firms and Financial Markets

Firms are owned by infinitely lived shareholders. Shareholders collect output, pay wages, post vacancies, and maximize the present value of profits. Their preferences take the same recursive form as in (6), with two differences. First, shareholders are infinitely lived, so $\zeta = 0$. Second, they are fully diversified and therefore bear no idiosyncratic risk. We capture this by setting $\gamma = 0$, which in our preference specification also implies an infinite elasticity of intertemporal substitution.

These assumptions imply that the price of a claim to a future cash-flow stream X is

$$P_t = \mathbb{E}_t \left\{ \sum_{\tau=1}^{\infty} \left(\prod_{s=1}^{\tau} \Lambda_{t+s} \right) X_{t+\tau} \right\}, \tag{11}$$

where the one-period stochastic discount factor is

$$\Lambda_{t+1} = \beta \exp\left\{ -\frac{1}{2} x_t^2 - x_t \varepsilon_{A,t+1} \right\}. \tag{12}$$

Thus, the shareholders' stochastic discount factor has the same form as in Lettau and Wachter (2007), with the risk-free rate pinned down by β.

Directed Search and Matching

Nonemployed workers decide each period whether to search for a job. Employed workers also receive an opportunity to search on the job, with probability

$$\chi(\Omega) = \frac{z^{\bar{\chi}_1}}{\bar{\chi}_0 + z^{\bar{\chi}_1}}. \tag{13}$$

Thus, the search pool consists of unemployed workers and a subset of employed workers who receive an on-the-job search opportunity.

Firms post vacancies directed at workers of a given productivity type. Labor markets are competitive: in each period, firms can freely enter any market. The per-period cost of posting a vacancy directed at worker type Ω is

$$\kappa_t(\Omega) = \bar{\kappa}_0 \, A_t \, h \, z^{\bar{\kappa}_1}. \tag{14}$$

As in Afrouzi, Blanco, Drenik, and Hurst (2024); Meeuwis et al. (2025), vacancy posting becomes more costly for more productive workers. Scaling vacancy costs with A and h ensures a nondegenerate limiting employment distribution, while allowing costs to rise with z helps the model match the weak dependence of job-finding rates on prior earnings in the data.

Vacancies are directed not only at worker type Ω but also at a promised lifetime utility value V. Labor markets are therefore organized into submarkets (Ω, V). Workers who search choose which submarket to enter, taking as given their current productivity type Ω and the utility value offered by each market.

In each submarket, matches depend on market tightness $\theta \equiv \nu/u$, where u is the measure of searchers and ν is the measure of posted vacancies. Following den Haan, Ramey, and Watson (2000), the matching function is

$$m(u, \nu) = \frac{u \, \nu}{(u^\alpha + \nu^\alpha)^{1/\alpha}}. \tag{15}$$

This specification implies bounded job-finding and vacancy-filling rates. In particular, the probability that a vacancy is filled is $q(\theta) = (1 + \theta^\alpha)^{-1/\alpha}$, and the probability that a searcher finds a job is $p(\theta) = \theta(1 + \theta^\alpha)^{-1/\alpha}$.

1.2 Labor Market Search and Contracting

Firms partially insure workers through dynamic labor contracts. A contract specifies the wage paid to an employed worker as a function of the worker's productivity history, which is observable to both parties and contractible. Worker search decisions on the job, however, are private information

and are not contractible. Workers and firms also have limited commitment: either side can walk away from a match that delivers less than its outside option.

Apart from these frictions, the contract is fully flexible, so the pass-through of aggregate and idiosyncratic shocks to wages is determined endogenously. Following Spear and Srivastava (1987), we summarize the relevant history by the worker's current promised utility, which keeps the state space tractable. Thus, the optimal contract solves a dynamic problem with forward-looking constraints and assigns a promised utility V to a worker of type Ω.

Worker Problem

Employed workers decide whether to keep their current match. Job seekers, whether unemployed or employed, choose which submarket to search in. Nonparticipants decide whether to enter the search pool next period.

Continuation decision. Consider a worker in an existing match that has not been exogenously destroyed. Let $V_t^O(\Omega)$ denote the worker's value of nonemployment. Limited worker commitment implies that a worker with continuation value V_t^C remains in the match if and only if $V_t^C \geq V_t^O(\Omega)$; otherwise, she leaves for nonemployment.

Worker directed search. Consider a worker who searches for a new job in the current search-and-matching stage. Her current outside option is V, which may be the value of unemployment or the value of her current job. If she is already employed, the current firm cannot match a successful outside offer ex post; instead, it can only affect the probability of retention in advance through the continuation value it promises. Given productivity type Ω, she chooses the submarket that maximizes the expected gain in lifetime utility:

$$R_t(\Omega, V) \equiv \sup_{\mathcal{V}} p(\theta_t(\Omega, \mathcal{V})) \frac{\mathcal{V}^{1-\gamma} - V^{1-\gamma}}{1-\gamma}. \tag{16}$$

Thus, the worker trades off the probability of finding a job against the utility gain conditional on finding one. We denote the solution by $\mathcal{V}_t^*(\Omega, V)$ and the associated job-finding probability by $p_t^*(\Omega, V) \equiv p(\theta_t(\Omega, \mathcal{V}_t^*(\Omega, V)))$. For a worker in an existing match, the implied retention probability for the firm is

$$\tilde{p}_t(\Omega, V) \equiv 1 - \chi(\Omega)\, p_t^*(\Omega, V). \tag{17}$$

Participation decision. Consider a worker of type Ω who ends the current period in nonemployment. Next period, she either remains out of the labor force, which yields value $V_t^N(\Omega)$, or enters the unemployment pool and searches, which yields value $V_t^U(\Omega)$. Her nonemployment value is therefore

$$V_t^O(\Omega) = \max\{V_t^U(\Omega), V_t^N(\Omega)\}. \tag{18}$$

A worker who remains out of the labor force receives the nonparticipation payoff (10) and, conditional on surviving, stays nonemployed next period. Her continuation value is

$$V_t^N(\Omega) = \Big\{ n_t(\Omega)^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\,V_{t+1}^O(\Omega')^{1-\gamma}\Big]\Big\}^{\frac{1}{1-\gamma}}. \tag{19}$$

An unemployed worker receives unemployment benefits (9) and, conditional on surviving, searches in the next search-and-matching stage. Her continuation value is

$$V_t^U(\Omega) = \Big\{ b_t(\Omega)^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{V_{t+1}^O(\Omega')^{1-\gamma} + (1-\gamma)\,R_{t+1}\Big(\Omega', V_{t+1}^O(\Omega')\Big)\Big\}\Big]\Big\}^{\frac{1}{1-\gamma}}. \tag{20}$$

Firm Problem

Due to the linear production technology, the firm's problem can be analyzed individually for each job.

Continuation decision. Let $J_t^{FC}(\Omega, V^C)$ denote the value to the firm of continuing a match in the second stage when the worker has promised utility V^C. Limited firm commitment implies that the firm terminates the match whenever $J_t^{FC}(\Omega, V^C) < 0$, its outside option. Combining this with limited worker commitment, the continuation decision is

$$\mathbb{1}_t^C(\Omega, V^C) = \begin{cases} 1 & \text{if } J_t^{FC}(\Omega, V^C) \geq 0 \text{ and } V^C \geq V_t^O(\Omega), \\ 0 & \text{otherwise.} \end{cases} \tag{21}$$

Randomization. In the third stage, the firm chooses the continuation utility V^S promised to a worker who stays rather than leaving through on-the-job search. The promise-keeping constraint requires that the worker's ex ante value equal the value promised previously. As in Thomas and Worrall (1990), the firm may randomize over continuation utilities in order to preserve concavity of the value function J. In particular, it chooses probabilities π_j and promised utilities V_j^S over a two-point lottery:

$$J_t^{FC}(\Omega, V^C) = \max_{\pi_j, V_j^S} \quad \sum_{j=1}^{2} \pi_j\, \tilde{p}_t(\Omega, V_j^S)\, J_t^{FS}(\Omega, V_j^S) \tag{22}$$

$$\text{s.t.} \quad V^C \leq \left(\sum_{j=1}^{2} \pi_j \Big\{ (V_j^S)^{1-\gamma} + \chi(\Omega)\,(1-\gamma)\,R_t(\Omega, V_j^S) \Big\} \right)^{\frac{1}{1-\gamma}}. \tag{23}$$

Vacancy posting. The value to a firm of posting a vacancy directed at workers of type Ω and promising utility V is

$$\Pi_t(\Omega, V) = -\kappa_t(\Omega) + q(\theta_t(\Omega, V)) J_t^{FS}(\Omega, V). \tag{24}$$

A firm posts such a vacancy if and only if $\Pi_t(\Omega, V)$ is weakly positive. With free entry, vacancies are posted until expected profits are nonpositive in every submarket:

$$\Pi_t(\Omega, V) \leq 0 \quad \forall (\Omega, V), \tag{25}$$

with equality whenever $\theta_t(\Omega, V) > 0$.

Optimal dynamic contract. Consider a worker who remains employed through the final stage of the period. Let V^S denote her current continuation value. The contract delivers this value through two objects: a current wage w, which the worker consumes today, and a state-contingent promised continuation value $V^{C\prime}$ for the next period, conditional on the match surviving until the end of the second stage. Promise-keeping requires that these objects satisfy

$$\begin{aligned} V^S = \Big\{ w^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{ V^O_{t+1}(\Omega')^{1-\gamma} + \\ + (1-s)\mathbb{1}^C_{t+1}(\Omega', V^{C\prime})\Big((V^{C\prime})^{1-\gamma} - V^O_{t+1}(\Omega')^{1-\gamma}\Big)\Big\}\Big]\Big\}^{\frac{1}{1-\gamma}}. \end{aligned} \tag{26}$$

The firm chooses the current wage w and the state-contingent promised utility $V^{C\prime}$ to maximize the value of the match:

$$J^{FS}_t(\Omega, V^S) = \max_{w, \{V^{C\prime}\}} \; y_t(\Omega) - w + (1-\zeta)(1-s)\,\mathbb{E}_{t,\Omega}\Big[\Lambda_{t+1}\mathbb{1}^C_{t+1}(\Omega', V^{C\prime})\,J^{FC}_{t+1}(\Omega', V^{C\prime})\Big], \tag{27}$$

subject to the promise-keeping constraint (26), which holds with equality along the equilibrium path.

1.3 Competitive Search Equilibrium

We construct a competitive search equilibrium in the spirit of Montgomery (1991) and Moen (1997). An equilibrium consists of firm and worker value functions, contract and search policies, market tightness in each submarket, and distributions of workers and vacancies across states and submarkets. The model admits a block-recursive equilibrium.

2 Model Mechanisms

This section uses the model's numerical solution to illustrate the mechanisms behind the main predictions.

2.1 Job-Finding and Separation Rates

Higher risk premia x reduce vacancy creation and lower job-finding rates (Figure A.1). They also raise both the separation threshold $z^*(x, \lambda)$ and the participation threshold $\underline{z}(x)$, increasing job destruction and reducing labor force participation (Figure A.2). The mechanism follows Kehoe et al. (2019, 2023); Meeuwis et al. (2025): employment payoffs are more backloaded than nonemployment payoffs, so an increase in risk premia lowers the value of a match relative to nonemployment. This effect is strongest for low-productivity (z) and low-experience (λ) workers, whose productivity is expected to grow the most and whose matches have the least surplus.

The novel element is worker risk aversion. Figure A.2 compares the baseline thresholds with a counterfactual in which workers are risk neutral over idiosyncratic shocks. In the numerical solution,

risk aversion raises the *level* of both thresholds. With incomplete markets and limited commitment, employment exposes workers to substantial idiosyncratic risk. Risk-averse workers therefore value the uncertain, backloaded payoffs from employment less than risk-neutral workers do, which lowers match surplus and shifts more workers into nonparticipation and more matches into separation.

Risk aversion also lowers the *slope* of the thresholds with respect to risk premia. Separation destroys the partial insurance that firms provide against both aggregate and idiosyncratic risk. When risk premia are high, nonemployment becomes riskier as well: unemployment spells are expected to last longer, duration uncertainty rises, human capital erodes more during nonemployment, and newly formed matches provide less insurance, especially for workers near the separation threshold. These forces make separation and nonparticipation especially costly for low-productivity workers when risk premia are elevated, which dampens the response of z^* and $\underline{z}$ to x relative to the risk-neutral case.

2.2 Firm Insurance and Wage Dynamics

A central feature of the model is that firms partially insure risk-averse workers through optimal dynamic contracts. Two forces shape the wage path. The first is an *insurance motive*: firms smooth wages to shield hand-to-mouth workers from idiosyncratic risk. The second is a *retention motive*: firms backload compensation to deter on-the-job poaching. We discuss each in turn and then illustrate how wages respond to shocks. Appendix A.2 provides further details.

Optimal wage path. Along any path on which neither limited commitment constraint binds—that is, $V^{C\prime} > V^{O}_{t+1}(\Omega')$ and $J^{FC}_{t+1}(\Omega', V^{C\prime}) > 0$—the firm's first-order conditions, together with the promise-keeping constraint (26), imply the following condition linking the current wage to the worker's continuation value:

$$w^{-\gamma} \;=\; \underbrace{\beta\,\Gamma_{t+1}[V^{C\prime}]^{-\gamma}}_{\substack{\text{marginal effect of}\\ V^{C\prime}\text{ on lifetime utility}\\ \text{(per unit probability)}}} \Bigg/ \underbrace{\left[-\Lambda_{t+1}\frac{\partial J^{FC}_{t+1}(\Omega', V^{C\prime})}{\partial V^{C\prime}}\right]}_{\substack{\text{marginal cost to}\\ \text{firm of providing } V^{C\prime}\\ \text{(per unit probability)}}} = \frac{[V^{C\prime}]^{-\gamma}}{\left[-\frac{\partial J^{FC}_{t+1}(\Omega', V^{C\prime})}{\partial V^{C\prime}}\right]}, \tag{28}$$

which (abstracting from randomization) can be rearranged to yield:

$$\frac{w'}{w} \;=\; \left[1 + \left(\frac{V^{S\prime}}{w}\right)^{\gamma} J^{FS}_{t+1}(\Omega', V^{S\prime})\,\frac{\partial \log \tilde{p}_{t+1}(\Omega', V^{S\prime})}{\partial V^{S\prime}}\right]^{1/\gamma}. \tag{29}$$

Equation (28) is a Borch-style risk-sharing condition within the match: the ratio of the worker's marginal utility of promised value to the firm's marginal cost of providing it is equalized across states in which neither limited commitment constraint binds. The first equality follows from the envelope theorem, and the second equality uses the fact that $\beta\,\Gamma_{t+1}/\Lambda_{t+1} = 1$. Thus, risk sharing within the match is efficient, subject to limited commitment. Because workers are risk averse over

both aggregate and idiosyncratic shocks while firms only price aggregate risk, firms can lower wage costs by insuring workers against idiosyncratic fluctuations. This is the insurance motive.

Equation (29) characterizes wage growth along the equilibrium path. The second term in brackets captures the *retention motive*: it depends on the elasticity of the retention probability (17) with respect to the worker's promised value. This elasticity is weakly positive, since higher continuation values reduce the probability that the worker receives a better outside offer. Because firms cannot counter a successful outside offer once it arrives, they must make the current match attractive enough in advance to discourage on-the-job search and limit poaching. This gives firms an incentive to backload wages. Without on-the-job search, the elasticity is zero and wages remain constant as long as the constraints do not bind, as in Thomas and Worrall (1990).

In sum, the optimal contract balances two forces: an insurance motive, which smooths wages, and a retention motive, which backloads compensation to deter outside offers. When neither limited commitment constraint binds, firms insure workers while still making the match attractive enough to limit poaching.

Behavior at the bounds and separation. At the limited commitment bounds, equation (29) no longer applies, and wages adjust to satisfy the promise-keeping constraint (26). When the worker's outside option binds ($V^{C\prime} = V^O_{t+1}(\Omega')$), the firm raises wages. When the firm's participation constraint binds ($J^{FC}_{t+1}(\Omega', V^{C\prime}) = 0$), the firm lowers wages. If $J^{FC}_{t+1}(\Omega', V^O_{t+1}(\Omega')) < 0$, no feasible contract satisfies both constraints, and the match is dissolved as in (21). Higher risk premia reduce match surplus, raise the separation threshold, and move more workers—especially low-productivity workers—into the region where no feasible contract satisfies both limited commitment constraints. Insurance therefore erodes precisely when workers are most vulnerable.

Illustration: response to shocks. Figure 1 uses the model's numerical solution to illustrate the qualitative forces at work. The purpose is to clarify how the insurance, retention, and separation margins map into continuation values and earnings responses. We focus on a worker with productivity $z = \bar{z}$, experience $\lambda = \lambda_H$, and promised utility at the midpoint of the limited commitment bounds, with aggregate risk premia fixed at their unconditional mean ($x = \bar{x}$).

The left panels show the response to human capital shocks h. Since separation risk does not depend on h, and because on-the-job search is limited for typical values of z, the insurance motive dominates. For small shocks, the firm fully insures the worker: wages remain flat while continuation values adjust within the nonbinding region. Once h' moves far enough from h, one of the limited commitment constraints binds and wages respond one-for-one. For higher values of z, on-the-job search becomes more important, so the retention motive also matters: wages respond to h even inside the bounds in order to deter outside offers, through the final term in (29).

The right panels show the response to productivity shocks z, which also affect separation and

outside offers. For small negative shocks, the firm absorbs the shock through lower continuation values, so wages fall only once the firm's limited commitment constraint binds. If z falls below the separation threshold, no feasible contract exists and the match dissolves. For large positive shocks, the retention motive becomes active. Because firms cannot respond to a successful outside offer once it arrives, they must discourage on-the-job search in advance by making the current match sufficiently attractive. Since the probability of an outside offer rises with z (see (13)), the firm raises wages and continuation values to reduce the worker's incentive to search and to deter poaching. For high-z workers, the threat of leaving comes mainly from successful on-the-job search rather than from separation into nonemployment, so the firm keeps continuation values close to—but not at—the upper bound (Figure 1b). The optimal contract therefore balances insurance against retention, smoothing wages as much as possible while limiting poaching.

Pass-through of idiosyncratic shocks to earnings. The bottom panels of Figure 1 summarize two patterns highlighted by the numerical solution of the model: pass-through varies systematically across workers in the earnings distribution, and for some shocks it also varies with aggregate risk premia x. Each line reports the coefficient from a regression of annual earnings on the realized annual idiosyncratic shock, estimated separately by aggregate shock ε_A and for incumbent workers in three prior-earnings bins based on the previous three years of labor income.

In the numerical solution, the pass-through of human capital shocks (Figure 1e) rises with prior earnings but is nearly flat across aggregate states. The cross-sectional pattern reflects the retention motive: for higher-value workers, firms raise wages more in response to increases in h in order to deter outside offers. By contrast, the model implies little state dependence in this pass-through. Because h is portable across jobs, unemployment, and nonemployment, shocks to h do not materially affect separation risk or the incentive to search for a new job. They therefore operate mainly through within-match contracting rather than through the extensive margin. Since risk premia in the model matter primarily by changing the value of backloaded employment relationships and the likelihood of separation, they have only a limited effect on the pass-through of h shocks.

The pass-through of shocks to z (Figure 1f) instead displays a U-shaped pattern across the earnings distribution. Pass-through is highest for top earners, lowest in the middle, and intermediate at the bottom. At the top, the retention motive raises pass-through on the intensive margin: workers with high z have better on-the-job search opportunities (equation (13)), so the contract tilts toward retention rather than insurance, making stayers' wages more responsive to shocks. In the middle, the retention motive is weak and limited commitment constraints rarely bind, so firms provide substantial insurance. At the bottom, pass-through is elevated through the extensive margin: low-productivity workers are close to the separation threshold, making total earnings—including nonemployment spells—highly sensitive to shocks.

A second implication of the model is that the pass-through of z shocks to worker earnings increases

when risk premia rise, especially for low-wage workers. The source of this state dependence is again the extensive margin. Adverse aggregate shocks raise risk premia and shift z^* upward (Figure A.2), so a given negative z shock is more likely to destroy the match. This effect is concentrated at the bottom of the earnings distribution, where workers are already close to the separation threshold, while pass-through at the top is shaped primarily by the retention motive and therefore changes less over the cycle.

In sum, the model highlights systematic heterogeneity in the pass-through of shocks to worker earnings across the earnings distribution. For some shocks, this pass-through also varies with aggregate risk premia x. These patterns guide our empirical analysis in the next section.

3 Mapping the Model to the Data

Our goal in this section is to quantitatively discipline the model. The previous section highlighted the mechanisms that govern earnings pass-through in the model: retention matters more for higher-earning workers, while endogenous separations matter more for lower-earning workers and become especially important when risk premia are elevated. Guided by these mechanisms, we focus on empirical moments that capture how the pass-through of shocks to worker earnings varies across workers and over time.

We proceed in two steps. First, we document in the data the cross-sectional and time-series variation in the pass-through of shocks to worker earnings. We then use these moments to discipline the quantitative analysis. Our key empirical object is how the pass-through of firm-specific shocks changes with aggregate risk premia. More broadly, we study how pass-through varies across the earnings distribution and over time with aggregate financial conditions.

To map the model to these empirical objects, we introduce a notion of firm-level shocks. Since decisions in the model are bilateral—between a single worker and a single firm—the boundary of the firm is indeterminate without further assumptions. We therefore assume that workers employed at the same firm are exposed to common components of the idiosyncratic shocks z and h. This extension lets the model speak directly to the pass-through of firm-level and aggregate shocks to worker earnings.

3.1 Data and Methodology

We study empirical moments that are informative about the model's main mechanisms by relating worker earnings to firm-specific and aggregate shocks. Our analysis uses a stratified random subsample of earnings records from the Longitudinal Employer–Household Dynamics (LEHD) database for 1990–2019, matched to firm data from Compustat. The matched panel follows incumbent U.S. workers employed by public firms and tracks their subsequent earnings. Appendix B.1 provides further details.

Our main outcome is cumulative, age-adjusted earnings growth, following Autor, Dorn, Hanson,

and Song (2014) and Guvenen et al. (2014):

$$g_{i,t:t+H} \equiv w_{i,t+1,t+H} - w_{i,t-2,t}, \qquad w_{i,\tau_1,\tau_2} \equiv \log\left(\frac{\sum_{\tau=\tau_1}^{\tau_2} \text{real wage earnings}_{i,\tau}}{\sum_{\tau=\tau_1}^{\tau_2} D(\text{age}_{i,\tau})}\right). \tag{30}$$

This measure emphasizes persistent changes by averaging earnings over multiple years while adjusting for the life-cycle earnings profile. We include workers employed by a Compustat firm in year t and then follow them regardless of subsequent employment status, so the measure captures both wage changes and nonemployment spells.

We relate worker earnings to two sources of shocks. The first is firm productivity growth, $\epsilon^{tfp}_{f,t+1}$, measured as annual revenue-based TFP following İmrohoroğlu and Tüzel (2014) (Appendix B.2). The second is an aggregate risk premium shock, ϵ^{rp}_{t+1}. We construct this series as the first principal component of AR(1) residuals from nine monthly financial indicators, including the excess bond premium of Gilchrist and Zakrajšek (2012), the VIX, the financial uncertainty index of Jurado, Ludvigson, and Ng (2015), and the risk appetite index of Bauer, Bernanke, and Milstein (2023), following Meeuwis et al. (2025). Because the underlying indicators are asset prices and risk spreads observed at daily or monthly frequency, they have a direct interpretation as time-varying compensation for risk rather than as proxies for slow-moving labor demand or credit conditions. This common factor explains 60% of the variance across the nine series. Because earnings are measured annually, we cumulate monthly risk premium shocks from mid-year t to mid-year $t+1$ to construct ϵ^{rp}_{t+1}. Appendix B.3 provides details.

3.2 Cross-Sectional Heterogeneity in Pass-Through to Workers

Before turning to the interaction between firm-specific shocks and aggregate risk premia, we revisit two empirical objects that are informative for the model's two main channels. First, the direct pass-through of risk premium shocks to worker earnings varies systematically with workers' prior earnings rank (Meeuwis et al., 2025). This object is informative because risk premia directly affect match surplus and the separation margin. Second, the pass-through of firm-specific shocks rises toward the top of the earnings distribution, consistent with Friedrich et al. (2019). Together, these estimates discipline the retention motive, which raises pass-through at the top of the earnings distribution, and endogenous separation, which raises exposure at the bottom.

We estimate the pass-through of firm productivity shocks and risk premium shocks to worker earnings and allow the coefficients to vary with the worker's prior earnings rank within the firm:

$$g_{i,t:t+H} = \sum_s \{\beta_s\, \epsilon^{tfp}_{f(i,t),t+1} + \gamma_s\, \epsilon^{rp}_{t+1}\}\, \mathbb{1}_{\tau(i,t)=s} + c'\mathbf{Z}_{i,t} + \eta_{i,t+H}. \tag{31}$$

Here, $f(i,t)$ denotes worker i's employer at time t, and $\mathbb{1}_{\tau(i,t)=s}$ indicates the worker's prior within-firm earnings bin. Because workers are ranked by prior earnings within the firm rather than economy-wide, the heterogeneity we estimate reflects differences across workers employed at the same firm,

not differences across firms. Controls include a third-order polynomial in prior log average earnings, the lagged risk premium index interacted with earnings-group dummies, industry × earnings-group fixed effects, and worker industry × age × gender fixed effects. Standard errors are clustered by worker and year. Appendix B.4 provides details.

The pass-through of firm TFP shocks to cumulative earnings growth is roughly flat—around 3 to 5 percent—across the bottom 90 percent of the within-firm earnings distribution, and then rises sharply to 15–20 percent for workers in the top few percentiles (Figure 2a). This pattern is consistent with Friedrich et al. (2019). Adverse firm shocks also raise nonemployment risk, especially for lower-rank workers, for whom firm downturns translate primarily into separations rather than wage cuts (Figure 2b). In the model, these patterns line up with the two main channels: retention raises pass-through at the top, while the extensive margin drives exposure at the bottom.

The direct pass-through of risk premium shocks to worker earnings is U-shaped at short horizons—high for both low- and high-wage workers, and low in the middle—and becomes more strongly decreasing with income at longer horizons as cumulative exposure at the bottom rises further (Figure 3a). Again, the strong exposure of low-wage workers primarily reflects job destruction rather than wage declines (Figure 3b). This pattern is useful for our purposes because, in the model, risk premia directly affect worker earnings by changing match surplus and the separation margin, with especially strong effects for workers near endogenous job destruction.

3.3 Time Variation in Pass-Through to Workers

We now turn to a robust implication of the model: the pass-through of firm-specific shocks to worker earnings should vary with aggregate risk premia, with the amplification concentrated among lower-paid workers (Figure 1f). In the model, higher risk premia tighten limited commitment constraints, weaken firm insurance, and move workers near the separation margin closer to endogenous job destruction. This mechanism is strongest at the bottom of the earnings distribution, where the extensive margin is most important. We begin by documenting the time variation in pass-through directly, estimating time-varying pass-through coefficients:

$$g_{i,t:t+H} = \beta_t \, \epsilon^{tfp}_{f(i,t),t+1} + \xi_{I(i,t),\tau(i,t),t} + c' \mathbf{Z}_{i,t} + \eta_{i,t+H}, \tag{32}$$

where $\xi_{I(i,t),\tau(i,t),t}$ are industry × earnings group × year fixed effects. Figure 4 shows that the degree of pass-through varies considerably over time, and that periods of high pass-through systematically coincide with increases in risk premia.

We then estimate a variant of equation (31) that interacts firm productivity shocks with risk premium shocks and allows the coefficients to vary with the worker's prior earnings rank:

$$g_{i,t:t+H} = \sum_s \{\beta_{0,s} \, \epsilon^{tfp}_{f(i,t),t+1} + \beta_{1,s} \, \epsilon^{tfp}_{f(i,t),t+1} \times \epsilon^{rp}_{t+1} + \gamma_s \, \epsilon^{rp}_{t+1}\} \, \mathbb{1}_{\tau(i,t)=s} + c' \mathbf{Z}_{i,t} + \eta_{i,t+H}. \tag{33}$$

Table 1 reports the estimated coefficients β_0 and β_1 for a three-year horizon. A key pattern emerges: although the average pass-through β_0 of firm-specific shocks increases with workers' prior earnings, the interaction coefficients β_1 are large and positive for lower-paid workers and decline sharply with prior earnings: higher risk premia amplify the pass-through of firm-specific shocks, with this amplification concentrated among low-wage workers. This finding is robust across alternative specifications: it is unchanged when controlling for aggregate output growth or the fraction of the year spent in NBER recessions, and it holds across horizons of one to five years (columns (1)–(3) of Table 2). The fact that our estimates are largely invariant to including these controls implies that the amplification we document is tied directly to fluctuations in risk premia rather than the business cycle.

In the model, the time variation in these earnings losses arises because higher risk premia shift the separation margin upward, so adverse firm shocks are more likely to destroy low-surplus matches. We therefore examine two direct implications on the extensive margin. First, using a nonemployment indicator—equal to one if the worker experiences at least one full quarter with zero wage earnings over the next H years—as the outcome variable, columns (4)–(6) of Table 2 show that the interaction between firm TFP shocks and risk premia also raises the probability of job loss, with the effect again concentrated among lower-paid workers.

Second, Table 3 reports estimates of equation (33) separately for movers (workers who leave their initial employer) and stayers (those who remain). The increased pass-through during periods of elevated risk premia is substantially larger for movers than for stayers, particularly among lower-paid workers, providing direct evidence that the amplification operates primarily through the extensive margin rather than through greater wage cyclicality among stayers.

In sum, the pass-through of firm-specific shocks to worker earnings varies systematically across both time and workers. Higher risk premia increase the pass-through of firm shocks to workers, especially for lower-paid workers, primarily through higher rates of job destruction. We next interpret this reduced-form fact through the lens of our structural model. Calibrating the model to the degree to which pass-through varies across workers and with the level of risk premia allows us to infer the extent to which increases in risk premia erode firm insurance.

3.4 Model Calibration

We now discipline the model quantitatively. We proceed in two steps. A subset of parameters is pinned down a priori using external evidence, while the rest is chosen jointly to match moments from asset markets, labor market dynamics, firm-level and worker-level volatility, and the cross-sectional and time-series pass-through of shocks to worker earnings. Table 4 summarizes the baseline calibration.

Parameters Calibrated a Priori

We first calibrate a subset of parameters using a priori information, listed in Panel A of Table 4. We set the mean μ_A and volatility σ_A of aggregate productivity growth to match the corresponding moments of aggregate labor productivity growth from the U.S. Bureau of Labor Statistics (BLS) between 1947 and 2019, which are 2.2 percent and 1.8 percent per year, respectively. The model is calibrated at a monthly frequency, and all values are converted accordingly.

We normalize the initial level $\bar{h}$ of permanent human capital h, the long-run mean of persistent worker productivity z, and the lower level of worker experience λ to one. The long-run growth rate of permanent human capital during employment, g_E, is 3.5 percent per year, following Kehoe et al. (2023). The persistence of the worker productivity process z is $\psi_z = 0.991$ at a monthly frequency, consistent with Menzio, Telyukova, and Visschers (2016), which implies a half-life of roughly six years. The volatility of persistent productivity shocks is $\sigma_z = 10.9\%$, as in Meeuwis et al. (2025), and the dispersion of initial worker productivity is $\sigma_{z0} = 0.666$, chosen to match the interquartile range of earnings at age 25 documented by Guvenen, Kaplan, Song, and Weidner (2022). We assume that it takes on average three years for a worker to progress to the upper level of λ, corresponding to a transition probability of $f = 1/36$ per month.

Finally, we choose the mortality rate ζ so that the average model lifespan of a worker equals 30 years, and we set the curvature parameter of the matching function to $\alpha = 0.407$, following Hagedorn and Manovskii (2008).

Parameters Calibrated to Targeted Moments

The second step of the calibration chooses the remaining parameters jointly to match the targeted moments listed in Panel B of Table 4. We target 91 empirical moments with 20 parameters. Many of these moments come from asset markets and aggregate labor market dynamics. In addition, we target heterogeneous pass-through coefficients and their interaction with risk premia. For exposition, we group the targeted moments and the parameters most closely associated with them into four categories, while recognizing that equilibrium outcomes reflect the joint interaction of all parameters.

Asset markets. We assume that dividends represent a levered claim on aggregate TFP, and choose the parameters governing the dynamics of the aggregate stochastic discount factor to match key moments of U.S. asset markets: the average real risk-free rate, the mean and persistence of the price–earnings ratio, and the mean and volatility of annual equity returns. Specifically, we set $\beta = 0.9988$ so that the real risk-free rate is 1.4% per year as in the data. The calibrated values $\mu_E = 0.08\%, \bar{x} = 0.079$, $\psi_x = 0.993$, and $\sigma_x = 10.7\%$ imply that fluctuations in the market price of risk are highly persistent and sufficiently volatile to reproduce the historical behavior of equity valuations and returns (Appendix A.3).

Labor market dynamics. We discipline the 7 parameters governing job flows, search frictions, and match surplus values to match key features of aggregate and cross-sectional labor market dynamics. We target the mean (6.5%) and volatility (1.4%) of the HP-filtered unemployment rate; the cyclicality of the labor force participation rate, measured by its regression beta on the unemployment rate; the mean and cyclicality of aggregate job-finding and separation rates constructed from CPS microdata (1978–2019) using Abowd–Zellner corrected transitions following Elsby, Hobijn, and Şahin (2015); Krusell, Mukoyama, Rogerson, and Şahin (2017) (see Appendix B.5); and the mean and cyclicality of relative job-finding and separation rates by prior earnings level, estimated from the Survey of Income and Program Participation (SIPP) between 1990 and 2019 (see Appendix B.6).

We set the exogenous separation rate s to 0.75%, targeting average separation rates of higher-wage workers. The unemployment benefit parameters ($\bar{b}_0 = 1.8, \bar{b}_1 = 0.77$) and the nonparticipation payoff $\bar{n} = 2.3$ determine the level of match surplus and the value of job search across worker productivity states z, and hence the rate of endogenous separations and the average unemployment rate. The growth rate of human capital during nonemployment, $g_O = 0.3\%$ per year, shapes time variation in endogenous separation and job-finding rates. This relative decline in labor productivity during nonemployment aligns with Kehoe et al. (2019, 2023) and with micro estimates of human capital depreciation (Couch and Placzek, 2010). Last, the vacancy cost parameters ($\bar{\kappa}_0 = 0.016, \bar{\kappa}_1 = 2.7$) are chosen to match average job-finding rates by prior earnings in the data.

Because the value of nonemployment is not directly observed, we infer it from workers' transitions into and out of unemployment. Its level is disciplined by separation rates across the earnings distribution, while the strong cyclicality of separation and job-finding rates requires the net value of employment to have a long duration, so that the separation threshold and labor market tightness are sensitive to discount rates. This long duration is driven by the gap between human capital growth in employment and nonemployment ($g_E - g_O = 3.2$ percent per year) and, for workers near the threshold, from mean reversion in z, a mechanism that Meeuwis et al. (2025) discuss in detail. The implied average payoff in nonemployment is 73 percent of the average wage in employment—within the range of values used in the literature (Shimer, 2005; Hagedorn and Manovskii, 2008; Chodorow-Reich and Karabarbounis, 2016) and somewhat lower than the 85 percent used by Rudanko (2011) in a model with hand-to-mouth workers. We discuss the labor market dynamics implied by the calibrated model and their fit to the targeted moments in Section 4.1.

Firm-level shock structure and volatility moments. The empirical analysis is organized around firm-level TFP shocks, while the model is written at the level of individual worker–firm matches. To connect the two, we assume that a component of individual productivity shocks is common to all workers currently employed by the same firm. Since production is linear, this extension induces comovement in worker outcomes within firms while preserving the tractability of ex-ante identical firms.

Formally, consider a worker i who is employed by firm j at the beginning of period t. Following Balke and Lamadon (2022), we assume that worker productivity shocks are partly driven by firm-wide productivity shocks. Specifically, shocks to permanent human capital h and persistent productivity z are given by

$$\varepsilon_{h,i,t} = \rho_h \, \tilde{\varepsilon}_{h,j,t} + \sqrt{1-\rho_h^2} \, \varepsilon^{\perp}_{h,i,t} \tag{34}$$

$$\varepsilon_{z,i,t} = \rho_z \, \tilde{\varepsilon}_{z,j,t} + \sqrt{1-\rho_z^2} \, \varepsilon^{\perp}_{z,i,t}, \tag{35}$$

where $\tilde{\varepsilon}_{h,j,t}$ and $\tilde{\varepsilon}_{z,j,t}$ denote firm-level standard normal shocks and $\varepsilon^{\perp}_{h,i,t}$ and $\varepsilon^{\perp}_{z,i,t}$ are independent worker-specific standard normal shocks. The parameters ρ_h and ρ_z determine the within-firm correlation of productivity shocks and thus the strength of firm-level comovement in worker productivity growth.

Let $I_{j,t}$ denote the set of incumbent workers at firm j in period t. We define firm-level TFP growth as

$$\epsilon^{tfp}_{j,t+1} = \log\left(\int_{I_{j,t}} y_{i,t+1} \, di\right) - \log\left(\int_{I_{j,t}} y_{i,t} \, di\right) + \sigma_k \, \tilde{\varepsilon}_{k,j,t+1}, \tag{36}$$

where $\tilde{\varepsilon}_{k,j,t+1}$ is an independent firm-level standard normal shock capturing residual firm-specific fluctuations that do not affect worker fundamentals, as well as measurement error.

We discipline the strength of the firm-level component by targeting the empirical standard deviation of firm TFP growth (24% per year), which implies a value of $\sigma_k = 5.9\%$. On the worker side, we target the average volatility of annual worker earnings growth (53%) reported by Guvenen et al. (2014), which implies $\sigma_h = 2.7\%$.

Pass-through regressions. Finally, we discipline the remaining 6 parameters using the pass-through moments documented above. These moments connect the model to the paper's central empirical object: the exposure of worker earnings growth and nonemployment to firm TFP growth, risk premium shocks, and their interaction at horizons of one and three years, separately for the five income groups used in the empirical analysis.

The parameter $\bar{\lambda}_H = 2.4$ governs the upper level of worker experience, which creates scope for firms to insure workers against shocks. The parameters $\bar{\chi}_0 = 44.4$ and $\bar{\chi}_1 = 2.7$ imply that the monthly probability of on-the-job search is approximately 2 percent for low-skill workers and increases with individual productivity z. This heterogeneity in search opportunities affects workers' outside options and, consequently, the degree of pass-through of productivity shocks to wages. The parameters $\rho_h = 27.3\%$ and $\rho_z = 25.0\%$ govern the correlation of worker productivity shocks within firms and are disciplined by the combined pass-through of firm shocks on both the intensive and extensive margins, as well as its variation across horizons. The risk aversion parameter $\gamma = 0.48$ is calibrated so that the model matches the pass-through of risk premium shocks to worker earnings.

While this value is somewhat lower than is typically assumed in the literature, our model abstracts away household self-insurance via savings, a channel that Souchier (2025) shows amplifies workers' tolerance for idiosyncratic earnings risk. Section 4.2 discusses the resulting model's fit to the empirical pass-through regression estimates.

4 Model Fit

We first assess the calibrated model against the labor market and pass-through moments targeted in calibration. We then examine the forces behind the fit and, finally, subject the model to a stricter test: can it reproduce realized fluctuations when driven by the empirical series of financial shocks?

4.1 Heterogeneous Labor Market Dynamics

Table 5 reports the aggregate labor market moments targeted in the calibration. The model matches the mean unemployment rate (6.5%) and its volatility (1.4%), as well as the cyclicality of labor force participation (Panel A). It also matches the mean and cyclicality of aggregate job-finding and separation rates (Panel B): job finding falls in recessions, while separations rise.

Figure A.3 turns to cross-sectional heterogeneity in worker flows, which is also targeted in calibration through the mean and cyclicality of relative job-finding and separation rates by prior earnings. The model matches these patterns closely. Job-finding rates are nearly flat across income groups, whereas separation rates are sharply concentrated among low-wage workers, both in level and in cyclicality.

Beyond these targeted moments, the model captures several additional features of labor market dynamics. Unemployment is highly persistent, labor market tightness is volatile and strongly procyclical, and the employment-to-population ratio comoves closely with the business cycle. The main discrepancy is labor force participation volatility, which the model overstates, likely because it abstracts from non-economic motives for non-participation. Panel C provides a further untargeted check: following Shimer (2005, 2012), the job-finding margin accounts for most unemployment fluctuations in both the data and the model. The model thus captures the central role of job creation while also generating realistically countercyclical separations.

4.2 Pass-Through of Productivity and Risk Premium Shocks

We now turn to the pass-through moments targeted in the calibration. We estimate regression (33) on model-simulated data and compare the resulting coefficients with their empirical counterparts from Section 3.

Figure 5 compares model-implied and empirical pass-through coefficients for earnings growth across horizons and income groups. The fit is close. Three patterns are common to the model and the data. First, the unconditional pass-through of firm TFP growth to earnings rises with

worker income (panel a), consistent with the retention motive and the greater wage sensitivity of top earners. Second, the model reproduces workers' direct exposure to risk premium shocks (panel b). Third, the interaction between firm TFP growth and risk premia is positive for low-wage workers and near zero for high-wage workers (panel c). For lower-paid workers, this interaction is modest at short horizons but stronger at medium horizons, in line with the data.

Figure 6 repeats the analysis using job destruction—an indicator for a zero-earnings quarter—as the outcome. The model again matches the data closely. Adverse TFP shocks raise the probability of job loss, especially for lower-paid workers (panel a), the other side of the U-shaped exposure documented in Section 3.2: earnings pass-through rises with income, whereas the separation margin is concentrated at the bottom. Rising risk premia also increase nonemployment risk disproportionately for low-wage workers (panel b). The amplification of firm TFP shocks when risk premia are high is again modest at short horizons but lines up with the data at longer horizons (panel c).

Although stayer–mover differences are not targeted in the calibration, Appendix Figure A.4 shows that the model also speaks to this margin. The model overstates the unconditional TFP response for low-wage movers—likely due to our assumption that firm shocks are correlated with worker productivity z—but it captures the central empirical pattern: earnings respond much more strongly to risk premium shocks—and to the TFP–risk premium interaction—for movers than for stayers, especially among lower-paid workers.

4.3 Why the Model Fits the Data

Two ingredients are central to the model's fit: worker heterogeneity and differential exposure to worker shocks.

Worker heterogeneity. Labor productivity has three idiosyncratic components—human capital h, current productivity z, and experience λ—all of which raise earnings. Figure A.5 shows that higher-paid workers therefore tend to rank higher along all three dimensions: they have accumulated more human capital, draw higher current productivity, and are more likely to have high experience. Human capital displays the steepest gradient—about 2.5 log points from bottom to top of the earnings distribution, compared with 1.6 for z. Yet substantial dispersion in both h and z remains within each earnings percentile, so no single component pins down a worker's rank. The share of high-experience workers rises from roughly 60% at the bottom to over 90% at the top.

Earnings exposures to productivity shocks. Figure A.6 quantifies the contribution of each shock to variation in worker earnings. At short horizons, current productivity shocks z dominate (panel a), both because they are more volatile ($\sigma_z = 10.9\%$, versus $\sigma_h = 2.7\%$) and because they directly affect separation risk (panel c). Their earnings impact is largest at the bottom and top of the distribution, but for different reasons: at the bottom because job loss is more likely, and at the top because stayers' earnings are more sensitive. The same asymmetry explains why exposure

to z shocks becomes more state dependent when risk premia rise (panel b), with the additional sensitivity concentrated among lower-wage workers. Human capital shocks are the second most important source of earnings variation and matter most for higher-paid workers. Experience and aggregate shocks matter less because their within-period volatility is smaller.

4.4 Model-Implied Fluctuations

Figure 7 traces the model's response to a negative aggregate shock. The shock raises risk premia and lowers TFP, reducing output and employment while increasing unemployment. Earnings decline, especially for low-income workers, because separations rise and wages of stayers fall.

We then ask whether the same mechanism accounts for realized fluctuations in labor market quantities and pass-through. To do so, we feed the model the empirical series of financial shocks, using the risk premium shocks ϵ^{rp}_{t+1} from Section 3.1 as proxies for the model's aggregate shocks $\varepsilon_{A,t+1}$, and compute the implied time series for labor market variables. Because of scale invariance, these series do not depend on the realized path of aggregate TFP A.

Figure 8 plots the resulting series. Consistent with its shared ingredients with Meeuwis et al. (2025), the model tracks aggregate labor market dynamics closely: the correlation between model-implied and observed unemployment is 64%. Episodes of elevated risk premia—most notably the early 2000s and the 2008–09 financial crisis—are accompanied by sharp increases in unemployment and slow recoveries in both the model and the data. Job-finding rates and labor market tightness (V/U) are strongly procyclical, whereas separations are countercyclical. The model-implied path of V/U closely tracks the data, capturing the sensitivity of vacancy creation to aggregate conditions emphasized by Shimer (2005). Total employment also comoves strongly with the data, although the model overstates the volatility of the employment-to-population ratio (Section 4.1).

More importantly, the model matches the paper's central new empirical fact: the pass-through of firm shocks to worker earnings rises with aggregate risk premia, with the strongest amplification among lower-income workers near the separation margin. Because the model is calibrated to unconditional and cross-sectional pass-through moments, the realized time-series comovement is a genuine out-of-sample prediction. We show this by estimating the model analog of equation (32) on the realized paths, obtaining time-varying pass-through coefficients for three-year earnings growth by income group (Figures 8g and 8h; series are demeaned). For low-income workers, model-implied pass-through tracks both risk premia and the empirical estimates closely, with a correlation of 69%, although the data exhibit somewhat larger swings. For high-income workers, neither the data nor the model display systematic cyclicality in pass-through.

In sum, the same mechanism that matches the cross section also accounts for the realized time-series variation in exposure. Because wages are endogenous, pass-through has a structural interpretation as firm-provided insurance. This allows us to quantify what these facts imply for

worker welfare and to evaluate labor market policies in a setting where wages respond endogenously to policy changes. Section 5 pursues these implications.

5 Model Implications

Section 4 showed that the model matches how workers' exposure to firm shocks varies with aggregate financial conditions, especially among lower-income workers near the separation margin. We now use the model to study the consequences of this state-dependent exposure, in two steps. Sections 5.1 and 5.2 first show that the model reproduces a broad set of untargeted facts about earnings risk and scarring. We then exploit our calibrated structural model in Sections 5.3 and 5.4 to quantify the welfare cost of idiosyncratic risk, the value of workers' human capital, and the welfare gains from counterfactual labor market policies.

5.1 Worker Earnings Risk

We first assess whether the model can quantitatively account for the main empirical facts on earnings risk.

Overall Distribution of Earnings Growth

Administrative data from the U.S. Social Security Administration show that individual earnings growth is far from Gaussian (Guvenen et al., 2021). The distribution is sharply peaked at zero, has fat tails, and is negatively skewed. Roughly one-third of workers experience almost no earnings change in a given year, yet large increases and large declines occur far more frequently than a Gaussian benchmark would predict. The log density declines approximately linearly in both tails, consistent with Pareto-like behavior.

The model reproduces these features quantitatively, even though idiosyncratic productivity shocks are Gaussian. None of these moments are targeted in the calibration. Figure 9a plots the log density of annual earnings growth, where a Gaussian benchmark would appear quadratic; Appendix Figure A.7a plots the density in levels. The model matches the sharp peak, the fat tails, and the negative skewness observed in the data. The kurtosis of annual earnings growth is 11.3 in the model versus 14.9 in the data.

Two channels in the model generate these non-Gaussian properties endogenously. On the intensive margin, firm insurance makes earnings a nonlinear function of productivity shocks. Inside the limited commitment region, optimal contracts smooth wages, while the retention motive generates some pass-through to deter on-the-job poaching. When the firm's or the worker's limited commitment constraint binds or a match dissolves, wages are reset discretely, producing large earnings changes even from small shocks. This nonlinear mapping from productivity to earnings generates the mass near zero and the fat tails. On the extensive margin, infrequent transitions between employment and

nonemployment produce temporary zero earnings and large, asymmetric earnings changes, much as in the job-ladder model of Hubmer (2018). What our framework adds is that both margins—wage resets and job destruction—become more frequent when aggregate risk premia are elevated.

Appendix Figure A.7 decomposes the distribution along these margins. Panel (b) compares all workers to those continuously employed in years t and $t+1$: conditioning on continuous employment sharply compresses the left tail but only slightly compresses the right tail, so nonemployment spells account for most large earnings declines but few large gains. Panel (c) compares stayers (same employer throughout t and $t+1$) to movers (everyone else): stayers are concentrated near zero, reflecting on-the-job wage smoothing under firm insurance, while movers' distribution is wide and roughly symmetric. Taken together, the two panels imply that large negative earnings changes are largely due to job loss, while large positive changes are largely driven by job-to-job transitions.

Earnings Risk across Workers

Earnings risk also varies systematically across workers (Guvenen et al., 2021). Lower-income workers face much higher variance and a more symmetric distribution of earnings growth, whereas higher-income workers face lower variance but more asymmetric and leptokurtic growth.

Figure 9b shows that the model captures this cross-sectional pattern well. Earnings volatility declines from the bottom to roughly the 90th percentile, although the model does not reproduce the sharp rise at the very top.[3] The model also reproduces the much greater volatility of movers relative to stayers, especially among low-income workers (Figure 9c). Appendix Figures A.8a–A.8d show that the model also captures the broad cross-sectional patterns in skewness and kurtosis, though the model matches these higher-order moments somewhat less closely than it matches the volatility of earnings across the income distribution.

In the model, these cross-sectional patterns reflect heterogeneity in exposure along both the extensive and intensive margins, as illustrated in Figure A.6. Low-income workers are closer to the separation margin and face higher separation risk, so modest negative shocks are more likely to translate into job loss and large earnings declines. Higher-income workers rarely separate endogenously but experience more earnings variation through wage adjustment, because on-the-job search intensity is higher at the top. The same logic explains why movers are much more volatile than stayers—job changes entail wage resets and, for some workers, nonemployment spells—and why wage growth volatility among stayers is higher for high-wage workers, who benefit less from wage smoothing.

Earnings shocks also vary in persistence. Guvenen et al. (2021) show that for low-income workers, negative earnings shocks are short-lived while positive shocks are highly persistent; among high-income workers, the pattern is reversed. Appendix Figure A.9 shows that the model reproduces the

[3] One candidate for the missing source of risk at the very top is firm-specific human capital and creative destruction, as emphasized by Green, Kogan, Papanikolaou, and Schmidt (2025).

low-income pattern quite closely: negative earnings changes mean-revert strongly, whereas positive changes do not. The model misses the heterogeneity across the earnings distribution: these model-implied patterns are broadly similar across worker groups.

Earnings Risk over Time

Guvenen et al. (2014) show that cyclical variation in earnings growth is driven mainly by higher-order moments rather than the variance. In recessions, the left tail expands and the right tail contracts—large negative earnings shocks become more frequent, while large positive shocks become rarer—producing strongly countercyclical skewness even as overall dispersion moves little.

Figure 10 illustrates the mechanism in impulse-response form: following an adverse aggregate shock, the left tail of earnings growth widens persistently, while the right tail initially compresses before recovering.

We then ask whether the model replicates the observed time series of left-tail risk. Using the empirical series of risk premium shocks as a direct proxy for the model shocks, as in Section 4.4, we compare the time path of earnings risk in the model and the data. The top panel of Figure 11 plots the difference between the median and the 10th percentile of earnings growth, capturing left-tail risk; the bottom panel plots the difference between the 90th percentile and the median, capturing right-tail risk. In both the model and the data, downturns coincide with a widening left tail and a narrowing right tail of the earnings growth distribution. Quantitatively, the model tracks these fluctuations well, with correlations of 57% for the left tail and 34% for the right tail.

5.2 Scarring and Lifetime Consequences

We next turn to the longer-run consequences of the same state-dependent exposure for displaced workers and recession entrants.

Earnings Losses from Job Displacement

Job displacement causes large and persistent earnings losses, especially in recessions (Davis and Von Wachter, 2011). We ask whether the model can reproduce these facts.

Following the empirical design of Davis and Von Wachter (2011), we focus on incumbent workers below age 50 with at least three years of tenure at their current employer. Because displacement in the data is concentrated among lower-tenure and lower-wage workers (see also Farber, 2017; Margolis and Montana, 2024), we also restrict attention to workers with pre-displacement earnings below the median. In the model, a worker is displaced in year y if she receives an exogenous separation shock during year $y-1$, so that she has positive earnings from her pre-displacement employer in $y-1$ but not in y. We then compare average earnings around the displacement event to those of otherwise similar workers who are not displaced.

Figure 12 shows that displacement generates large and highly persistent earnings losses, with a

sharp deterioration in the first few years after job loss and only partial recovery at long horizons. These losses reflect both nonemployment spells and lower wages upon reemployment. As in the data, earnings losses are substantially larger when displacement occurs in recessions, defined in the model as periods in which the unemployment rate exceeds 8%, than in expansions. The model somewhat understates the long-run difference in earnings losses between expansions and recessions. A likely reason is the absence of a job ladder, a complementary channel that helps explain this pattern in recent work (see, e.g., Huckfeldt, 2022; Jarosch, 2023; Acabbi et al., 2026).

These are earnings losses, not welfare losses: after displacement, unemployment benefits and the value of nonemployment replace part of the lost employment income. Section 5.3 quantifies this distinction and shows that it reshapes both the magnitude and the cross-sectional incidence of the cost of job loss.

Entering the Labor Market in a Recession

Entering the labor market in a recession also leaves persistent scars. Kahn (2010) and Schwandt and von Wachter (2019) document sizable and long-lived earnings shortfalls for cohorts that begin working lives in downturns.

In the model, workers who enter in recessions face weaker hiring prospects, lower-surplus initial matches, and less insurance in newly formed employment relationships. Early setbacks in employment and earnings therefore cumulate into persistent losses. These effects are strongest for workers who enter with low initial productivity z_0, consistent with the evidence that recession-entry costs are larger for less skilled workers.

To quantify this effect, we compute average earnings by horizon for newborn workers who enter the labor market and compare those who enter in recessions to those who enter in expansions. Figure 13 shows that entering when unemployment is high has a large initial effect on earnings; the effect fades only gradually and remains visible for up to 10 years. The magnitudes are broadly comparable to the estimates of Schwandt and von Wachter (2019). In particular, the model matches the large initial decline in earnings but understates the persistence of the losses at medium horizons.

Lifetime Earnings and Employment across Workers

A central finding in Guvenen et al. (2021) is that lifetime labor market outcomes are highly dispersed. Cumulative earnings growth between ages 25 and 55 rises steeply with lifetime earnings rank, and years employed over the life cycle vary widely across workers. To fit these patterns, their reduced-form specification requires the risk of nonemployment to be strongly worker- and time-dependent.

Figures 9d and 9e show that the model reproduces these broad patterns well. Figure 9d plots cumulative earnings growth between ages 25 and 55 against workers' lifetime earnings percentiles. The upward slope is partly mechanical, but the degree of heterogeneity is large and close to the data: average earnings grow by 154% at the median, 635% at the top 5%, and 1024% at the

top 1%. Figure 9e shows substantial dispersion in lifetime employment rates, although the model understates the share of workers with very long nonemployment spells. To complement this analysis, Figure 9f plots the cross-sectional variance of log earnings by age, a moment the model also matches reasonably well.

These lifetime patterns arise because short-run nonlinear earnings risk cumulates over time. Low-productivity workers are closer to the separation threshold, so modest negative shocks are more likely to induce job loss, and their subsequent job prospects are weaker. As a result, the model generates persistent nonemployment spells and substantial dispersion in lifetime earnings and employment. The model therefore provides a structural rationale for the state-dependent nonemployment risk in Guvenen et al. (2021). This mechanism shares an important element with Jung and Kuhn (2019): workers with weaker labor market prospects recover more slowly after adverse shocks. In our model, however, this heterogeneity arises from worker productivity, endogenous separation risk, and firm insurance rather than from job-ladder dynamics.

Taken together, these results imply that the aggregate state x governs both short-run earnings risk and longer-run scarring. When risk premia are high, lower-productivity workers are more likely to separate, and, conditional on separation, their subsequent earnings losses are larger and more persistent. These nonlinearities cumulate into substantial heterogeneity in lifetime earnings and employment.

5.3 Valuation Implications

We now turn from the dynamics of labor income to its valuation: what do the risks documented above imply for workers and for the firms that partially insure them? These exercises share a common logic: they require an equilibrium model in which the earnings risk that arises endogenously is consistent with the value that workers place on nonemployment—an object that is not observed from earnings data alone.

Cost of Job Loss

We begin by assessing the cost of job loss, the normative counterpart to the positive displacement results of Section 5.2. Evaluating the welfare consequences of earnings losses around displacement—or any policy that would mitigate them—requires more than the observed income process: after displacement, unemployment benefits and home production replace part of the worker's lost employment income, so focusing on observed earnings losses alone overstates the welfare impact of job loss.

We therefore contrast two present-value measures of losses due to an exogenous separation. The first is the loss in the present value of the worker's earnings, discounted with the aggregate SDF ($\gamma = 0$)—an earnings-based measure in the spirit of the displacement literature (Davis and Von Wachter,

2011). The second is the loss in worker welfare, measured in certainty-equivalent consumption,

$$CE_{i,t} = (1 - \beta(1 - \zeta))^{1/(1-\gamma)} V_{i,t}, \tag{37}$$

where $V_{i,t}$ is the worker's expected lifetime utility; welfare aggregates her entire flow-utility stream and so credits the value of nonemployment she receives after job loss. The two measures give very different answers.

Panels (a) and (b) of Figure 14 plot both losses by the worker's state—current productivity z and experience λ—at average risk premia. Near the separation margin the two measures starkly diverge: displaced workers with low z forgo a sizable share of future labor earnings but the welfare loss remains close to zero. An experienced low-z worker, for example, loses about 13% of the present value of her earnings but under 1% in welfare terms. Although she is displaced by an exogenous shock, she was close to leaving for her outside option in any case, so the lost job is worth almost nothing to her, and inferring her welfare loss from her earnings loss is misleading. As z rises and workers are further away from the separation threshold, the welfare loss starts tracking the earnings loss more closely; both losses are larger for experienced (λ_H) workers, who have more valuable matches to lose, than for inexperienced (λ_L) workers.

This welfare loss has a direct interpretation: it is the compensation a worker would require to accept separation, or equivalently the amount she would pay to keep her job. Near the separation margin this willingness to pay is small, and for a worker who separates endogenously it is exactly zero—at the threshold she is indifferent between employment and her outside option. This logic resolves a tension that arises when earnings losses are read as welfare losses. If job loss were as costly as the earnings loss implies, separating rather than accepting a pay cut would leave large gains from trade unexploited, an apparent symptom of wage rigidity (Bewley, 1999). In our model there is no such puzzle: the marginal worker's flow value from employment already lies below her outside option—she stays only for the backloaded value of continued employment—so when that surplus erodes, separation is privately efficient rather than a failure to adjust wages.

Panels (c) and (d) of Figure 14 show how these patterns translate into losses across the earnings distribution and by aggregate state. Again, welfare losses are much smaller than earnings losses, especially at the bottom: in recessions, displaced workers in the bottom decile lose about 15 percent of the present value of their earnings but only 2 percent in welfare terms, because nonemployment payoffs replace a large share of their employed consumption. Moreover, the cross-sectional incidence reverses: earnings losses are roughly flat across the earnings distribution, whereas welfare losses rise steeply with earnings—from about 2 percent at the bottom to 10 percent at the top—because the nonemployment payoff replaces a smaller share of high earners' consumption. Observed earnings losses therefore misrepresent both the size and the incidence of the welfare cost of job loss, and recovering it requires the value of nonemployment that the model infers from observed behavior.

Welfare Cost of Idiosyncratic Risk

We now move to the welfare cost of *all* idiosyncratic risk, comparing the baseline economy with calibrated risk aversion, $\gamma = 0.48$, to a counterfactual valuation with $\gamma = 0$; the two valuations share the same labor and consumption dynamics and differ only in how workers value idiosyncratic risk. The gap between the certainty equivalents under the baseline and counterfactual valuations is our measure of the welfare cost of idiosyncratic risk.

From behind the veil of ignorance, overall welfare is the expectation of welfare for newborn workers over the unconditional distribution of the aggregate state x and initial productivity z_0. Certainty-equivalent consumption in the baseline valuation is 25% lower than under the counterfactual valuation with $\gamma = 0$. Equivalently, workers would require a permanent increase in consumption of about one third to be as well off as under risk neutrality with respect to idiosyncratic shocks. These large welfare losses reflect the substantial lifetime variation in consumption generated by the model's earnings dynamics and the assumption that workers are hand-to-mouth.

Figure 15 plots the welfare loss as a function of initial productivity z_0 (black line), integrating over the unconditional distribution of x. Losses are sizeable everywhere, ranging from 22 to 27 percent across initial productivity types, but the pattern is nonmonotonic: losses are largest for workers around the middle and at the top of the productivity distribution rather than at the very bottom. This pattern reflects a key advantage of our framework: labor supply is endogenous, and the value of nonemployment is inferred from behavior rather than imposed. Workers at the very bottom of the productivity distribution often choose nonparticipation and rely on relatively safe home production, so their welfare is less exposed to labor market risk than their low realized earnings might suggest. Workers near the participation and separation margins, by contrast, remain attached to the labor market but face large fluctuations in employment and reemployment prospects. Because part of observed earnings volatility reflects optimal choice rather than uninsured risk—and our welfare estimates account for these decisions—rankings by observed earnings losses need not coincide with rankings by certainty-equivalent losses.

To gauge how much this outside option matters, we recompute the welfare loss under a standard exogenous-income assumption. We leave the dynamics of workers' employment status and wages unchanged from the baseline economy but replace the model-implied nonemployment payoffs—$b_t(\Omega)$ for the unemployed and $n_t(\Omega)$ for nonparticipants—with a single fixed payoff that is independent of worker type:

$$b_t^{\text{fix}} = n_t^{\text{fix}} = \rho\, \overline{w/A} \times A_t, \tag{38}$$

where $\overline{w/A}$ is average earnings per unit of aggregate productivity among the employed, so nonemployment income is a constant fraction ρ of the average wage. Following Guvenen et al. (2024), we set ρ to their benchmark and high-floor values, 6.75% and 22%. This is the natural analogue of the

safety net in incomplete-markets models, in which the nonemployed consume a low, productivity-independent transfer with no separate value of home production or leisure.[4]

This change in nonemployment payoffs sharply raises the calculated welfare cost of idiosyncratic risk (colored lines in Figure 15). The overall loss rises from 25 percent in the baseline to 29 percent at $\rho = 22\%$ and 38 percent at $\rho = 6.75\%$—in the range of existing exogenous-income estimates: 33 percent in Guvenen et al. (2024), 26 percent in De Nardi et al. (2020), and 36–39 percent in Constantinides (2024). The discrepancy is largest precisely where the model infers the most valuable outside option: for the lowest-z_0 workers, who spend the most time nonemployed, the fixed-payoff assumption nearly doubles the welfare cost—from 22 to 41 percent—and inverts its cross-sectional gradient. Calculations that take the income process as given therefore overstate the welfare cost of idiosyncratic risk, most severely for low-paid workers, because they underestimate the value of nonemployment that the model infers from observed labor market flows.

One caveat for interpreting these estimates is that our model shuts down self-insurance through savings, which Souchier (2025) shows increases workers' tolerance for idiosyncratic risk. We partially compensate for this by calibrating $\gamma = 0.48$, roughly four times smaller than the benchmark value of 2 used in much of the literature, to reflect that workers' marginal propensity to consume is well below one. Even with this adjustment, the model likely understates heterogeneity in welfare losses: if higher-income households can smooth idiosyncratic shocks through savings while lower-income households remain closer to hand-to-mouth, welfare losses at the top would be smaller than in our benchmark, and cross-sectional dispersion in welfare losses correspondingly larger.

Human Capital Discount Rates

We now turn from worker welfare to the value of workers' human capital—the present value of their future labor income—and to the discount rate this value implies. Appendix Figure A.11 reports workers' private value of human capital as a ratio to annualized current consumption, an analog of the price–earnings ratio for equity. In the baseline model this ratio averages about 15, well below its counterfactual value at $\gamma = 0$, because idiosyncratic earnings risk discounts future labor income more heavily, on top of the discounting already applied for aggregate wage and employment risk (see also Huggett and Kaplan, 2016, for a related discussion).

We summarize this discounting with an implied human capital discount rate: the internal rate of return that equates the present value of a worker's expected future labor income to her private value of human capital. This rate is the sum of three pieces: a risk-free rate; a premium for exposure to aggregate risk, which is priced even under $\gamma = 0$ through the time-varying aggregate price of risk x; and a premium for exposure to idiosyncratic risk—the gap between the $\gamma > 0$ and $\gamma = 0$

[4] An alternative approach would be to assume that workers receive a fixed replacement of labor productivity when unemployed. Doing so would magnify how the differences we highlight below vary across the income distribution.

valuations. Human capital is risky both because labor income covaries with priced aggregate shocks and because workers cannot diversify their idiosyncratic earnings risk.

Figure 16 reports these human capital discount rates. Among incumbent workers, the average rate under risk neutrality with respect to idiosyncratic shocks ($\gamma = 0$) is about 7% per year and embeds a 2.1% premium for aggregate risk—similar to the 2.4% estimate in Lustig, Van Nieuwerburgh, and Verdelhan (2013). Pricing idiosyncratic risk raises the average to around 9% per year, adding roughly two percentage points. Figure 16a shows that the aggregate human capital discount rate is countercyclical and peaks following the Great Recession.

The cross section is also informative. For incumbent workers sorted by prior earnings (Figure 16b), the additional discount induced by idiosyncratic risk is largest for low-wage workers, who face the largest earnings risk and the highest expected consumption growth: for them, pricing idiosyncratic risk raises the discount rate from about 13 to 16 percent per year. For all workers sorted by labor productivity (Figure 16c), the key difference appears in the left tail and is driven by the nonemployed. Among the least productive workers, discount rates are similar under $\gamma = 0$ and $\gamma > 0$, since nonparticipating workers have payoffs that are not subject to idiosyncratic productivity shocks and have relatively low expected consumption growth. At the far right of the distribution, the two lines converge because two effects offset each other (Appendix Figure A.12): a higher subjective risk premium and a lower subjective risk-free rate as productivity is expected to decline and firm insurance becomes weaker.

Discount rates also vary along the experience dimension—a rough proxy for worker age in the model.[5] Figure 16d plots the average human capital discount rate by current productivity z, separately for inexperienced (λ_L) and experienced (λ_H) workers. At every level of z, inexperienced workers face higher discount rates: they sit closer to the separation margin, have more volatile earnings, and receive less firm-provided insurance, so their human capital is more equity-like.

Value of Firm Insurance

We next value wage contracts from the firm's perspective, computing the net present value of the compensation stream embedded in the contract. This differs from the calculations above: it is the present value of wages—discounted with the shareholders' stochastic discount factor ($\gamma = 0$)—rather than the valuation of a worker's lifetime consumption or human capital. Figure A.13 reports the present-value difference between the optimal contract of a new hire and a piece-rate benchmark that holds allocations fixed but constrains wages to be proportional to productivity. The piece rate is set so that the worker derives the same value from the newly formed match as under the optimal contract. A negative present-value difference means that the optimal contract is cheaper for the firm to provide.

[5] Workers in our model have finite lives, but the perpetual youth assumption makes it hard to capture the relevant notion of worker age in the model: young workers have earnings of longer duration than older workers; differences in worker experience partly capture that, since low-experience workers have earnings that are expected to grow faster than experienced workers.

Providing insurance to risk-averse workers allows firms to pay lower wages. That is, over most of the z distribution, insurance provision significantly reduces the NPV of compensation. In the right tail the sign reverses, because the retention motive dominates the insurance motive. To retain its most productive workers, the firm offers strongly nonlinear, backloaded pay, making stayers' earnings highly sensitive to z shocks at high productivity levels (Figure A.6e). Because these valuation differences are driven primarily by idiosyncratic rather than aggregate risk, they are not very sensitive to risk premia x: higher x reduces insurance gains for low-z workers by raising endogenous separation risk, while it shrinks the retention premium for high-z workers by lowering poaching risk.

Welfare Cost of On-the-Job Poaching

The optimal wage contract trades off insurance against retention: firms backload compensation to deter on-the-job poaching, which limits the insurance they can provide to risk-averse workers. We quantify the welfare cost of this distortion from optimal risk sharing by comparing the baseline economy to a counterfactual in which the on-the-job search probability $\chi(\Omega)$ is set to zero, holding other parameters constant. Figure A.14 reports the certainty-equivalent welfare gains as a function of initial productivity z_0, integrating over the unconditional distribution of x. Welfare gains are uniformly positive and monotonically increasing in z_0, rising from about 2 percent at the bottom to about 7 percent at the top. These are sizable magnitudes: removing on-the-job search recovers 10 to 25 percent of the welfare cost of idiosyncratic risk in Figure 15, depending on initial productivity.

The cross-sectional pattern reflects heterogeneity in baseline on-the-job search intensity. At low z, $\chi(\Omega)$ is small—about 2 percent per month—so the retention motive barely distorts the contract. As z rises, $\chi(\Omega)$ rises with it, and firms backload compensation more aggressively to deter poaching, exposing high-productivity workers to more idiosyncratic risk; the welfare benefits of muting the retention motive grow accordingly. Two features of our model explain the uniformly positive sign of these effects. First, firms in our model are homogeneous, so on-the-job search plays no productive role: worker mobility does not improve match quality. Second, free entry competes away firm rents, so the retention premium paid to high-z workers in the baseline model reflects limited firm insurance rather than a transfer from firms to workers. We find that expanding the scope for firm insurance has large welfare benefits.

5.4 Policy Counterfactuals

The valuation results suggest large gains from interventions that either provide insurance in unemployment or support employment relationships. We therefore consider two labor market policies that are activated only in recessions. Both are financed by a flat consumption tax chosen so that the net present value of transfers equals the net present value of tax revenue, computed using the shareholders' SDF.

We define recession states using the level of risk premia x_t, the aggregate state variable that

drives business cycles in the model. Specifically, we define the recession indicator $\mathbb{1}_t^R$ that marks periods in which labor markets are slack as

$$\mathbb{1}_t^R = \begin{cases} 1 & \text{if } \log x_t > \log \bar{x} + \sigma_x^{unc}\, \Phi^{-1}(1-\delta) \\ 0 & \text{otherwise,} \end{cases} \tag{39}$$

where Φ is the standard normal cdf and δ is the unconditional frequency of recessions.

The first policy is an unemployment subsidy that raises the flow value of unemployment when $\mathbb{1}_t^R = 1$. The total payoff in unemployment becomes

$$b_t(\Omega) = \left(1 + \xi_b\, \mathbb{1}_t^R\right) A_t\, h\, (\bar{b}_0 + \bar{b}_1\, z), \tag{40}$$

so that in recessions the worker receives a fraction ξ_b more than under the baseline policy.

The second policy is an employment subsidy that pays the firm a transfer equal to a fraction ξ_y of output when $\mathbb{1}_t^R = 1$. It is intended to capture policies that subsidize the cost of employment in downturns—for example, a temporary payroll tax cut or a wage subsidy.

The two policies operate on different margins. The unemployment subsidy provides insurance to workers near the separation margin exactly when job-finding rates are low and marginal utility is high, but it weakens employment incentives. The employment subsidy raises the value of employment to the firm and therefore supports job creation and retention, but it is a less targeted insurance device because it reallocates resources toward more productive workers, who are more likely to remain employed. We first study these policies at empirically motivated magnitudes, then trace out the welfare–employment frontier when their sizes are chosen optimally subject to a fiscal constraint.

Pre-Specified Labor Market Policies

We first study policies at magnitudes roughly matching what has been used in practice. We set $\xi_b = 20\%$ to capture historical UI benefit extensions, and $\xi_y = 6\%$ to roughly match the employer share of Social Security contributions that a payroll tax holiday would suspend.

Table 6 reports results for each policy in isolation and combined, at two activation thresholds $\delta = 10\%$ and $\delta = 20\%$ that bracket the empirical frequency of NBER recession months (13% over 1960–2019). The tax rate τ needed to fund these policies ranges from 1.0 to 3.1 percent. Despite the balanced-budget requirement, both policies raise ex-ante welfare: workers value transfers in bad states more than shareholders value the taxes that fund them. Welfare gains are larger when intervention is rarer. At $\delta = 10\%$, welfare rises by 1.7% under the unemployment subsidy, 0.7% under the employment subsidy, and 2.2% when the two are combined; at $\delta = 20\%$, the corresponding gains are 1.4%, 0.4%, and 1.7%. The combination delivers the largest welfare gains in either case.

These welfare gains do not come with gains in output or consumption, and more frequent intervention amplifies the effects on average quantities. The unemployment subsidy reduces average

GDP by 2.3% at $\delta = 10\%$ and 4.6% at $\delta = 20\%$, and raises the average unemployment rate by 3.5 and 7.6 percentage points. The employment subsidy moves in the opposite direction: GDP rises by 1.2 to 1.6%, and average unemployment falls by 0.6 to 0.8 percentage points. The combined policy delivers the largest welfare gains, but output and employment still fall relative to the baseline: the unemployment subsidy's disincentive effects outweigh the employment subsidy's incentive effects.

These endogenous responses to the policies are central to the welfare gains. Holding employment and wage dynamics fixed (the "no behavioral response" rows of Table 6), the policies look both cheaper and less valuable: at $\delta = 10\%$, the unemployment subsidy requires a tax of 0.3 rather than 1.2 percent and raises welfare by 1.1 rather than 1.7 percent, and the combined policy delivers 1.7 rather than 2.2 percent. With behavior frozen, the subsidy induces no additional take-up, so the policy reaches fewer workers. Its budget therefore balances at a lower tax, and, transferring fewer resources to the states where doing so raises welfare, it also delivers smaller welfare gains. Taking the income process as given thus understates both what these policies cost and what they achieve—like the valuation exercises above, evaluating policy requires a model in which employment and earnings respond to the policy environment rather than being held fixed.

Figure 17 decomposes the welfare gains by initial productivity z_0 for $\delta = 10\%$. The unemployment subsidy mainly benefits workers in the left tail of the distribution, with the largest gains for workers near the employment margin, who value additional unemployment insurance the most. The employment subsidy instead benefits workers in the right tail, where its incidence is higher because these workers are more likely to be employed; for the far left tail, its welfare effect is slightly negative. The combined policy delivers positive welfare gains throughout the distribution, peaking among middle-productivity workers, who benefit from both insurance at the bottom and employment support to the right.

Optimal Labor Market Policies

We now solve for the policy magnitudes (ξ_b, ξ_y) that maximize ex-ante welfare subject to two constraints: the consumption tax τ that finances the policy cannot exceed 10%, and the average employment-to-population ratio is fixed at a target $\bar{e}$. Varying $\bar{e}$ traces out a welfare–employment frontier. We hold the recession threshold at $\delta = 10\%$ throughout.

Figure 18 plots this frontier; Table 7 reports the underlying policies and outcomes at four points along it. The frontier is hump-shaped: welfare rises with employment on the left, peaks, then declines. Column (1) sits at the peak—the welfare-maximizing policy when no employment constraint binds, with average employment 2.4 percentage points below the no-policy baseline. Columns (2)–(4) lie on the right side of the peak, with average employment at the baseline, 2 percentage points above it, and 4 percentage points above it. The fiscal constraint binds throughout: $\tau = 10\%$ in every column.

Welfare gains are large everywhere on the frontier, ranging from 9.5% to 12.7%—four to six times larger than the 2.2% gain delivered by the combined pre-specified policy at the same threshold.

The composition of the optimal policy shifts sharply along the frontier. At the welfare-maximizing point, the policy is heavily insurance-loaded: the unemployment subsidy nearly doubles the flow payoff in unemployment during recessions ($\xi_b = 92\%$), while the employment subsidy is much smaller ($\xi_y = 18\%$). As the employment target rises along the frontier, the composition shifts: ξ_b falls from 92% to 47%, while ξ_y rises from 18% to 52%. At the high-employment end of the frontier, the employment subsidy is the dominant instrument.

Three features of the frontier stand out. First, the welfare–employment trade-off is mild near the welfare maximum: the planner preserves almost all of the welfare gain at column (1) while bringing average employment back to the no-policy baseline (column (2)). The trade-off steepens further out: pushing average employment 4 percentage points above the baseline (column (4)) costs 3.2% of welfare. Second, average consumption is below the no-policy baseline everywhere on the frontier, by 7.8% at column (1) and 3.7% at column (4). Optimal welfare-improving policies do not raise mean consumption: the gains come from reallocating consumption toward bad aggregate states and toward workers near the separation margin, where marginal utility is highest. Third, the effect on the volatility of employment and earnings growth flips sign along the frontier. At the welfare-maximizing point, income growth volatility rises by 4.5 percentage points relative to the no-policy baseline; by column (4), it falls by 2.3 percentage points. The unemployment subsidy reduces consumption risk by transferring resources to bad states, but it also raises employment risk by weakening labor supply incentives; the employment subsidy offsets this disincentive by directly supporting employment.

Discussion and Conclusion

We show how aggregate financial conditions shape not only asset prices but also labor income risk. In our model, higher risk premia reduce match surplus, tighten limited commitment constraints, and erode the insurance firms provide to workers. The result is state-dependent pass-through of firm shocks to earnings, with countercyclical amplification concentrated among lower-paid workers near the separation margin. The data support this prediction: the pass-through of firm-specific TFP shocks to worker earnings rises when risk premia are elevated, especially through job destruction at the bottom of the earnings distribution. The calibrated model matches these targeted moments and accounts for a broader set of untargeted facts about income risk, including non-Gaussian earnings growth, countercyclical left-tail risk, and persistent scarring from displacement and recession entry. It also implies that the welfare cost of idiosyncratic risk is large, but smaller than in exogenous-income models, especially for low-wage workers most exposed to separation.

We close by drawing out several broader implications of the framework.

Financial conditions and monetary policy. A large literature argues that risk premia are central to monetary transmission and that central banks should stabilize financial conditions (Caballero and Simsek, 2020). Our framework points to a complementary channel: because risk premia govern the extent of firm insurance and the incidence of job destruction, policies that stabilize financial conditions would also stabilize the labor income risk faced by workers, which could in turn stabilize aggregate demand.

Human capital valuation and portfolio choice. Because workers cannot diversify their idiosyncratic income risk, they discount their future labor income at rates well above the risk-free rate—especially for low-wage and low-experience workers, whose earnings are the most volatile and least insured. Further, in our model younger workers' human capital is more exposed to variation in risk premia, and hence stock market fluctuations, than the human capital of older workers. This implication runs against the textbook lifecycle view that early-career human capital is bond-like and weakens the standard case for tilting young workers' portfolios toward equities.

Measuring the cost of job loss. A large literature gauges the cost of job loss by the size and persistence of displaced workers' earnings losses (Davis and Von Wachter, 2011). Because the value of nonemployment replaces much of the income lost to a separation, these earnings losses potentially overestimate the welfare cost of displacement and misrank its incidence: welfare losses are far smaller than earnings losses, and—since the wedge is largest at the bottom—the welfare cost of job loss rises with earnings even as earnings losses stay roughly flat. For workers who separate endogenously, the welfare cost is minimal, since they are essentially at their outside option—which is also why firms let low-surplus matches end rather than retain workers at lower pay (Bewley, 1999).

Designing countercyclical labor market policy. State-contingent transfers deliver substantial welfare gains, and the mechanism shapes their optimal design. When risk premia rise, firm-provided insurance erodes exactly when marginal utility is highest, so public transfers are most valuable in those states and for the low-wage workers who face the largest idiosyncratic income risk. Unemployment subsidies supply this insurance but blunt employment incentives; employment subsidies support jobs and output but reach workers who are less likely to separate, and so provide insurance less precisely. The welfare-maximizing policy leans heavily on insurance, while sustaining higher employment shifts the mix toward employment subsidies—a tradeoff that is invisible in models where wages and employment do not respond to policy.

References

Acabbi, E. M., A. Alati, and L. Mazzone (2026). The collapse of human capital ladders in recessions. Working paper, University of Mannheim.

Afrouzi, H., A. Blanco, A. Drenik, and E. Hurst (2024, December). A theory of how workers keep up with inflation. Working Paper 33233, National Bureau of Economic Research.

Ai, H. and A. Bhandari (2021). Asset pricing with endogenously uninsurable tail risk. *Econometrica 89*(3), 1471–1505.

Autor, D. H., D. Dorn, G. H. Hanson, and J. Song (2014). Trade adjustment: Worker-level evidence. *Quarterly Journal of Economics 129*(4), 1799–1860.

Balke, N. and T. Lamadon (2022, July). Productivity shocks, long-term contracts, and earnings dynamics. *American Economic Review 112*(7), 2139–77.

Bauer, M. D., B. S. Bernanke, and E. Milstein (2023). Risk appetite and the risk-taking channel of monetary policy. *The Journal of Economic Perspectives 37*(1), pp. 77–100.

Bewley, T. F. (1999). *Why Wages Don't Fall during a Recession*. Cambridge, MA: Harvard University Press.

Caballero, R. J. and A. Simsek (2020, 03). A Risk-Centric Model of Demand Recessions and Speculation. *The Quarterly Journal of Economics 135*(3), 1493–1566.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Catherine, S. (2022). Countercyclical labor income risk and portfolio choices over the life cycle. *Review of Financial Studies 35*(9), 4016–4054.

Chan, M., M. Xu, and S. Salgado (2019). Heterogeneous passthrough from tfp to wages. 2019 Meeting Papers 1447, Society for Economic Dynamics.

Chodorow-Reich, G. and L. Karabarbounis (2016). The cyclicality of the opportunity cost of employment. *Journal of Political Economy 124*(6), 1563–1618.

Constantinides, G. M. (2024). Welfare costs of idiosyncratic and aggregate consumption shocks. *Review of Asset Pricing Studies 14*(4), 596–638.

Couch, K. A. and D. W. Placzek (2010). Earnings losses of displaced workers revisited. *The American Economic Review 100*(1), 572–589.

Davis, S. and T. Von Wachter (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity 42*(2 (Fall)), 1–72.

De Nardi, M., G. Fella, and G. Paz-Pardo (2020). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association 18*(2), 890–926.

den Haan, W. J., G. Ramey, and J. Watson (2000, June). Job destruction and propagation of shocks. *American Economic Review 90*(3), 482–498.

Diamond, P. A. (1982). Wage determination and efficiency in search equilibrium. *The Review of Economic Studies 49*(2), 217–227.

Elsby, M. W., B. Hobijn, and A. Şahin (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics 72*, 64–82.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969.

Farber, H. S. (2017). Employment, hours, and earnings consequences of job loss: US evidence from the displaced workers survey. *Journal of Labor Economics 35*(S1), S235–S272.

Friedrich, B., L. Laun, C. Meghir, and L. Pistaferri (2019, April). Earnings dynamics and firm-level shocks. Working Paper 25786, National Bureau of Economic Research.

Gilchrist, S. and E. Zakrajšek (2012, June). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

Green, B., L. Kogan, D. Papanikolaou, and L. D. W. Schmidt (2025). Winners and losers: Competition, creative destruction, and labor income risk. Working paper, MIT Sloan School of Management.

Guiso, L., L. Pistaferri, and F. Schivardi (2005). Insurance within the firm. *Journal of Political Economy 113*(5), 1054–1087.

Guvenen, F., G. Kaplan, J. Song, and J. Weidner (2022). Lifetime Earnings in the United States over Six Decades. *American Economic Journal: Applied Economics 14*(4), 446–79.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and R. Madera (2024). Consumption and welfare effects of earnings shocks: Non-Gaussian features and hand-to-mouth households. *Journal of Economic Dynamics and Control 163*, 104876.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Hagedorn, M. and I. Manovskii (2008, September). The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review 98*(4), 1692–1706.

Hall, R. E. (2017, February). High discounts and high unemployment. *American Economic Review 107*(2), 305–30.

Hansen, L. and T. J. Sargent (2001, May). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

Hubmer, J. (2018). The job ladder and its implications for earnings risk. *Review of Economic Dynamics 29*, 172–194.

Huckfeldt, C. (2022). Understanding the scarring effect of recessions. *American Economic Review 112*(4), 1273–1310.

Huggett, M. and G. Kaplan (2016). How large is the stock component of human capital? *Review of Economic Dynamics 21*, 23–40.

İmrohoroğlu, A. and Ş. Tüzel (2014). Firm-level productivity, risk, and return. *Management Science 60*(8), 2073–2090.

Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica 91*(3), 855–898.

Jung, P. and K. Kuester (2015). Optimal Labor-Market Policy in Recessions. *American Economic Journal: Macroeconomics 7*(2), 124–156.

Jung, P. and M. Kuhn (2019). Earnings losses and labor mobility over the life cycle. *Journal of the European Economic Association 17*(3), 678–724.

Jurado, K., S. C. Ludvigson, and S. Ng (2015, March). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216.

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics 17*(2), 303–316.

Kehoe, P. J., P. Lopez, V. Midrigan, and E. Pastorino (2023). Asset Prices and Unemployment Fluctuations: A Resolution of the Unemployment Volatility Puzzle. *Review of Economic Studies 90*(3), 1304–1357.

Kehoe, P. J., V. Midrigan, and E. Pastorino (2019). Debt constraints and employment. *Journal of Political Economy 127*(4), 1926–1991.

Kekre, R. (2023). Unemployment Insurance in Macroeconomic Stabilization. *The Review of Economic Studies 90*(5), 2439–2480.

Krebs, T. (2007, June). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

Kroft, K. and M. J. Notowidigdo (2016). Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence. *The Review of Economic Studies 83*(3), 1092–1124.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin (2017, November). Gross worker flows over the business cycle. *American Economic Review 107*(11), 3447–76.

Landais, C., P. Michaillat, and E. Saez (2018a). A Macroeconomic Approach to Optimal Unemployment Insurance: Applications. *American Economic Journal: Economic Policy 10*(2), 182–216.

Landais, C., P. Michaillat, and E. Saez (2018b). A Macroeconomic Approach to Optimal Unemployment Insurance: Theory. *American Economic Journal: Economic Policy 10*(2), 152–181.

Lettau, M. and J. A. Wachter (2007). Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium. *The Journal of Finance 62*(1), 55–92.

Ljungqvist, L. and T. J. Sargent (1998). The European unemployment dilemma. *Journal of political Economy 106*(3), 514–550.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan (2013). The Wealth-Consumption Ratio. *The Review of Asset Pricing Studies 3*(1), 38–94.

Malgieri, C. and L. Citino (2025). Wage contracts and financial frictions. Job Market Paper, January 26, 2025.

Margolis, D. N. and J. Montana (2024). Who gets to stay? How mass layoffs reshape firms' skills structure. IZA Discussion Paper 17426, Institute of Labor Economics (IZA).

McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics 88*, 1–14.

McKay, A. and R. Reis (2021). Optimal Automatic Stabilizers. *The Review of Economic Studies 88*(5), 2375–2406.

Meeuwis, M., D. Papanikolaou, J. L. Rothbaum, and L. D. Schmidt (2025). Time-varying risk premia and heterogeneous labor market dynamics. Working paper.

Menzio, G., I. A. Telyukova, and L. Visschers (2016). Directed search over the life cycle. *Review of Economic Dynamics 19*, 38–62.

Michelacci, C. and V. Quadrini (2009). Financial markets and wages. *Review of Economic Studies 76*(2), 795–827.

Mitman, K. and S. Rabinovich (2015). Optimal Unemployment Insurance in an Equilibrium Business-Cycle Model. *Journal of Monetary Economics 71*, 99–118.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy 105*(2), 385–411.

Montgomery, J. D. (1991). Equilibrium wage dispersion and interindustry wage differentials. *The Quarterly Journal of Economics 106*(1), 163–179.

Mortensen, D. T. (1982). The matching process as a noncooperative bargaining game. *The Economics of Information and Uncertainty*, 233–258.

Moscarini, G. and F. Postel-Vinay (2018, September). On the Job Search and Business Cycles. IZA Discussion Papers 11853, Institute of Labor Economics (IZA).

Mukoyama, T., C. Patterson, and A. Şahin (2018, January). Job search behavior over the business cycle. *American Economic Journal: Macroeconomics 10*(1), 190–215.

Nattinger, M. B. and J. Rothbaum (2025). The Effects of Firm Volatility on Workers. Working Paper.

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review 75*(4), 676–690.

Postel-Vinay, F. and J.-M. Robin (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica 70*(6), 2295–2350.

Rudanko, L. (2011). Aggregate and idiosyncratic risk in a frictional labor market. *American Economic Review 101*(6), 2823–2843.

Schmidt, L. D. W. (2025). Climbing and falling off the ladder: Asset pricing implications of labor market event risk. *Journal of Financial Economics 172*, 104131.

Schwandt, H. and T. von Wachter (2019). Unlucky cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets. *Journal of Labor Economics 37*(S1), S161–S198.

Shimer, R. (2005, March). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review 95*(1), 25–49.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics 15*(2), 127–148.

Skiadas, C. (2013). Scale-invariant uncertainty-averse preferences and source-dependent constant relative risk aversion. *Theoretical Economics 8*(1), 59–93.

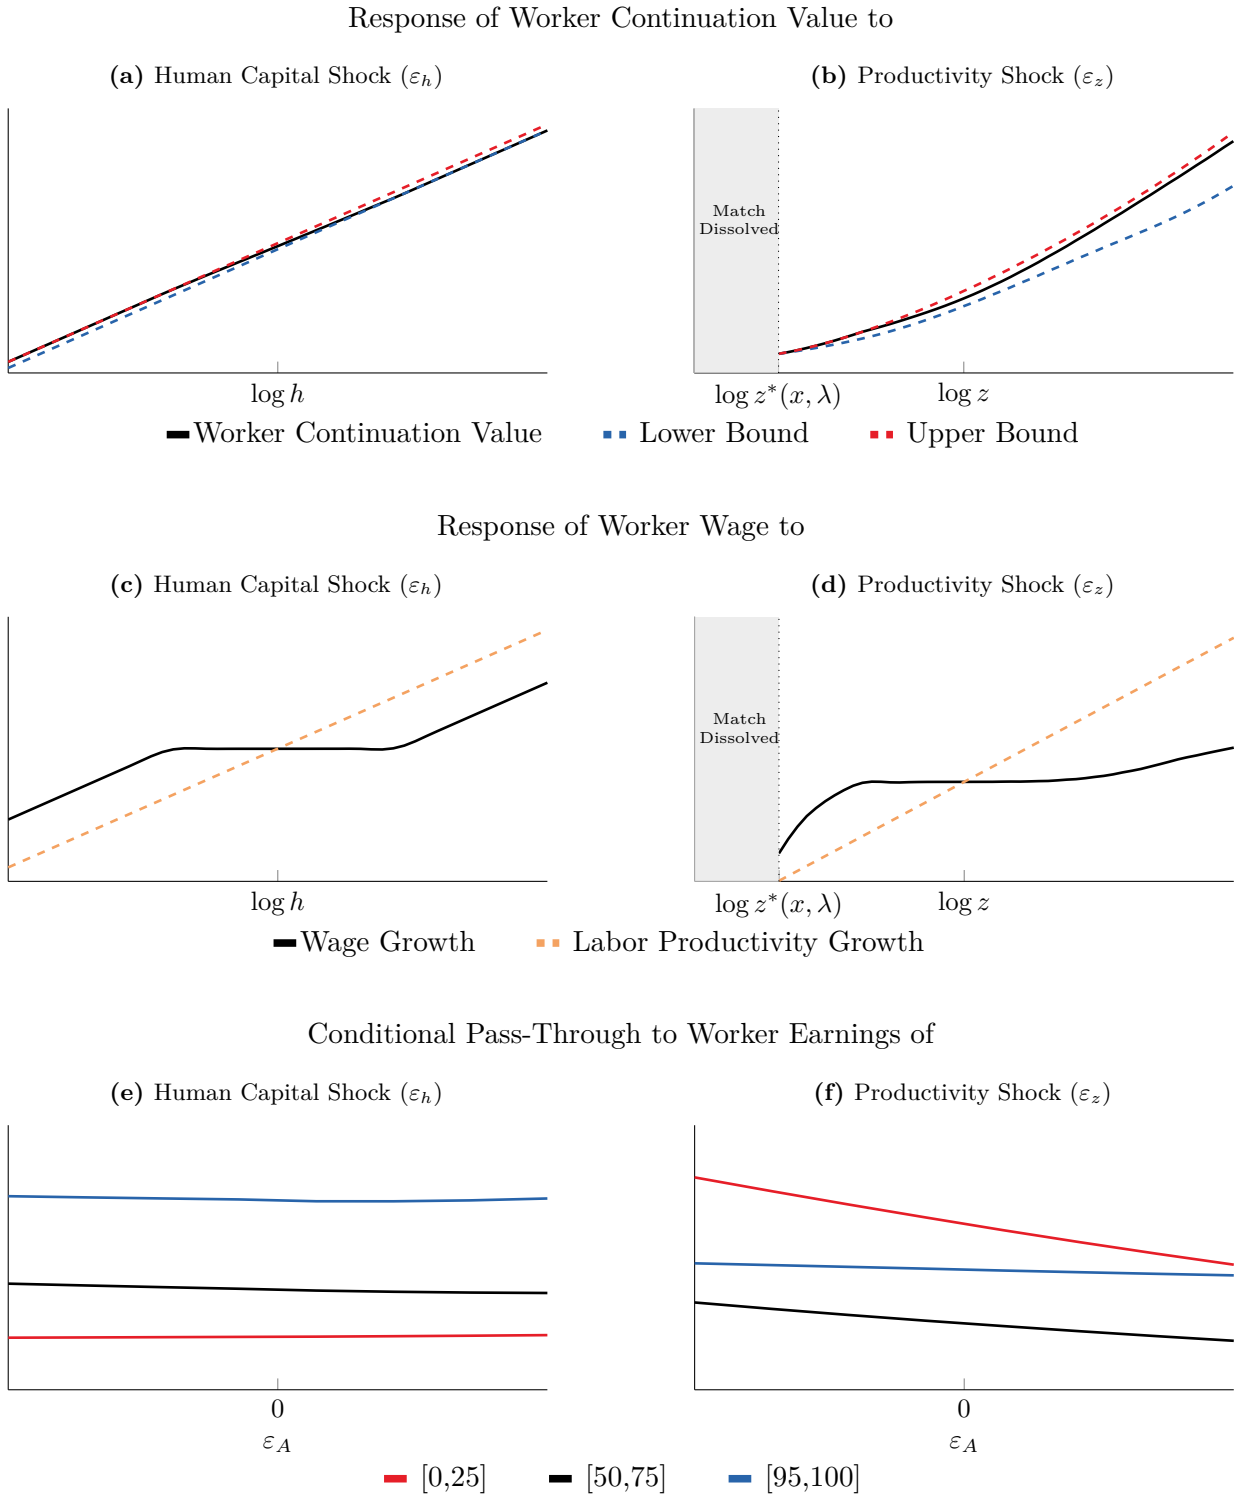
Souchier, M. (2023). The pass-through of productivity shocks to wages and the cyclical competition for workers. Working paper, June 4, 2023.

Souchier, M. (2025). Insurance inside and outside the firm. Working paper, Wharton School, University of Pennsylvania.

Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *The Review of Economic Studies 54*(4), 599–617.

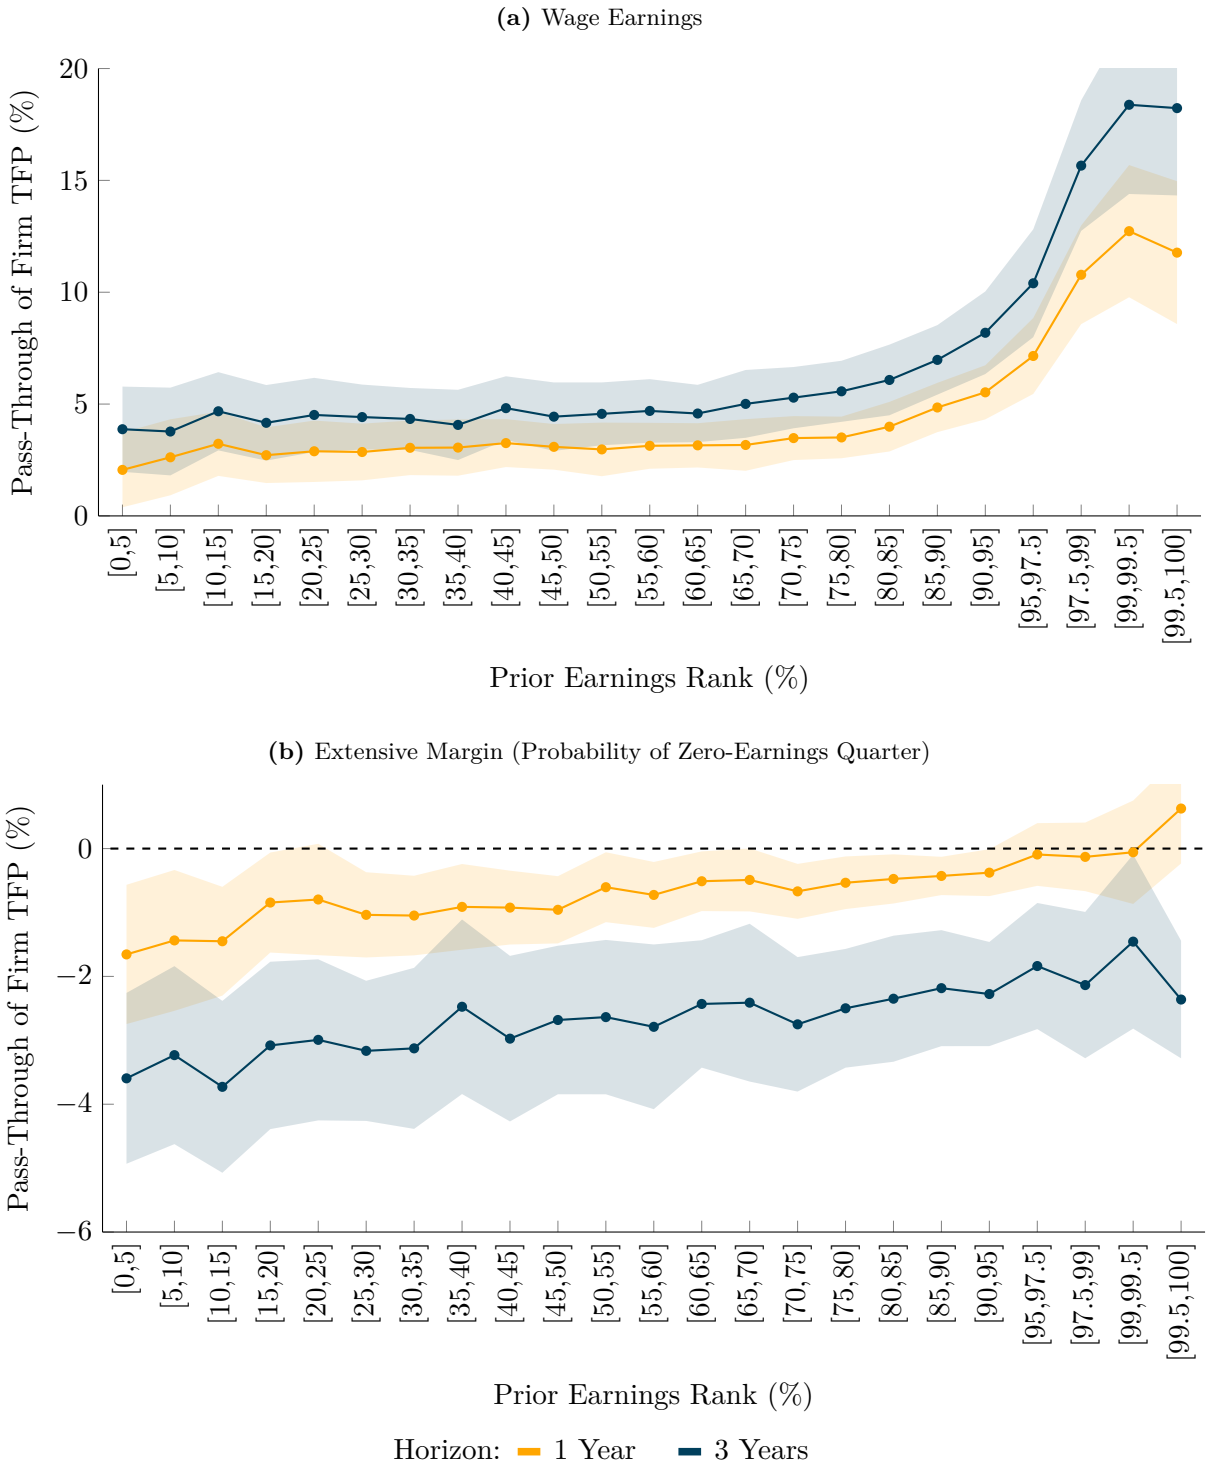
Thomas, J. and T. Worrall (1990, August). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory 51*(2), 367–390.

Figure 1: Wage and Earnings Dynamics in Model



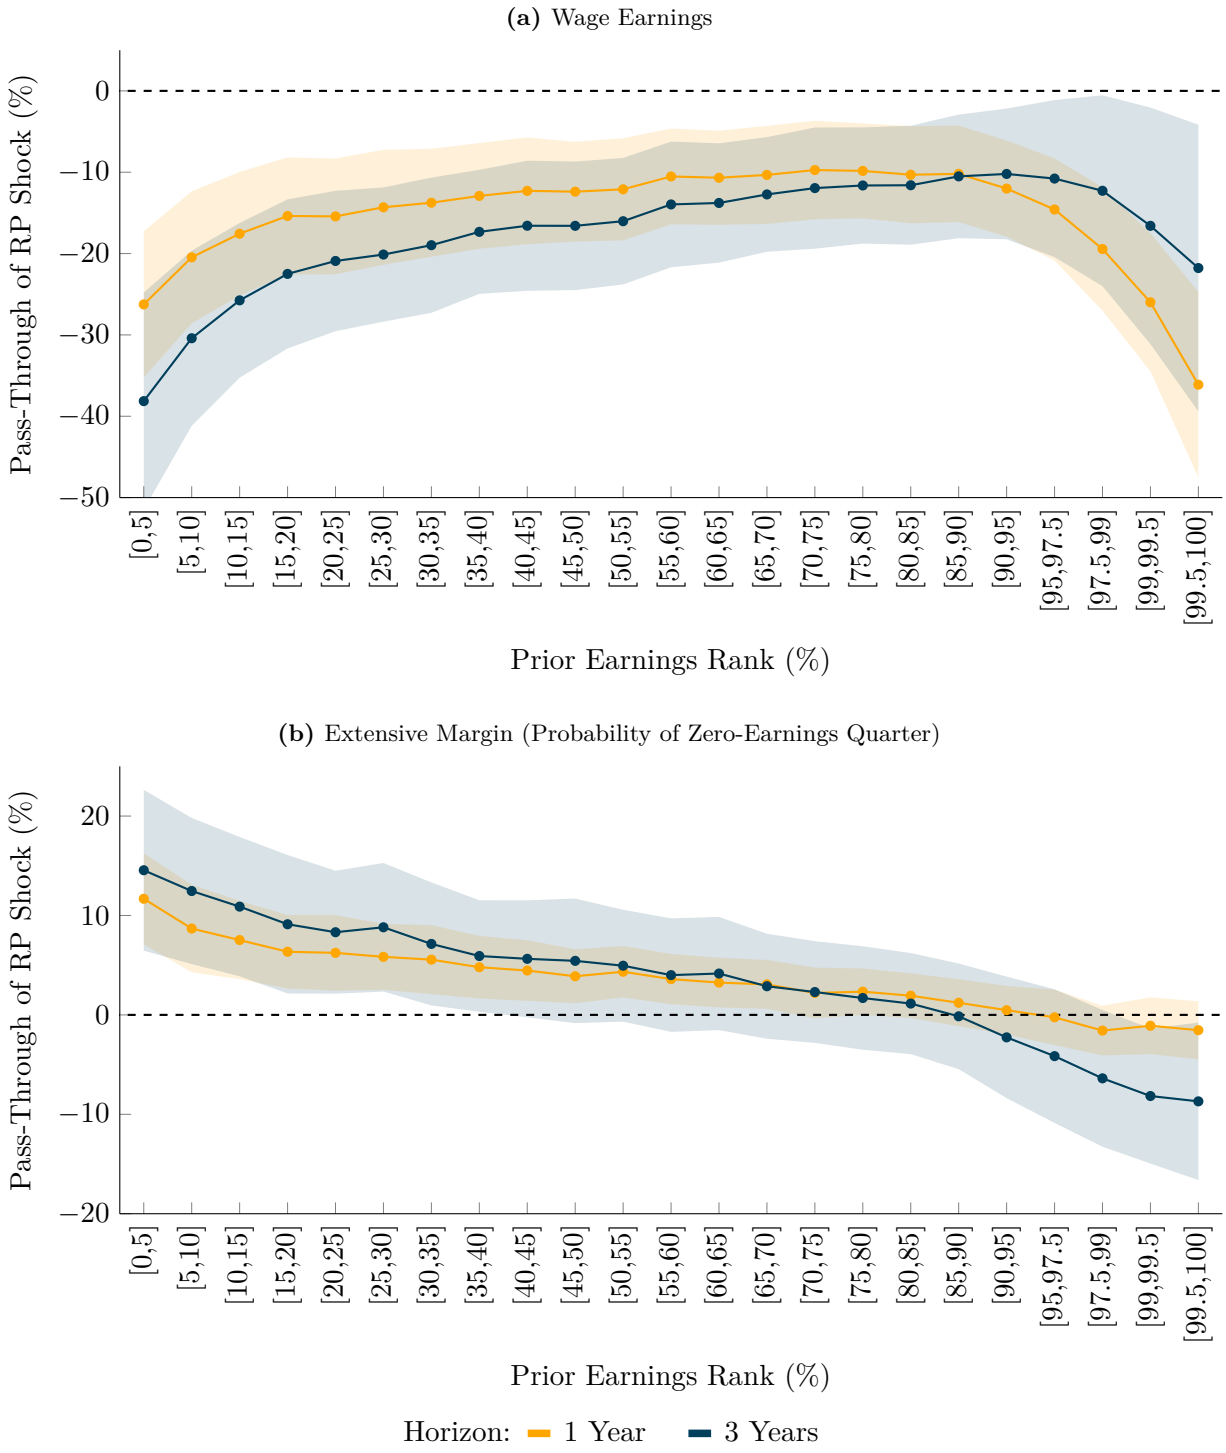
Panels (a) and (b) plot worker continuation values against the realized idiosyncratic shock, with h shocks on the left and z shocks on the right; dashed lines show the limited commitment bounds. Panels (c) and (d) plot wage growth against shock size; the dashed line shows labor productivity growth. These are for a worker with $z = \bar{z}$, $\lambda = \bar{\lambda}_H$, $x = \bar{x}$, and promised utility at the midpoint of the limited commitment bounds. Panels (e) and (f) plot the average pass-through of annual idiosyncratic shocks to annual worker earnings, by prior earnings percentile and conditional on the aggregate shock ε_A.

Figure 2: Pass-Through of Firm TFP Growth to Workers in Data



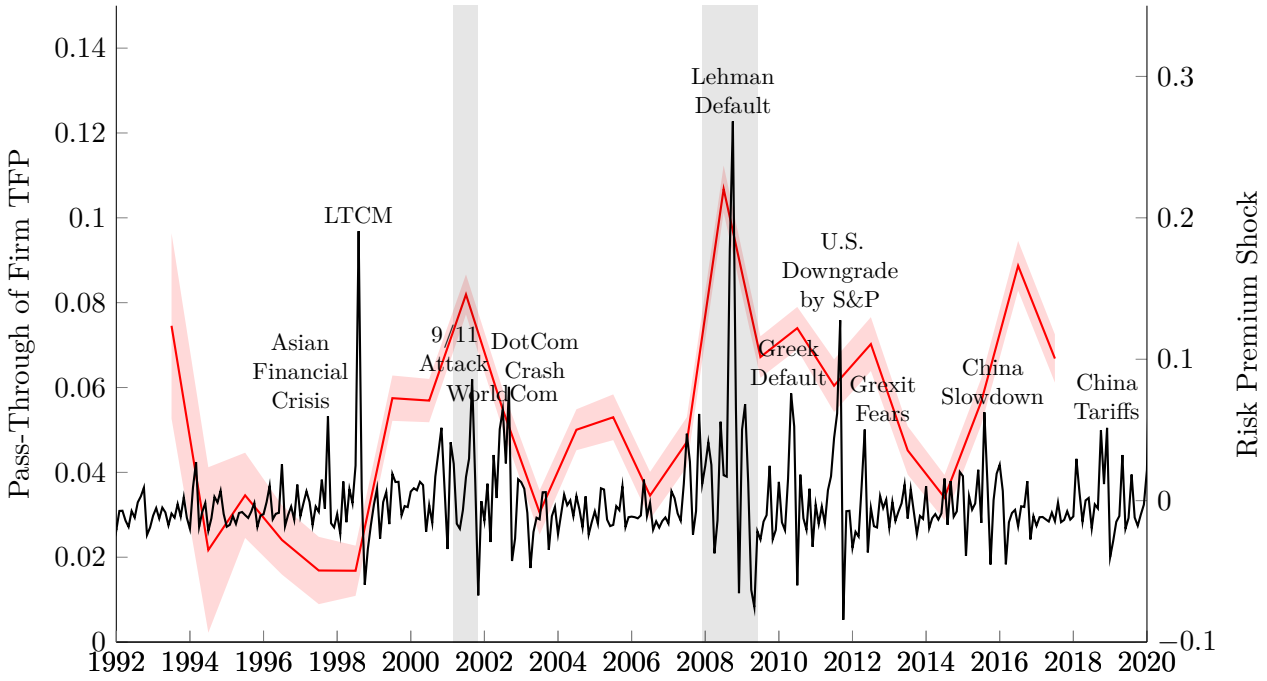
This figure plots estimated coefficients β_s from regression (31) by within-firm prior earnings percentile at horizons of one and three years. Panel (a) uses cumulative age-adjusted earnings growth $g_{i,t:t+H}$ as the outcome; panel (b) uses an indicator for at least one zero-earnings quarter over the next H years. The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019. Shaded areas are 95% confidence intervals based on standard errors clustered by worker and year.

Figure 3: Pass-Through of Risk Premium Shocks to Workers in Data



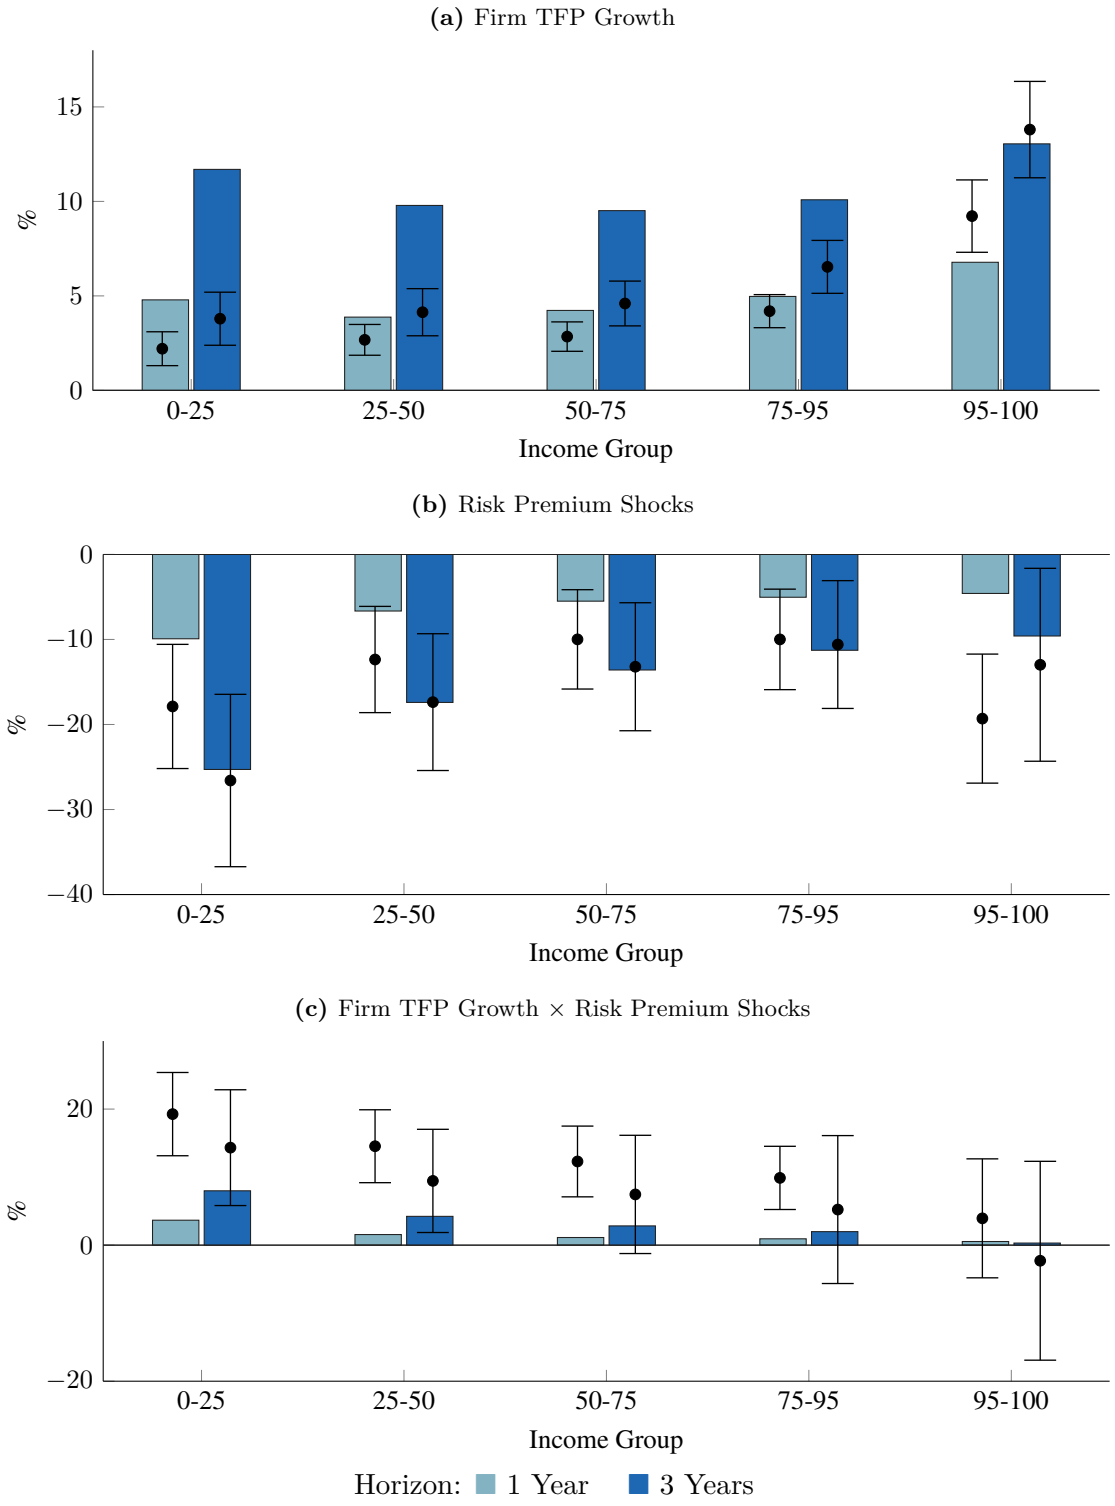
This figure plots estimated coefficients γ_s from regression (31) by within-firm prior earnings percentile at horizons of one and three years. Panel (a) uses cumulative age-adjusted earnings growth $g_{i,t:t+H}$ as the outcome; panel (b) uses an indicator for at least one zero-earnings quarter over the next H years. The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019. Shaded areas are 95% confidence intervals based on standard errors clustered by worker and year.

Figure 4: Time-Varying Pass-Through of Firm TFP Growth in Data



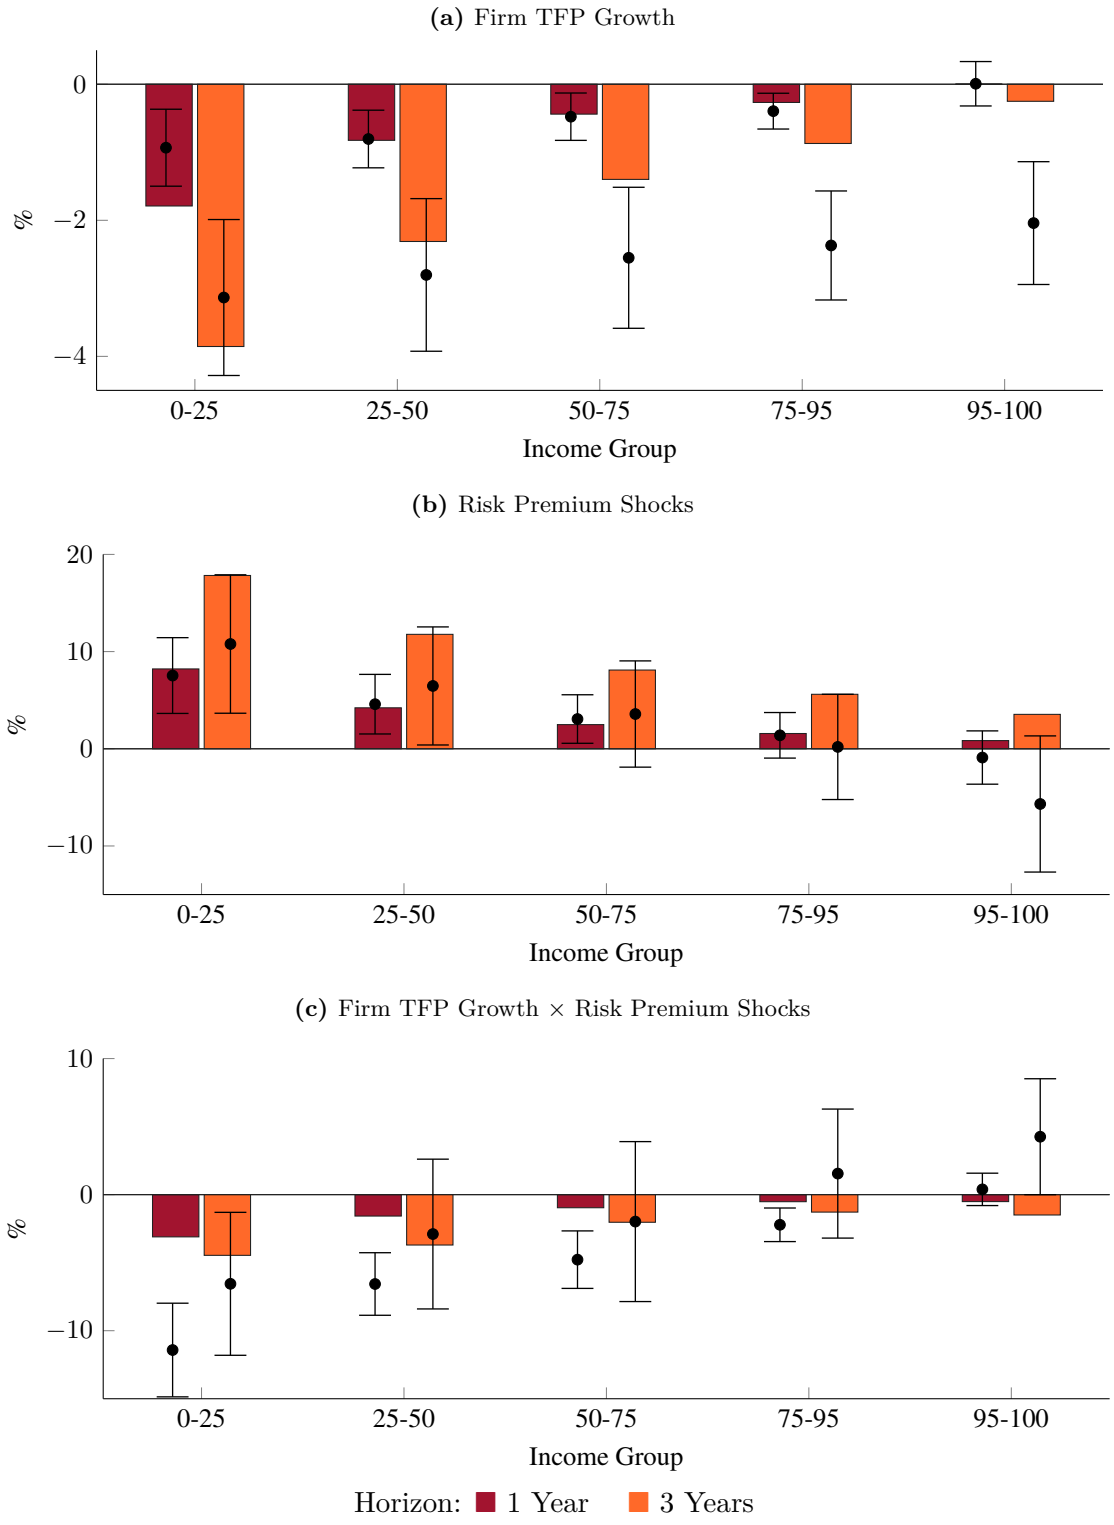
This figure plots estimated year-by-year pass-through coefficients β_t from regression (32) of three-year cumulative age-adjusted earnings growth on firm TFP growth (red, left axis), together with monthly risk premium shocks ϵ^{rp} (black, right axis). The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019. The shaded area is a 95% confidence interval based on standard errors clustered by worker and year. Gray bands indicate NBER recession dates.

Figure 5: Response of Wage Earnings: Model vs. Data



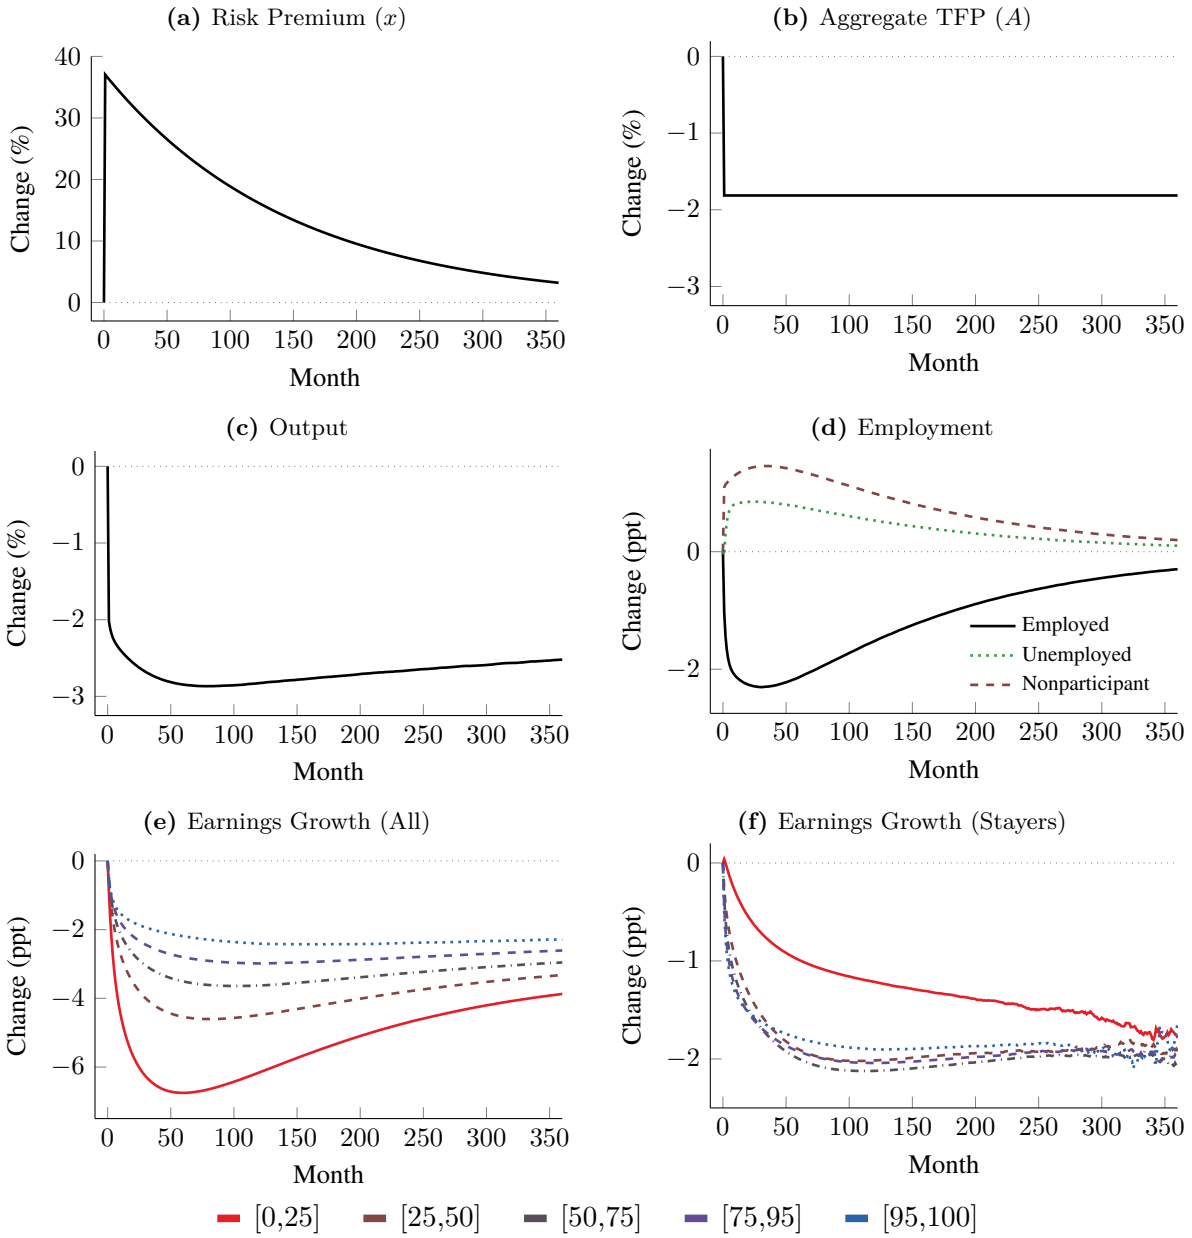
This figure compares estimated coefficients from regression (33) in the data (dots with 95% confidence intervals) and in model-simulated data (bars), by earnings group at horizons of one and three years. The outcome variable is cumulative age-adjusted earnings growth. Panel (a): pass-through of firm TFP growth ($\beta_{0,s}$). Panel (b): pass-through of risk premium shocks (γ_s). Panel (c): interaction of firm TFP growth with risk premium shocks ($\beta_{1,s}$).

Figure 6: Response of Job Destruction: Model vs. Data



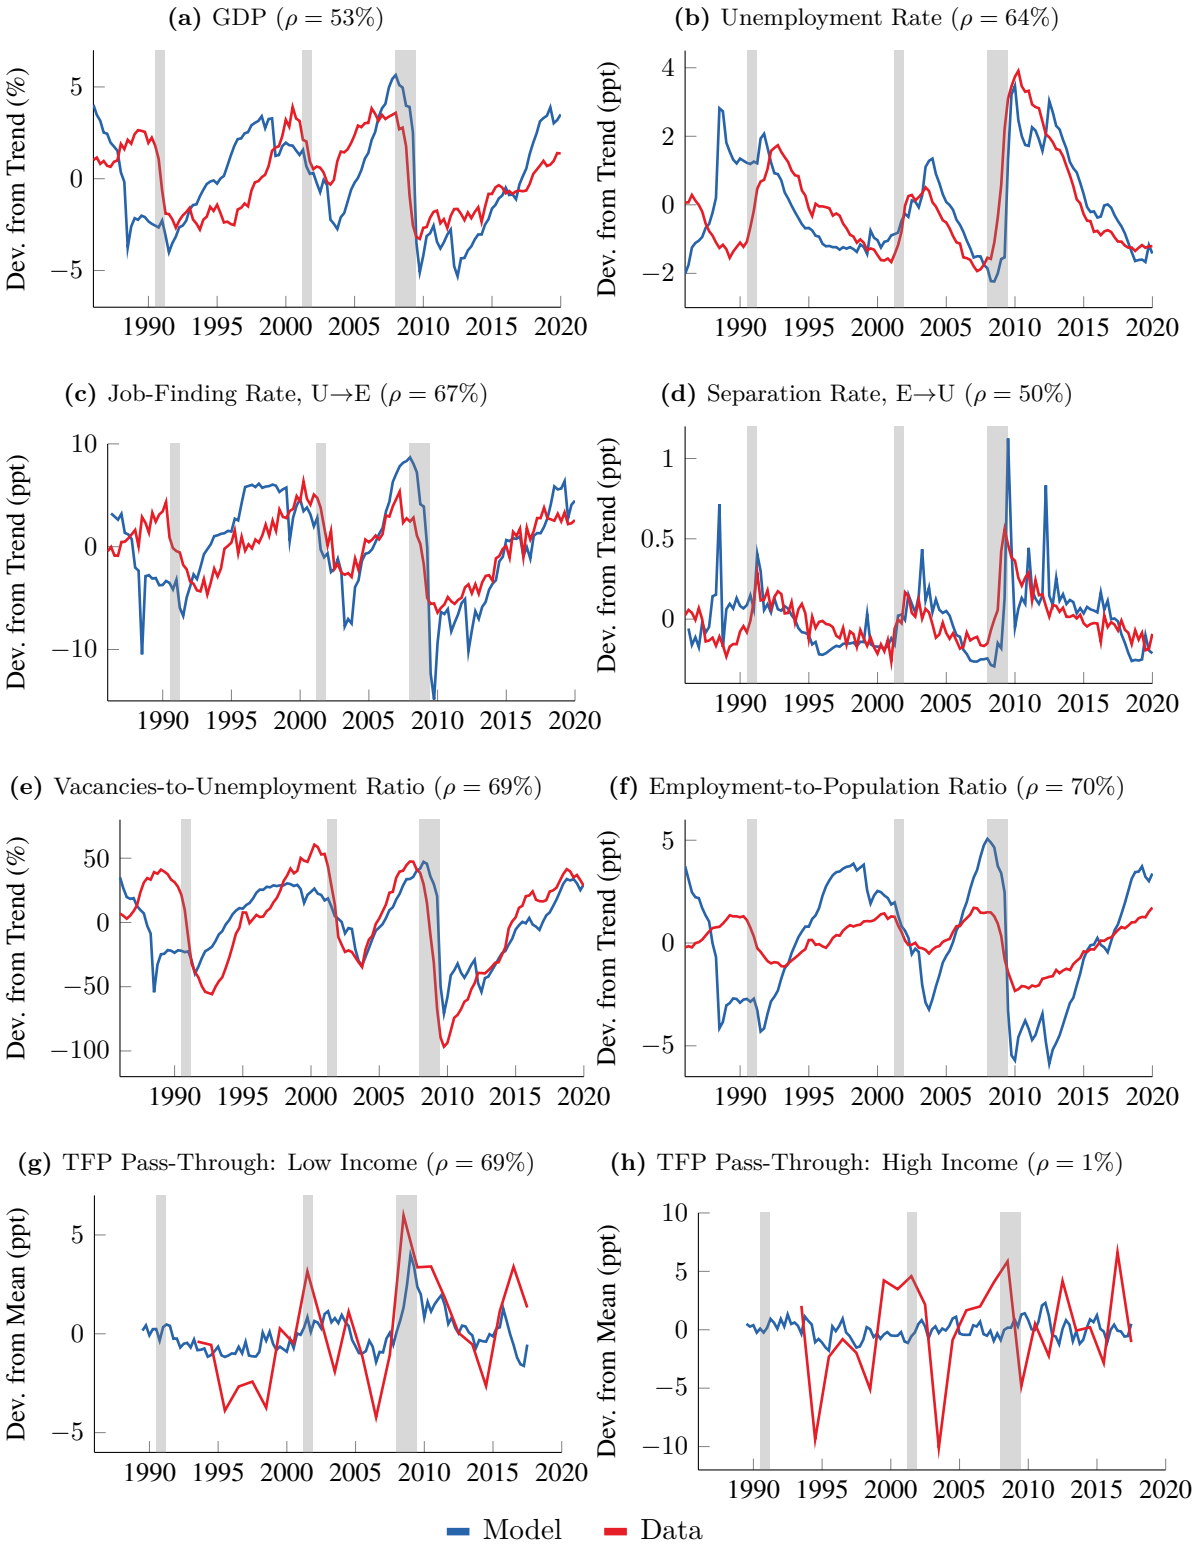
This figure compares estimated coefficients from regression (33) in the data (dots with 95% confidence intervals) and in model-simulated data (bars), by earnings group at horizons of one and three years. The outcome variable is an indicator for at least one zero-earnings quarter over the next H years. Panel (a): pass-through of firm TFP growth ($\beta_{0,s}$). Panel (b): pass-through of risk premium shocks (γ_s). Panel (c): interaction of firm TFP growth with risk premium shocks ($\beta_{1,s}$).

Figure 7: Impulse Responses to Aggregate Shock in Model



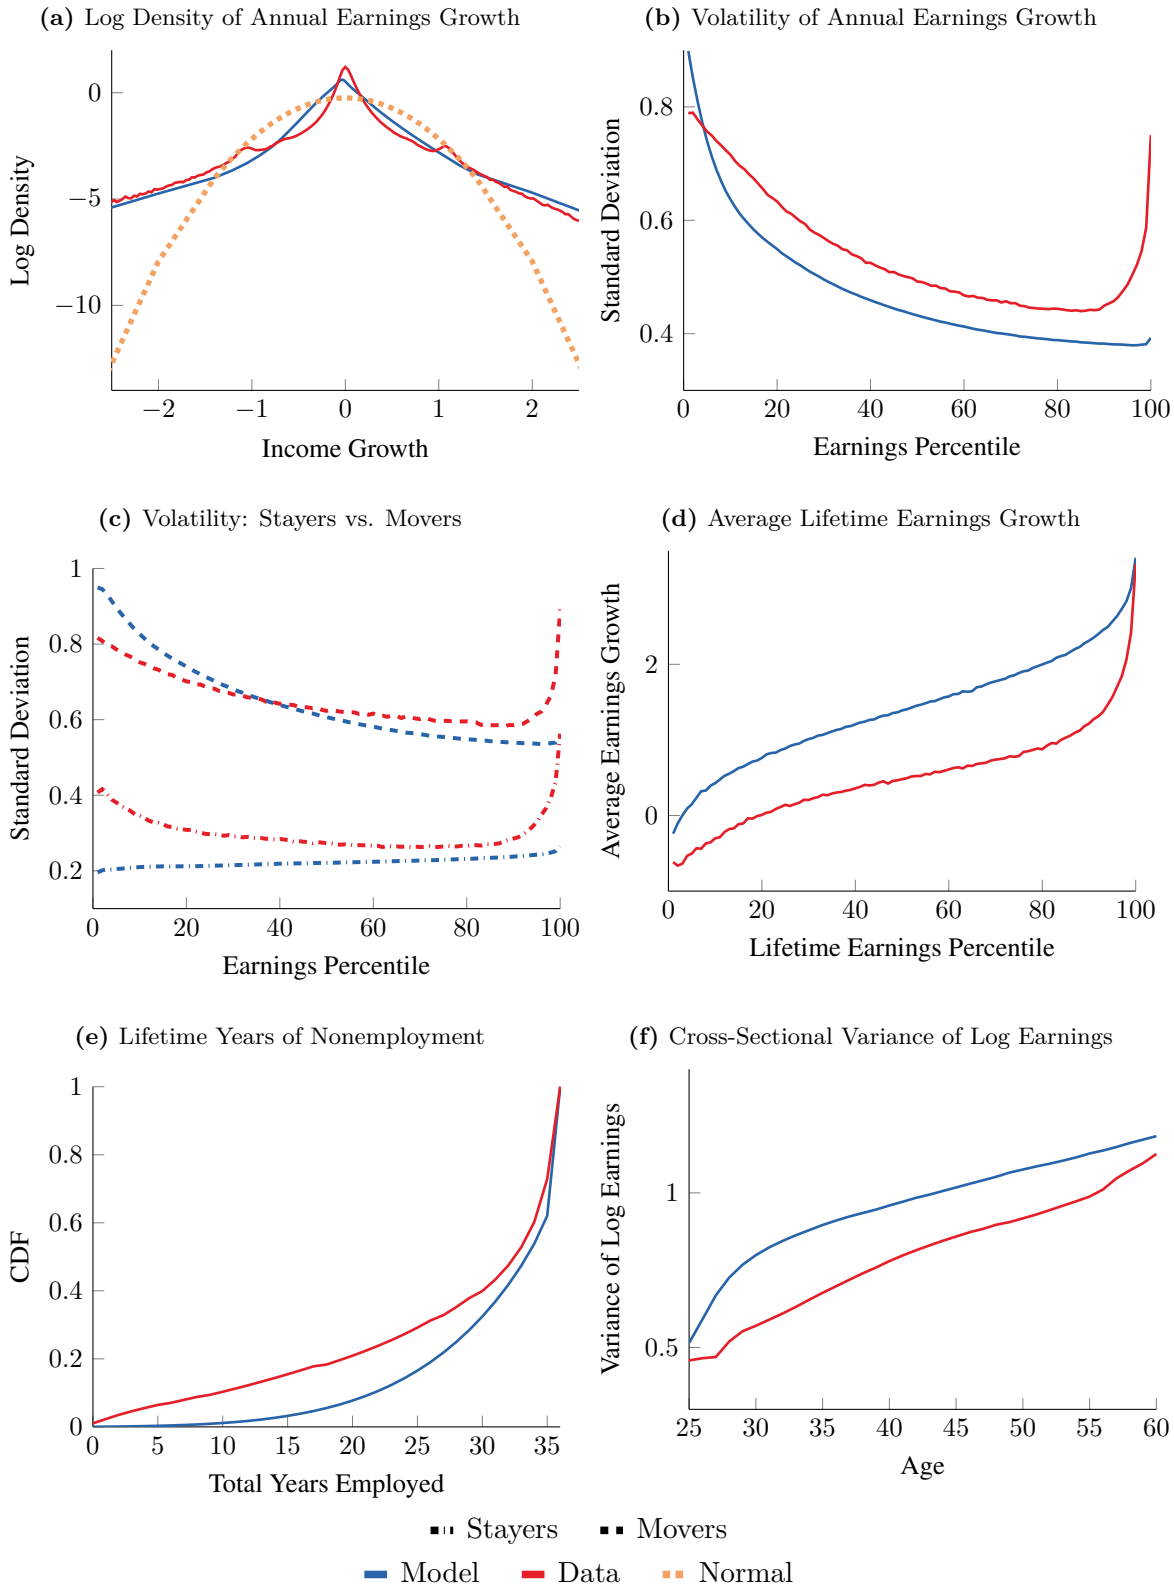
This figure shows the impulse responses of key model quantities following an aggregate shock of one annual standard deviation.

Figure 8: Realized Fluctuations in Labor Markets: Model vs. Data



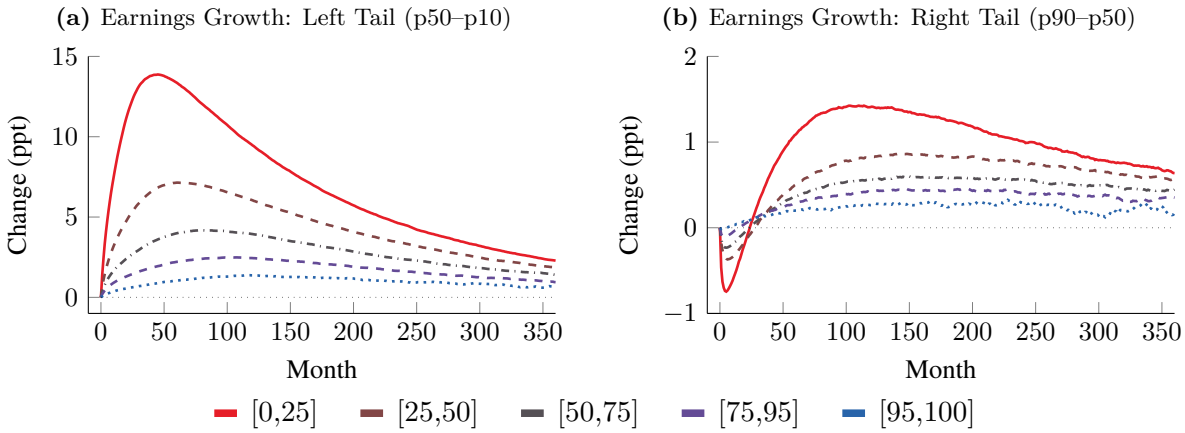
This figure compares the realized paths of key variables between the model and the data. We directly feed into the model our (scaled) empirical measures of risk premium shocks ϵ^{rp}. We detrend all macroeconomic series using an HP filter with quarterly smoothing parameter 10^5.

Figure 9: Labor Income Risk: Model vs. Data



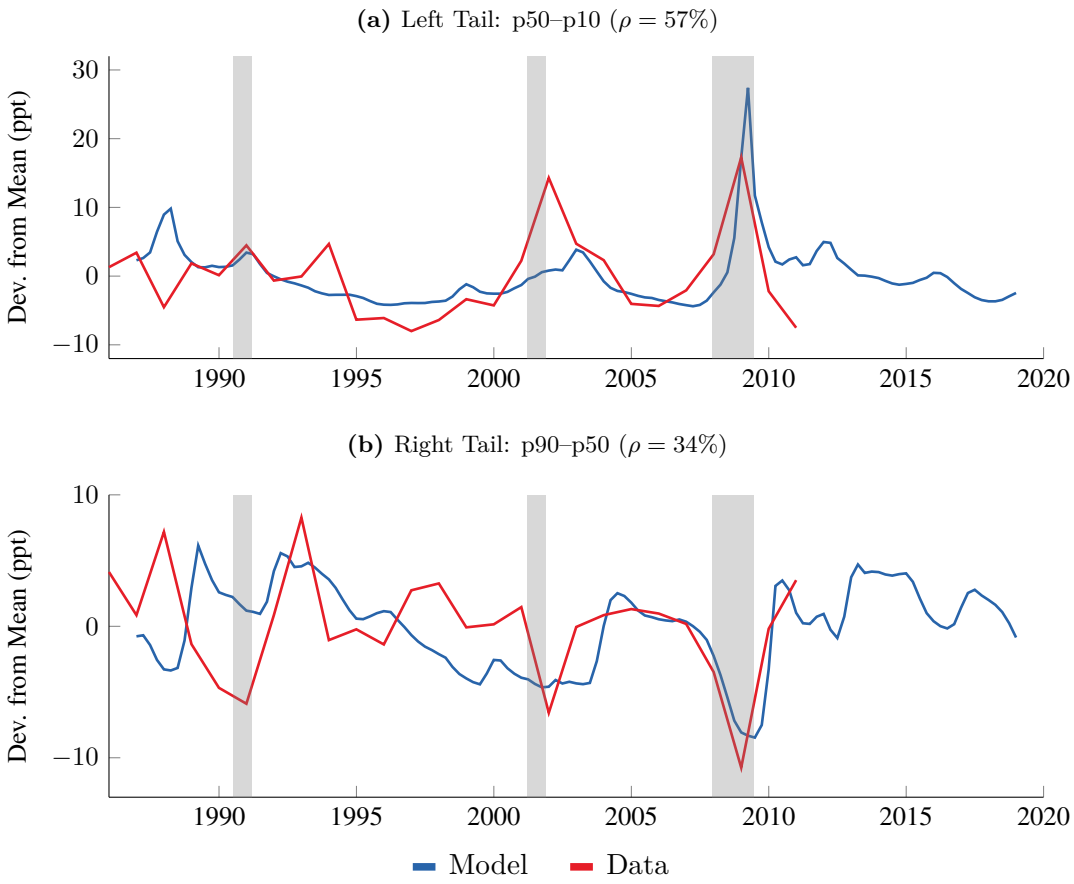
This figure compares key measures of worker earnings risk in the model and in the data. The data series are from Guvenen et al. (2021).

Figure 10: Impulse Responses of Earnings Growth Distribution to Aggregate Shock in Model



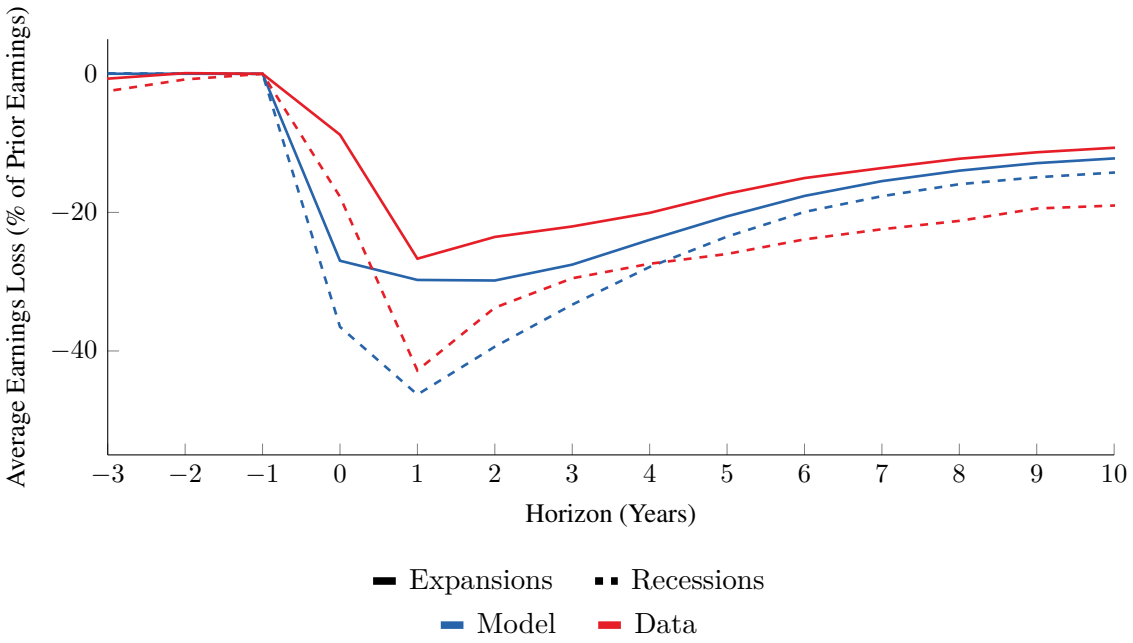
This figure shows the impulse responses of the left and right tails of earnings growth following an aggregate shock of one annual standard deviation.

Figure 11: Realized Fluctuations in Earnings Risk: Model vs. Data



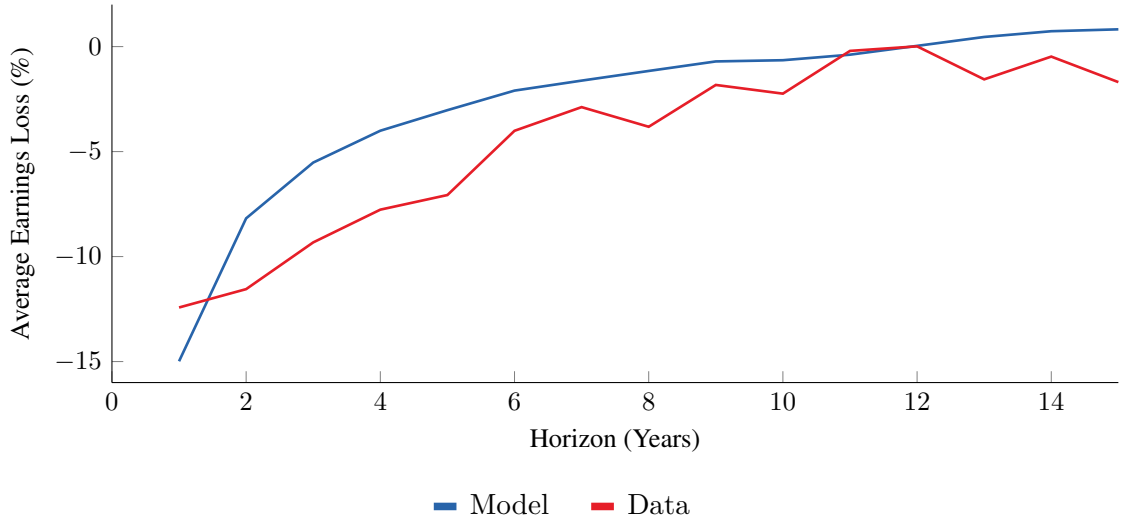
This figure compares model-implied and empirical measures of earnings risk over time, feeding the model with realized risk premium shocks ϵ^{rp}. Panel (a) plots the difference between the median and the 10th percentile of annual earnings growth (left-tail risk); panel (b) plots the difference between the 90th percentile and the median (right-tail risk). Data series are from Guvenen et al. (2014).

Figure 12: Earnings Losses Due to Job Displacement: Model vs. Data



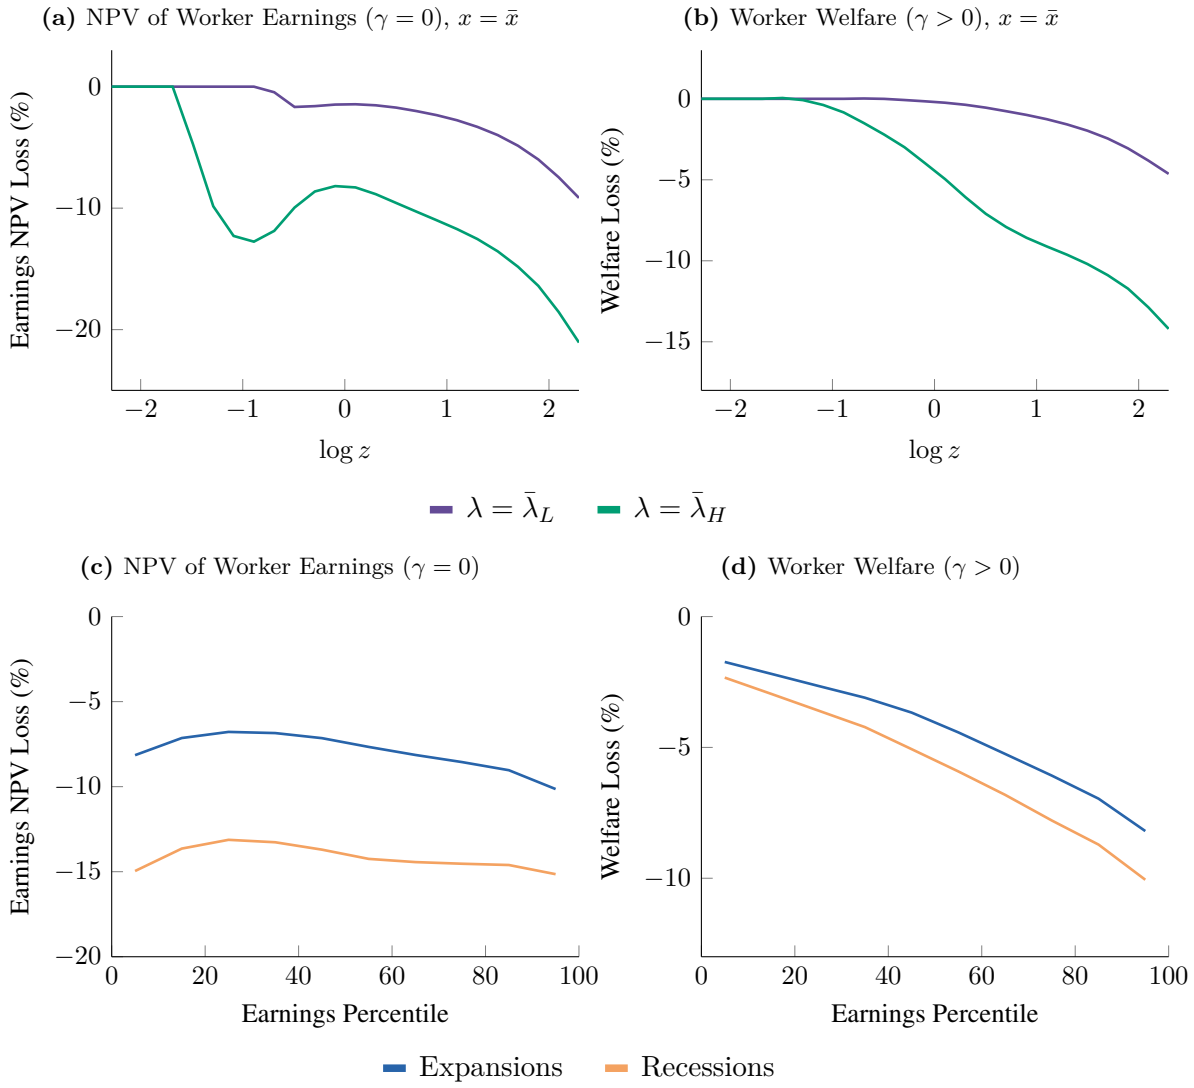
This figure plots average earnings around a displacement event relative to a control group of otherwise similar workers who are not displaced, following the empirical design of Davis and Von Wachter (2011). The sample is restricted to incumbent workers below age 50 with at least three years of tenure and pre-displacement earnings below the median. A worker is classified as displaced in year y if she receives an exogenous separation shock in year $y-1$. Recessions in the model are periods where the unemployment rate exceeds 8%.

Figure 13: Earnings Losses from Entering the Labor Market in a Recession: Model vs. Data



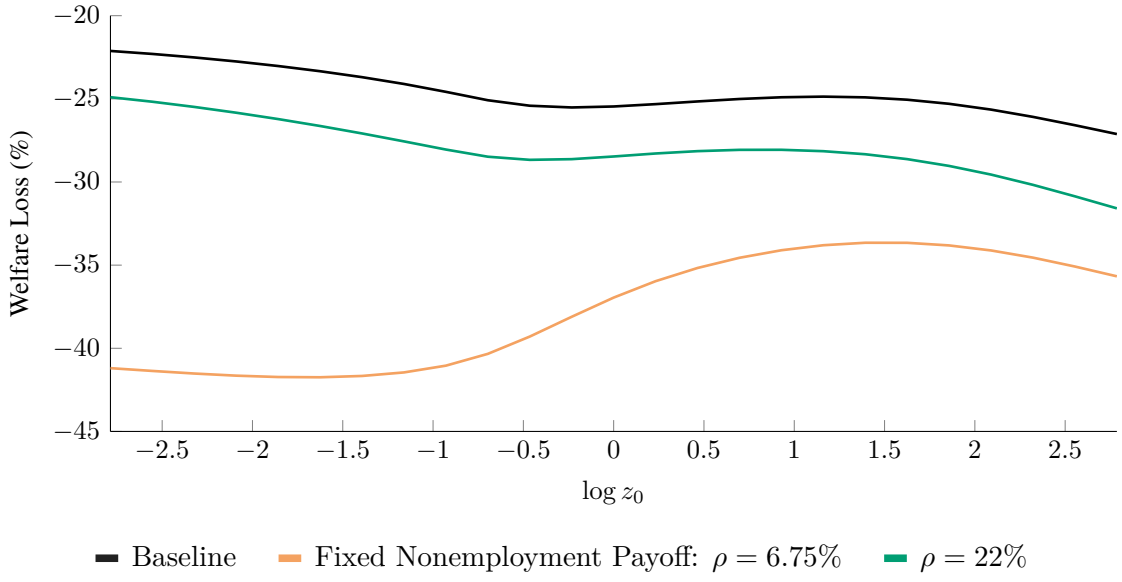
This figure plots average earnings by years since labor market entry for workers who enter in recessions versus expansions, comparing the model to the estimates of Schwandt and von Wachter (2019). Recessions in the model are periods where the unemployment rate exceeds 8%.

Figure 14: Losses Due to Job Displacement Across Workers



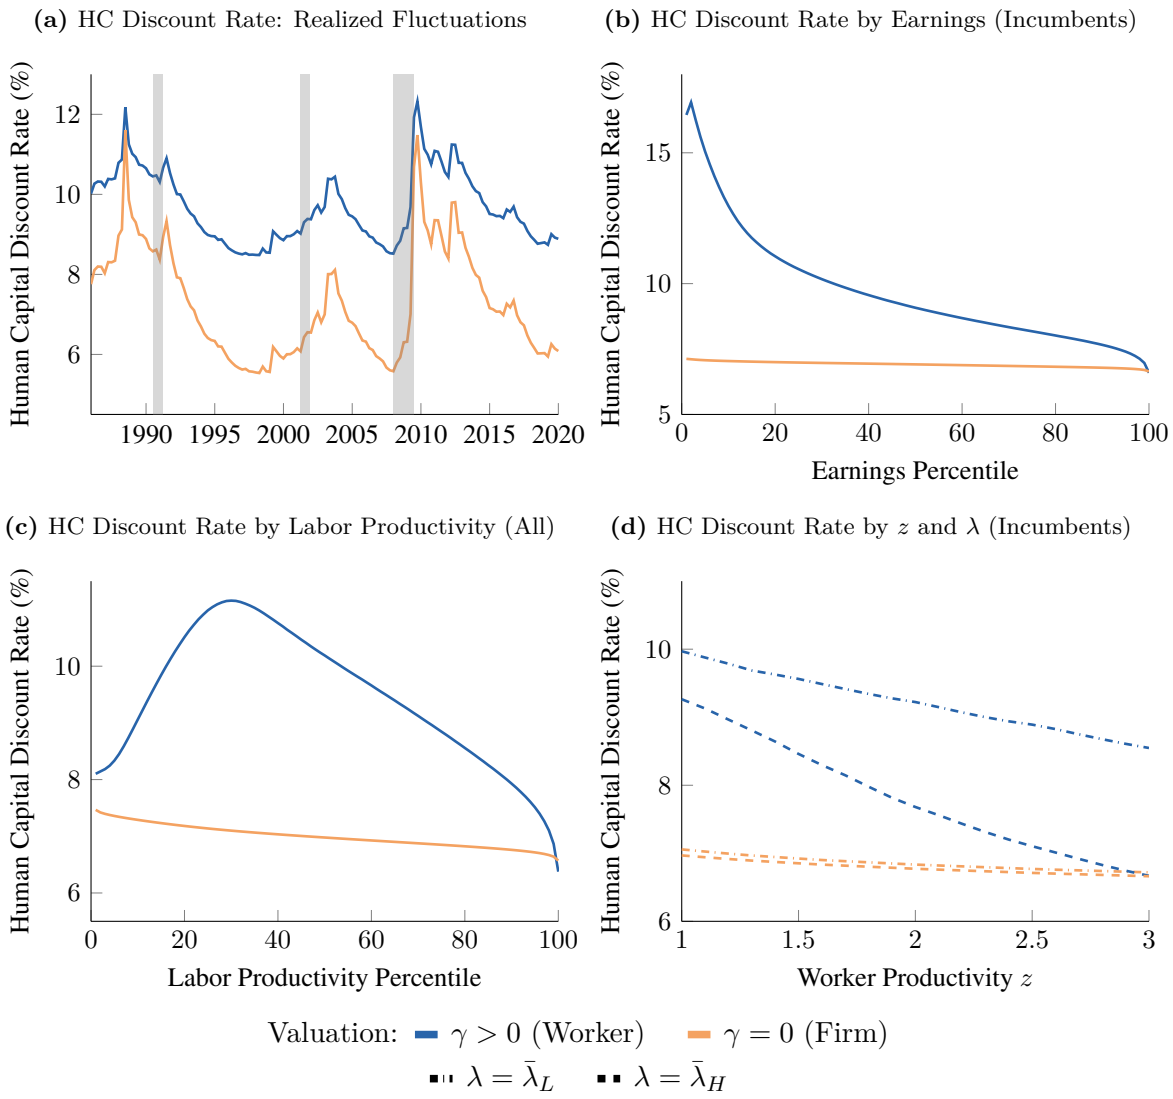
This figure plots present value losses due to an exogenous separation. The left panels (a) and (c) report the loss in the net present value of earnings, discounted with the aggregate SDF ($\gamma = 0$); the right panels (b) and (d) report the loss in worker welfare, measured in certainty-equivalent consumption ($\gamma > 0$). Panels (a) and (b) plot losses by current productivity z, separately for inexperienced ($\bar{\lambda}_L$) and experienced ($\bar{\lambda}_H$) workers, at average risk premia ($x_t = \bar{x}$). Panels (c) and (d) plot losses across the earnings distribution and compare expansions and recessions, where recessions in the model are periods in which the unemployment rate exceeds 8%.

Figure 15: Welfare Loss Due to Idiosyncratic Risk



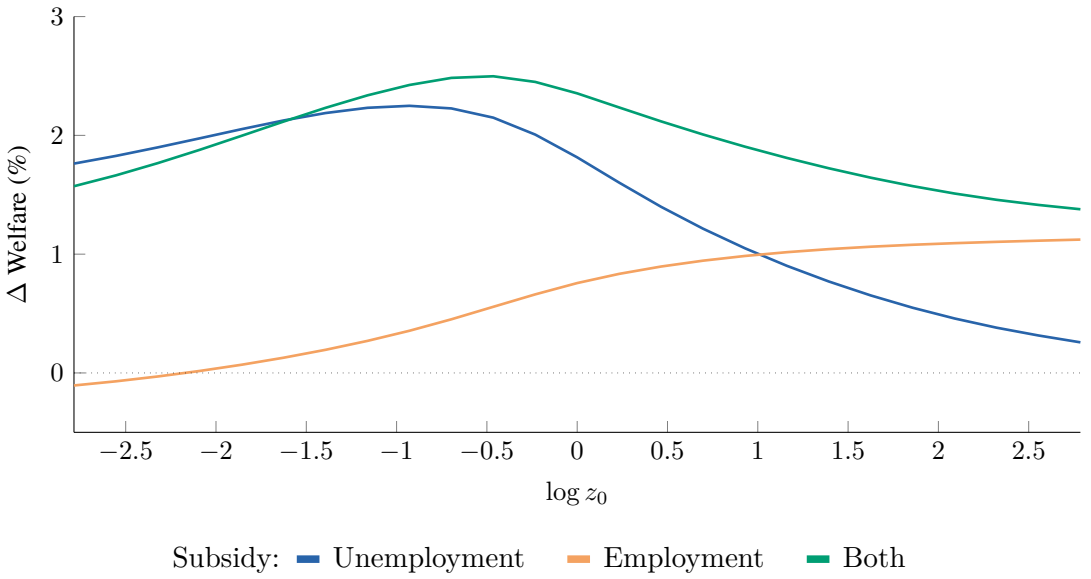
This figure plots welfare losses from idiosyncratic risk as a function of initial productivity z_0, measured as the difference between certainty-equivalent consumption under the baseline worker valuation ($\gamma > 0$) and under a counterfactual valuation with $\gamma = 0$. The baseline (black) uses the model-implied value of nonemployment. The colored lines hold employment and wage dynamics fixed but replace the model-implied nonemployment payoffs with a fixed payoff equal to a constant fraction ρ of the average wage, $b_t^{\text{fix}} = n_t^{\text{fix}} = \rho\, \overline{w/A} \times A_t$.

Figure 16: Human Capital Discount Rates in Model



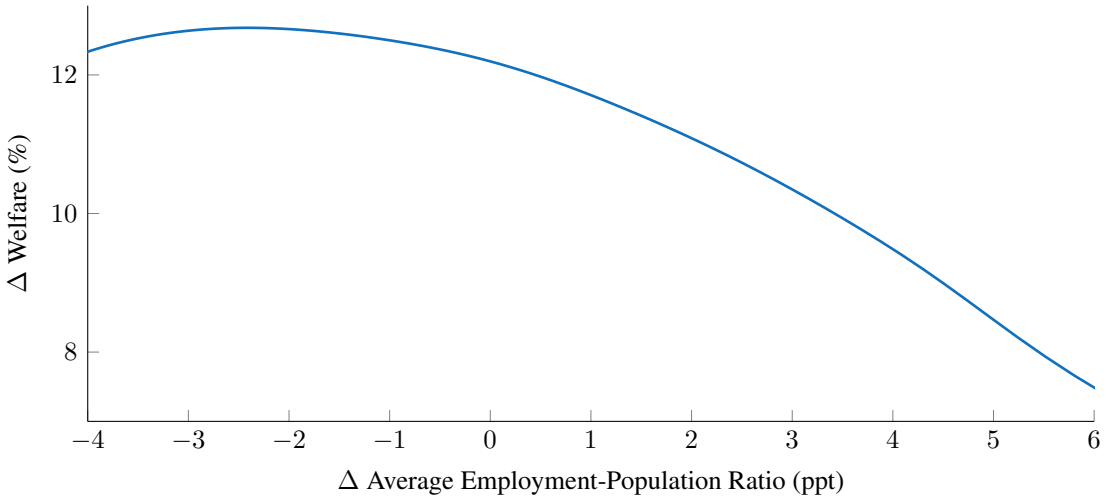
This figure plots the human capital discount rate under the worker ($\gamma > 0$) and firm ($\gamma = 0$) valuations, defined as the internal rate of return: for each worker, we compute expected payoffs by horizon, then find the per-period discount rate such that the resulting net present value equals the value of human capital. Panel (a) plots the time series of the aggregate annual rate; panels (b)–(d) plot average annual rates for incumbent workers sorted by prior earnings; for all workers sorted by labor productivity; and for incumbent workers sorted by productivity z and experience λ, respectively.

Figure 17: State-Contingent Labor Market Policies: Welfare Gains by Initial Productivity



This figure plots welfare gains from state-contingent labor market subsidies as a function of initial productivity z_0, for a recession threshold of $\delta = 0.1$. The unemployment subsidy raises the flow payoff in unemployment by 20%; the employment subsidy pays firms 6% of output; the combined policy implements both. Welfare gains are measured as the percentage change in certainty-equivalent consumption for newborn workers (integrated over aggregate states) relative to the no-policy baseline.

Figure 18: State-Contingent Labor Market Policies: Welfare–Employment Frontier



This figure plots the welfare–employment frontier of state-contingent labor market policies for a recession threshold of $\delta = 0.1$. Each point corresponds to the policy (ξ_b, ξ_y) that maximizes ex-ante welfare subject to a 10% cap on the financing consumption tax τ and a fixed value for the average employment-to-population ratio. The change in the average employment-to-population ratio and the welfare gain (certainty-equivalent consumption for newborn workers) are relative to the no-policy baseline.

Table 1: Heterogeneity in Pass-Through of Firm TFP Growth

	(1)	(2)	(3)	(4)	(5)	(6)
ϵ^{tfp}						
× Worker Earnings, 0–25th Percentile	4.21	3.79	4.29	3.97	6.32	3.53
	(5.01)	(5.29)	(31.15)	(28.33)	(22.83)	(23.33)
× Worker Earnings, 25–50th Percentile	4.41	4.14	4.42	4.21	6.59	3.80
	(6.23)	(6.50)	(39.48)	(37.16)	(29.73)	(30.93)
× Worker Earnings, 50–75th Percentile	4.81	4.60	4.87	4.71	6.78	4.33
	(7.20)	(7.61)	(45.46)	(43.27)	(31.31)	(36.70)
× Worker Earnings, 75–95th Percentile	6.69	6.54	6.81	6.69	9.26	6.19
	(8.66)	(9.15)	(54.29)	(52.43)	(36.33)	(44.82)
× Worker Earnings, 95–100th Percentile	13.74	13.80	13.93	13.97	17.48	13.32
	(10.03)	(10.61)	(44.45)	(43.72)	(26.53)	(38.45)
$\epsilon^{tfp} \times \epsilon^{rp}$						
× Worker Earnings, 0–25th Percentile		14.33		11.52	5.76	5.61
		(3.30)		(11.46)	(4.93)	(4.42)
× Worker Earnings, 25–50th Percentile		9.44		7.52	1.67	1.98
		(2.44)		(9.28)	(1.76)	(1.90)
× Worker Earnings, 50–75th Percentile		7.46		5.92	0.82	0.81
		(1.68)		(7.80)	(0.92)	(0.83)
× Worker Earnings, 75–95th Percentile		5.23		4.25	-2.05	-2.59
		(0.94)		(4.86)	(-1.99)	(-2.30)
× Worker Earnings, 95–100th Percentile		-2.29		-1.44	-10.08	-10.27
		(-0.31)		(-0.67)	(-3.94)	(-3.74)
Controls:						
Earn Grp × ϵ^{rp}	✓	✓				
Earn Grp × ϵ^{tfp} × GDP Growth					✓	
Earn Grp × ϵ^{tfp} × NBER Recession						✓
Fixed Effects:						
NAICS2 × Age × Gender	✓	✓	✓	✓	✓	✓
NAICS2 × Earn Grp	✓	✓				
NAICS2 × Earn Grp × Year			✓	✓	✓	✓
Observations	23.1m	23.1m	23.1m	23.1m	23.1m	23.1m

This table reports estimated coefficients from regression (33) at a three-year horizon. The outcome variable is cumulative age-adjusted earnings growth $g_{i,t:t+3}$. The upper panel reports the pass-through $\beta_{0,s}$ of firm TFP growth by within-firm prior earnings group; the lower panel reports the interaction $\beta_{1,s}$ with risk premium shocks. Columns (1)–(2) include industry × earnings group fixed effects; columns (3)–(6) replace these with industry × earnings group × year fixed effects. Columns (5) and (6) additionally control for the interaction of firm TFP growth with GDP growth and the NBER recession indicator, respectively. t-statistics based on standard errors clustered by worker and year are in parentheses. The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019.

Table 2: Pass-Through of Firm TFP Growth across Horizons

	Cumulative Earnings Growth			Probability of Nonemployment		
	1 Year	2 Years	5 Years	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)	(5)	(6)
ϵ^{tfp}						
× Worker Earnings, 0–25th Percentile	2.21	3.48	3.98	-0.93	-2.45	-3.14
	(4.82)	(6.18)	(4.71)	(-3.23)	(-5.33)	(-5.36)
× Worker Earnings, 25–50th Percentile	2.67	3.84	4.41	-0.81	-2.04	-2.80
	(6.43)	(7.67)	(5.56)	(-3.73)	(-4.79)	(-4.90)
× Worker Earnings, 50–75th Percentile	2.85	4.25	4.89	-0.48	-1.84	-2.55
	(7.17)	(8.86)	(6.41)	(-2.69)	(-4.87)	(-4.82)
× Worker Earnings, 75–95th Percentile	4.19	6.08	6.86	-0.40	-1.60	-2.37
	(9.39)	(10.55)	(7.80)	(-2.95)	(-5.62)	(-5.80)
× Worker Earnings, 95–100th Percentile	9.22	13.04	14.37	0.01	-1.14	-2.04
	(9.44)	(11.23)	(9.46)	(0.04)	(-3.74)	(-4.43)
$\epsilon^{tfp} \times \epsilon^{rp}$						
× Worker Earnings, 0–25th Percentile	19.26	17.84	9.87	-11.42	-11.84	-6.55
	(6.17)	(5.59)	(1.95)	(-6.50)	(-4.15)	(-2.45)
× Worker Earnings, 25–50th Percentile	14.54	12.68	5.79	-6.57	-7.29	-2.90
	(5.32)	(4.34)	(1.26)	(-5.60)	(-2.72)	(-1.03)
× Worker Earnings, 50–75th Percentile	12.30	10.29	3.69	-4.78	-4.86	-1.98
	(4.64)	(3.21)	(0.68)	(-4.43)	(-2.00)	(-0.66)
× Worker Earnings, 75–95th Percentile	9.88	7.66	1.77	-2.22	-1.52	1.55
	(4.17)	(1.91)	(0.26)	(-3.52)	(-0.82)	(0.64)
× Worker Earnings, 95–100th Percentile	3.94	1.87	-8.47	0.39	1.72	4.26
	(0.88)	(0.31)	(-1.10)	(0.64)	(1.19)	(1.96)
Controls:						
Earn Grp × ϵ^{rp}	✓	✓	✓	✓	✓	✓
Fixed Effects:						
NAICS2 × Age × Gender	✓	✓	✓	✓	✓	✓
NAICS2 × Earn Grp	✓	✓	✓	✓	✓	✓
Observations	24.9m	24.0m	21.3m	24.9m	24.0m	23.1m

This table reports estimated coefficients from regression (33) across horizons. Columns (1)–(3) use cumulative age-adjusted earnings growth $g_{i,t:t+H}$ as the outcome; columns (4)–(6) use an indicator for at least one zero-earnings quarter over the next H years. The upper panel reports the pass-through $\beta_{0,s}$ of firm TFP growth by within-firm prior earnings group; the lower panel reports the interaction $\beta_{1,s}$ with risk premium shocks. t-statistics based on standard errors clustered by worker and year are in parentheses. The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019.

Table 3: Pass-Through of Firm TFP Growth across Horizons: Stayers versus Movers

	1 Year		3 Years	
	Stayer	Mover	Stayer	Mover
	(1)	(2)	(3)	(4)
ϵ^{tfp}				
× Worker Earnings, 0–25th Percentile	1.95	-2.23	2.75	0.16
	(5.46)	(-2.46)	(5.58)	(0.20)
× Worker Earnings, 25–50th Percentile	2.54	-0.97	3.27	0.94
	(7.39)	(-1.15)	(8.24)	(1.12)
× Worker Earnings, 50–75th Percentile	2.99	-0.78	4.11	1.65
	(9.42)	(-0.82)	(10.94)	(1.95)
× Worker Earnings, 75–95th Percentile	4.58	0.19	6.69	3.07
	(12.56)	(0.20)	(12.55)	(3.31)
× Worker Earnings, 95–100th Percentile	10.74	1.05	16.17	7.52
	(11.99)	(0.77)	(12.74)	(5.99)
$\epsilon^{tfp} \times \epsilon^{rp}$				
× Worker Earnings, 0–25th Percentile	11.44	23.03	4.64	27.23
	(5.62)	(4.30)	(2.36)	(3.25)
× Worker Earnings, 25–50th Percentile	10.04	24.28	3.99	24.91
	(5.41)	(3.60)	(2.04)	(4.09)
× Worker Earnings, 50–75th Percentile	8.96	22.72	4.87	20.13
	(4.65)	(3.40)	(2.12)	(2.78)
× Worker Earnings, 75–95th Percentile	9.37	16.19	5.67	14.16
	(4.88)	(3.26)	(1.44)	(1.71)
× Worker Earnings, 95–100th Percentile	3.39	12.44	-0.22	5.20
	(0.82)	(1.97)	(-0.03)	(0.56)
Controls:				
Earn Grp × ϵ^{rp}	✓	✓	✓	✓
Fixed Effects:				
NAICS2 × Age × Gender	✓	✓	✓	✓
NAICS2 × Earn Grp	✓	✓	✓	✓
Observations	19.9m	4.2m	13.3m	8.9m

This table reports estimated coefficients from regression (33) at horizons of one and three years, estimated separately for stayers (workers who remain with their initial employer) and movers (workers who leave). The outcome variable is cumulative age-adjusted earnings growth $g_{i,t:t+H}$. The upper panel reports the pass-through $\beta_{0,s}$ of firm TFP growth by within-firm prior earnings group; the lower panel reports the interaction $\beta_{1,s}$ with risk premium shocks. t-statistics based on standard errors clustered by worker and year are in parentheses. The sample consists of incumbent workers at Compustat firms in the LEHD, 1990–2019.

Table 4: Calibrated Parameters

A. Parameters Calibrated Ex Ante	Symbol	Value
Average TFP growth (%)	μ_A	0.18
Volatility of TFP growth (%)	σ_A	0.52
Mortality rate (%)	ζ	0.28
Matching function elasticity	α	0.41
Initial value of h	$\bar{h}$	1
Growth of h in employment (%)	g_E	0.29
Persistence of z	ψ_z	0.991
Long-run mean of z	$\bar{z}$	1
Volatility of z (%)	σ_z	10.9
Volatility of initial z (%)	σ_{z0}	66.6
Low value of λ	$\bar{\lambda}_L$	1
Transition probability of λ (%)	f	2.8
B. Parameters Calibrated to Asset and Labor Markets	**Symbol**	**Value**
Time preference parameter	β	0.999
Average dividend growth (%)	μ_E	0.08
Persistence of price of risk	ψ_x	0.993
Average price of risk	$\bar{x}$	0.079
Volatility of price of risk (%)	σ_x	10.7
Vacancy posting cost, scale	$\bar{\kappa}_0$	0.016
Vacancy posting cost, elasticity to z	$\bar{\kappa}_1$	2.7
Exogenous separation rate (%)	s	0.75
Unemployment benefit, intercept	$\bar{b}_0$	1.8
Unemployment benefit, dependence on z	$\bar{b}_1$	0.77
Nonparticipation benefit	$\bar{n}$	2.3
Growth of h in nonemployment (%)	g_O	0.03
Volatility of h (%)	σ_h	2.7
High value of λ	$\bar{\lambda}_H$	2.4
Worker risk aversion	γ	0.48
On-the-job search, intercept	$\bar{\chi}_0$	44.4
On-the-job search, dependence on z	$\bar{\chi}_1$	2.7
Correlation of h in firm (%)	ρ_h	27.3
Correlation of z in firm (%)	ρ_z	25.0
Volatility of orthogonal firm TFP shocks (%)	σ_k	5.9

This table reports the parameter values in our baseline calibration of the model. The model is calibrated at a monthly frequency. See Section 3.4 for details.

Table 5: Labor Market Dynamics: Model vs. Data

	Volatility		Autocorrelation		Cyclicality	
A. *Labor Market Indicators*	Model	Data	Model	Data	Model	Data
Unemployment rate (%)	1.43	1.44	0.94	0.97	1.00	1.00
Long-term unemployment share (%)	6.05	5.78	0.75	0.97	2.61	3.45
Employment-population ratio (%)	2.53	1.08	0.96	0.97	-1.59	-0.72
Labor force participation rate (%)	1.65	0.35	0.95	0.91	-0.85	-0.07
Labor market tightness (log V/U ratio, %)	23.89	37.71	0.86	0.97	-14.91	-25.32
B. *Job Flows*						
Job-finding rate (%)	4.07	2.93	0.80	0.92	-2.10	-1.91
Separation rate into unemployment (%)	0.16	0.17	0.47	0.83	0.05	0.10
C. *Decomposition of Unemployment Rate*						
Unemployment rate w/ constant separations (%)	0.82	0.79	0.95	0.97	0.55	0.51
Unemployment rate w/ constant job finding (%)	0.45	0.61	0.82	0.94	0.21	0.40

This table reports key labor market moments in the model and in the data. We report the volatility and persistence (autocorrelation) of these series, together with their cyclicality—the slope coefficient (beta) of a regression of each series on the unemployment rate. Panel C reports the moments of counterfactual unemployment rate series that hold either the separation rate or the job-finding rate constant.

Table 6: Pre-Specified Labor Market Policies

	Labor Market Subsidy					
	Unemployment		Employment		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
Threshold parameter (δ, %)	10	20	10	20	10	20
Unemployment subsidy (ξ_b, %)	20	20	0	0	20	20
Employment subsidy (ξ_y, %)	0	0	6	6	6	6
Tax rate (τ, %)	1.2	1.8	1.0	1.6	2.0	3.1
Tax rate, no behavioral response (τ, %)	0.3	0.5	1.0	1.6	1.3	2.1
Δ welfare (%)	1.7	1.4	0.7	0.4	2.2	1.7
Δ welfare, no behavioral response (%)	1.1	1.1	0.6	0.2	1.7	1.3
Δ average GDP (%)	-2.3	-4.6	1.2	1.6	-1.0	-2.5
Δ average consumption (%)	-1.7	-2.7	-0.2	0.1	-1.7	-2.3
Δ employment-population ratio mean (ppt)	-1.8	-3.6	1.0	1.4	-0.7	-2.0
Δ employment-population ratio vol (ppt)	1.5	2.2	-0.4	-0.4	1.0	1.5
Δ unemployment rate mean (ppt)	3.5	7.6	-0.6	-0.8	2.7	6.3
Δ unemployment rate vol (ppt)	4.8	6.3	-0.4	-0.5	4.2	5.5
Δ income growth vol (ppt)	1.6	3.0	-0.6	-0.9	0.8	1.6

This table reports the effects of state-contingent labor market subsidies that are active when risk premia x_t exceed a threshold (recession frequency δ). The unemployment subsidy raises the flow payoff in unemployment by ξ_b; the employment subsidy pays firms a fraction ξ_y of output. Both are financed by a flat consumption tax τ such that the policy is budget-balanced in net present value (using the shareholders' SDF). Welfare is measured as certainty-equivalent consumption for newborn workers. The rows labeled "no behavioral response" recompute the tax and the welfare gain holding employment and wage dynamics fixed at their no-intervention values.

Table 7: Labor Market Policies on the Welfare–Employment Frontier

	(1)	(2)	(3)	(4)
Threshold parameter (δ, %)	10	10	10	10
Unemployment subsidy (ξ_b, %)	92	81	66	47
Employment subsidy (ξ_y, %)	18	34	44	52
Tax rate (τ, %)	10	10	10	10
Δ welfare (%)	12.7	12.2	11.1	9.5
Δ average GDP (%)	-4.6	-0.7	2.1	4.7
Δ average consumption (%)	-7.8	-6.3	-5.0	-3.7
Δ employment-population ratio mean (ppt)	-2.4	0.0	2.0	4.0
Δ employment-population ratio vol (ppt)	5.1	3.0	1.4	-0.2
Δ unemployment rate mean (ppt)	5.9	3.7	2.0	0.4
Δ unemployment rate vol (ppt)	8.8	6.6	4.7	2.9
Δ income growth vol (ppt)	4.5	2.0	-0.2	-2.3

This table reports four policies on the welfare–employment frontier and their effects, for $\delta = 0.1$ and subject to a 10% tax cap. Column (1) is the welfare-maximizing policy without an employment constraint; columns (2)–(4) fix average employment at the no-policy baseline, 2 percentage points above it, and 4 percentage points above it.

Online Appendix

A Model Appendix

A.1 A Microfoundation for Worker Preferences

The worker preferences in equation (6) arise from Epstein-Zin utility combined with time-varying ambiguity aversion over aggregate risk, in the spirit of Hansen and Sargent (2001) and Skiadas (2013a).

Consider an agent with Epstein-Zin preferences and certainty equivalent $\mathcal{R}_t$:

$$V_{i,t} = \left\{ c_{i,t}^{1-1/\phi} + \beta\,(1-\zeta)\,\left(\mathcal{R}_t(V_{i,t+1})\right)^{1-1/\phi} \right\}^{\frac{1}{1-1/\phi}}. \tag{A.1}$$

Under standard expected utility, $\mathcal{R}_t(V) = \mathbb{E}_t[V^{1-\gamma}]^{1/(1-\gamma)}$. We modify the certainty equivalent to incorporate ambiguity aversion over aggregate risk, while maintaining standard risk aversion γ over idiosyncratic risk. Under an approximation described below, this modification yields the worker preferences in equation (6).

Entropic penalty for ambiguity. The modified certainty equivalent evaluates $V_{i,t+1}^{1-\gamma}$ under a worst-case measure Q^* rather than the reference measure P:

$$\mathcal{R}_t(V_{i,t+1}) = \mathbb{E}_t^{Q^*}[V_{i,t+1}^{1-\gamma}]^{1/(1-\gamma)}. \tag{A.2}$$

Following Hansen and Sargent (2001), Q^* minimizes the agent's certainty-equivalent continuation utility, penalized by its Kullback–Leibler divergence from P:

$$Q^* = \arg\min_{Q\in\mathcal{Q}} \left\{ \frac{1}{1-\gamma}\,\mathbb{E}_t^Q\left[V_{i,t+1}^{1-\gamma}\right] + \Psi_{i,t}\,\mathbb{E}_t^Q\left[\log\frac{dQ}{dP}\right] \right\}. \tag{A.3}$$

The set of measures $\mathcal{Q}$ distort only the distribution of the aggregate shock $\varepsilon_{A,t+1}$, leaving idiosyncratic shocks unchanged. This restriction ensures that ambiguity aversion is *source-dependent* in the sense of Skiadas (2013a): the agent is ambiguity-averse toward aggregate risk but evaluates idiosyncratic risk under standard expected utility with risk aversion γ.

Following Maenhout (2004), we scale the penalty coefficient by the value function,

$$\Psi_{i,t} = V_{i,t}^{1-\gamma}\,\frac{\sigma_A}{x_t}, \tag{A.4}$$

to preserve homotheticity. The $\mathcal{F}_t$-measurable process $x_t > 0$ governs the time-varying degree of ambiguity aversion: higher x_t reduces the penalty for deviating from the reference measure, so the agent's worst-case distortion is larger.

Optimal distortion. The first-order condition of (A.3) yields an optimal distortion that depends on the agent's continuation value and hence on her idiosyncratic state:

$$\frac{dQ^*}{dP} \propto \exp\left\{-\frac{(V_{i,t+1}/V_{i,t})^{1-\gamma}}{(1-\gamma)\,\sigma_A}\,x_t\right\}. \tag{A.5}$$

The degree of pessimism would therefore be endogenous to the agent's state. We approximate by assuming that $(V_{i,t+1}/V_{i,t})^{1-\gamma}$ is a function of the aggregate shock alone and is given by the following log-linear expansion:

$$\left(\frac{V_{i,t+1}}{V_{i,t}}\right)^{1-\gamma} \approx k + (1-\gamma)\,\sigma_A\,\varepsilon_{A,t+1}. \tag{A.6}$$

Substituting into the first-order condition, the constant k is absorbed into the normalization $\mathbb{E}_t^P[dQ^*/dP] = 1$, and the optimal distortion reduces to

$$\frac{dQ^*}{dP} = \exp\left\{-\frac{1}{2}x_t^2 - x_t\,\varepsilon_{A,t+1}\right\} \equiv \Gamma_{t+1}. \tag{A.7}$$

Equation (A.7) is an approximation that would be exact if $\gamma = 1$ and every agent's consumption were proportional to A_t state-by-state. Workers are hand-to-mouth and their consumption equals their wage, which is not exactly proportional to A_t; however, all wages and flow payoffs are cointegrated with aggregate productivity, so the approximation is reasonable. This is in the spirit of Campbell and Cochrane (1999) and Kehoe et al. (2023), who formulate habit preferences over the exogenous endowment or productivity process rather than endogenous consumption.

Under the distorted measure Q^*, the aggregate shock is distributed as $\varepsilon_{A,t+1} \sim N(-x_t, 1)$—a pessimistic shift proportional to the degree of ambiguity aversion.

Properties of the distortion. This specification satisfies properties that Skiadas (2013a,b) show are necessary for scale-invariant, source-dependent ambiguity-averse preferences:

1. *Proper density:* $\mathbb{E}_t^P[\Gamma_{t+1}] = 1$, so Q^* is a valid probability measure.

2. *Source dependence:* Γ_{t+1} depends only on the aggregate shock $\varepsilon_{A,t+1}$ and $\mathcal{F}_t$-measurable parameters, not on idiosyncratic shocks.

3. *Time variation:* The degree of distortion is governed by x_t, which varies over time following the process in (8), so both the mean and conditional variance of the worst-case distortion vary with the aggregate state.

4. *Controlled divergence:* The Kullback–Leibler divergence of Q^* from P equals

$$KL(Q^*\|P) = \mathbb{E}_t^{Q^*}[\log\Gamma_{t+1}] = \mathbb{E}_t^P[\Gamma_{t+1}\log\Gamma_{t+1}] = \tfrac{1}{2}x_t^2, \tag{A.8}$$

   which is increasing in the degree of ambiguity aversion x_t.

From entropic penalty to multiplicative distortion. Since $\mathcal{R}_t$ is the CES certainty equivalent evaluated under Q^*, and Q^* has Radon-Nikodym derivative Γ_{t+1}, we have $\mathcal{R}_t(V_{i,t+1}) = \mathbb{E}_t[\Gamma_{t+1} V_{i,t+1}^{1-\gamma}]^{1/(1-\gamma)}$: the distortion enters multiplicatively inside the CES aggregator. Substituting into the Epstein-Zin recursion (A.1) yields

$$V_{i,t} = \left\{ c_{i,t}^{1-1/\phi} + \beta\,(1-\zeta)\,\mathbb{E}_t\left[\Gamma_{t+1}\, V_{i,t+1}^{1-\gamma}\right]^{\frac{1-1/\phi}{1-\gamma}} \right\}^{\frac{1}{1-1/\phi}}. \tag{A.9}$$

In the main text, we set $\phi = 1/\gamma$, under which $\frac{1-1/\phi}{1-\gamma} = 1$ and (A.9) simplifies to equation (6). This restriction ties the elasticity of intertemporal substitution to risk aversion; the Epstein-Zin formulation nests alternative calibrations where $\phi \neq 1/\gamma$.

A.2 Optimal Wage Contracting

At the final stage of the period, given current promised utility V^S, the firm chooses the optimal wage and state-contingent future promised utility to solve the following dynamic optimization problem:

$$J_t^{FS}(\Omega, V^S) = \max_{w, \{V^{C\prime}\}} \; y_t(\Omega) - w + (1-\zeta)(1-s)\,\mathbb{E}_{t,\Omega}\left[\Lambda_{t+1} \mathbb{1}_{t+1}^C(\Omega', V^{C\prime})\, J_{t+1}^{FC}(\Omega', V^{C\prime})\right] \tag{A.10}$$

$$\text{s.t.} \quad V^S \leq \Bigg\{ w^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{ V_{t+1}^O(\Omega')^{1-\gamma} \tag{A.11}$$

$$+ (1-s)\mathbb{1}_{t+1}^C(\Omega', V^{C\prime})\Big((V^{C\prime})^{1-\gamma} - V_{t+1}^O(\Omega')^{1-\gamma}\Big)\Big\}\Big]\Bigg\}^{\frac{1}{1-\gamma}}.$$

The indicator $\mathbb{1}_{t+1}^C(\Omega', V^{C\prime})$ captures limited commitment: the contract will be terminated by the worker if the worker's continuation value $V^{C\prime}$ is less than the outside option $V_{t+1}^O(\Omega')$ or if the firm's continuation value $J_{t+1}^{FC}(\Omega', V^{C\prime})$ is negative. Without loss of generality, we can impose that the firm always offers a contract that is not terminated by the worker, and the match will be terminated if and only if the firm's maximum continuation value when evaluated at the worker's outside option is negative. Thus, we can rewrite the firm's optimization problem as:

$$J_t^{FS}(\Omega, V^S) = \max_{w, \{V^{C\prime}\}} \; y_t(\Omega) - w + (1-\zeta)(1-s)\,\mathbb{E}_{t,\Omega}\left[\Lambda_{t+1}\, \mathbb{1}_{t+1}^C(\Omega')\, J_{t+1}^{FC}(\Omega', V^{C\prime})\right] \tag{A.12}$$

$$\text{s.t.} \quad V^S = \Bigg\{ w^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{ V_{t+1}^O(\Omega')^{1-\gamma} \tag{A.13}$$

$$+ (1-s)\mathbb{1}_{t+1}^C(\Omega')\Big((V^{C\prime})^{1-\gamma} - V_{t+1}^O(\Omega')^{1-\gamma}\Big)\Big\}\Big]\Bigg\}^{\frac{1}{1-\gamma}}$$

$$V^{C\prime} \geq V_{t+1}^O(\Omega') \tag{A.14}$$

$$\mathbb{1}_{t+1}^C(\Omega')\, J_{t+1}^{FC}(\Omega', V^{C\prime}) \geq 0, \tag{A.15}$$

where the continuation indicator $\mathbb{1}^C_{t+1}(\Omega')$ does not depend on promised utility and is given by

$$\mathbb{1}^C_{t+1}(\Omega') = \begin{cases} 1 & \text{if } J^{FC}_{t+1}(\Omega', V^O_{t+1}(\Omega')) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{A.16}$$

If the continuation indicator is zero, there is no feasible continuation value that satisfies the limited commitment constraint on both ends, so that any promised utility in this state will lead to a termination of the match. If there exists a feasiable continuation value, it is always optimal for the firm to offer a continuation value to the worker that is at least as high as her outside option.

First-order conditions and risk-sharing. Let ξ denote the Lagrange multiplier on the promise-keeping constraint. In future states where neither the firm's nor the worker's limited commitment constraint binds, the first-order conditions are

$$1 = \xi \, (V^S)^\gamma \, w^{-\gamma}, \tag{A.17}$$

$$-\Lambda_{t+1} \frac{\partial J^{FC}_{t+1}(\Omega', V^{C\prime})}{\partial V^{C\prime}} = \xi \, \beta \, \Gamma_{t+1} \, (V^S)^\gamma \, (V^{C\prime})^{-\gamma}, \tag{A.18}$$

for the wage w and state-contingent continuation value $V^{C\prime}$, respectively. The envelope condition gives

$$\frac{\partial J^{FS}_t(\Omega, V^S)}{\partial V^S} = -\xi. \tag{A.19}$$

Combining (A.17) and (A.19):

$$-\frac{\partial J^{FS}_t(\Omega, V^S)}{\partial V^S} = \left(\frac{V^S}{w}\right)^{-\gamma}. \tag{A.20}$$

Substituting $\xi = w^\gamma (V^S)^{-\gamma}$ from (A.17) into (A.18) and using $\Lambda_{t+1} = \beta \, \Gamma_{t+1}$:

$$-\frac{\partial J^{FC}_{t+1}(\Omega', V^{C\prime})}{\partial V^{C\prime}} = \left(\frac{V^{C\prime}}{w}\right)^{-\gamma}. \tag{A.21}$$

Together, (A.20) and (A.21) yield the risk-sharing condition (28): the ratio of worker marginal utility to the firm's marginal cost of providing value is equalized across all states where neither limited commitment constraint binds.

Wage dynamics without on-the-job search. Without OJS ($\chi = 0$), the continuation and post-search values coincide ($V^{C\prime} = V^{S\prime}$) and $J^{FC} = J^{FS}$. Combining (A.20) (evaluated at $t+1$) with (A.21) yields

$$\left(\frac{V^{C\prime}}{w}\right)^{-\gamma} = \left(\frac{V^{C\prime}}{w'}\right)^{-\gamma} \quad \Rightarrow \quad w' = w. \tag{A.22}$$

Wages are constant as long as neither constraint binds, as in Thomas and Worrall (1990). When the worker's limited commitment constraint (A.14) binds ($V^{C\prime} = V^O_{t+1}(\Omega')$), the FOC (A.18) holds

as an inequality and the same logic yields $w' > w$: wages rise to satisfy the worker's outside option. Symmetrically, when the firm's participation constraint (A.15) binds, $w' < w$.

Wage growth with on-the-job search. With OJS, J^{FC} and J^{FS} are linked through the randomization stage. Abstracting from the lottery, the relationships at $t+1$ are

$$J^{FC}_{t+1}(\Omega', V^{C\prime}) = \tilde{p}_{t+1}(\Omega', V^{S\prime})\, J^{FS}_{t+1}(\Omega', V^{S\prime}), \tag{A.23}$$

$$(V^{C\prime})^{1-\gamma} = (V^{S\prime})^{1-\gamma} + \chi(\Omega')\,(1-\gamma)\,R_{t+1}(\Omega', V^{S\prime}), \tag{A.24}$$

where (A.24) is the promise-keeping constraint (23) of the randomization stage. Differentiating (A.23) with respect to $V^{C\prime}$ via the chain rule:

$$\frac{\partial J^{FC}_{t+1}(\Omega', V^{C\prime})}{\partial V^{C\prime}} = \frac{\frac{\partial \tilde{p}_{t+1}(\Omega', V^{S\prime})}{\partial V^{S\prime}}\, J^{FS}_{t+1}(\Omega', V^{S\prime}) + \tilde{p}_{t+1}(\Omega', V^{S\prime})\, \frac{\partial J^{FS}_{t+1}(\Omega', V^{S\prime})}{\partial V^{S\prime}}}{(V^{C\prime})^{\gamma}\left[(V^{S\prime})^{-\gamma} + \chi(\Omega')\, \frac{\partial R_{t+1}(\Omega', V^{S\prime})}{\partial V^{S\prime}}\right]}. \tag{A.25}$$

The envelope theorem applied to the worker's search problem gives $\partial R_{t+1}(\Omega', V^{S\prime})/\partial V^{S\prime} = -p(\theta_{t+1}(\Omega', \mathcal{V}^*(\Omega', V^{S\prime})))\,(V^{S\prime})^{-\gamma}$. The bracketed term in the denominator of (A.25) therefore simplifies to $\tilde{p}_{t+1}(\Omega', V^{S\prime})\,(V^{S\prime})^{-\gamma}$, by the definition (17) of the retention probability.

Substituting (A.21) for the left-hand side and (A.20) (applied at $t+1$) for $\partial J^{FS}_{t+1}/\partial V^{S\prime}$, and rearranging,

$$(w')^{\gamma} = w^{\gamma} + (V^{S\prime})^{\gamma}\, \frac{\partial \log \tilde{p}_{t+1}(\Omega', V^{S\prime})}{\partial V^{S\prime}}\, J^{FS}_{t+1}(\Omega', V^{S\prime}), \tag{A.26}$$

which yields the wage growth equation (29) in the main text. The second term captures the retention motive: since $\partial \log \tilde{p}_{t+1}/\partial V^{S\prime} \geq 0$ and $J^{FS}_{t+1} \geq 0$, wages are weakly increasing (backloaded) even when neither limited commitment constraint binds.

Behavior at the bounds. When the worker's limited commitment constraint binds, the post-search value $V^{S\prime}$ is determined by inverting the promise-keeping constraint (A.24) at $V^{C\prime} = V^{O}_{t+1}(\Omega')$, and the wage w' follows from the envelope condition (A.20) evaluated at this $V^{S\prime}$. When the firm's participation constraint binds, the post-search value $V^{S\prime}$ is determined by inverting the promise-keeping constraint (A.24) at the value $V^{C\prime}$ that solves $J^{FC}_{t+1}(\Omega', V^{C\prime}) = 0$, and the wage w' again follows from the envelope condition (A.20).

A.3 Calibration of the Stochastic Discount Factor

We choose the parameters that govern the dynamics of the stochastic discount factor (SDF) to match asset pricing moments, as in Meeuwis et al. (2025). To do so, we make the common assumption that corporate earnings E_t represent a levered claim on aggregate productivity,

$$\Delta \log E_{t+1} = \mu_E + L\,\sigma_A\,\varepsilon_{A,t+1}, \tag{A.27}$$

where μ_E is expected earnings growth and L is the leverage parameter. Based on the average value of nonfinancial corporate business debt as a percentage of the market value of corporate equity

between 1952 and 2019 from the Flow of Funds, which is 49%, we assume a leverage parameter L equal to 1.49. The total value of the stock market is given by the present value of aggregate earnings as specified in (11).

To calibrate the price of risk process x_t, we follow a strategy similar to that of Lettau and Wachter (2007), with two important distinctions: first, we assume that x_t follows an AR(1) process in logs instead of levels, so that it cannot get negative; second, we assume that aggregate productivity shocks and risk premium shocks are perfectly negatively correlated so that risk premium shocks are also priced.

We simulate the model at a monthly frequency and aggregate all financial variables to an annual frequency to compute annual moments. We choose β, μ_E, $\bar{x}$, ψ_x, and σ_x to target the average risk-free rate, the average price–earnings ratio, the autocorrelation of the log price–earnings ratio, and the mean and volatility of aggregate stock market returns.

Table A.1 shows that our calibration ($\beta = 0.999$, $\mu_E = 0.08\%$, $\bar{x} = 0.079$, $\psi_x = 0.993$, $\sigma_x = 0.107$) closely matches the targeted moments. The maximum monthly Sharpe ratio that can be attained in financial markets is

$$\frac{\sqrt{\text{Var}_t[\Lambda_{t+1}]}}{\mathbb{E}_t[\Lambda_{t+1}]} = \sqrt{\exp\{x_t^2\} - 1} \approx x_t. \tag{A.28}$$

When x_t is at its long-run mean $\bar{x}$, the maximum monthly Sharpe ratio is 0.079, or approximately 0.27 per year.

We assume that our empirical measure of risk premium shocks ϵ_{t+1}^{rp} corresponds to the aggregate shock $-\varepsilon_{A,t+1}$ in the model, which jointly drives aggregate productivity and risk premia. In quantitative comparisons of the model with the data, we therefore assume that ϵ_{t+1}^{rp} is proportional to $\varepsilon_{A,t+1}$. Given that the empirical distribution of ϵ_{t+1}^{rp} is positively skewed and leptokurtic, we calibrate the proportionality coefficient such that the interpercentile range (p99–p1) of monthly risk premium shocks matches between the model and the data: $\epsilon_{t+1}^{rp} = -0.045 \times \varepsilon_{A,t+1}$. Under this assumption, the sample moments of model-implied quantities given the realized risk premium shock series are similar to the unconditional moments. Appendix Figure A.15 further shows that the resulting model-implied price–earnings ratio captures the major swings in Shiller's CAPE ratio over time. We maintain the timing assumption from Section 3.1 in linking financial shocks to labor market outcomes.

B Additional Details on the Empirical Analysis

Here, we provide further details on the data construction and empirical analysis.

B.1 Worker Earnings Data

Our main data are employer–employee linked data from the Longitudinal Employer–Household Dynamics (LEHD) database. The LEHD contains earnings and employer information for U.S. workers, collected from state unemployment insurance filings. The LEHD data start in 1990, although many states joined the sample in later years as coverage became more complete. By the mid- to late-1990s, the LEHD covers the majority of jobs. We use data for years until 2019; only a few states drop

out of the sample for years before then. The LEHD data are based on firms' unemployment insurance filings to the state and contain total gross wages and other taxable forms of compensation as a measure of earnings. For the state–quarters in the LEHD, coverage of private sector jobs is nearly 100%. We link worker earnings to demographic information such as age and gender and convert all nominal earnings measures to real figures by deflating with the consumer price index (CPI).

The data allow us to track the incomes of individual workers over time and across employers. Our sample in year t covers individuals between ages 25 and 60 who live in a state in year t that is in the LEHD between years $t-2$ and $t+5$ and who have labor earnings in years t, $t-1$, and $t-2$ that exceed a minimum annual threshold as in Guvenen et al. (2014): the federal minimum wage times 20 hours times 13 weeks (1885 dollars in 2019). We merge leads and lags of individual annual labor earnings to the base year, where individuals without any earnings are assigned zero wage earnings for that year.

In addition to total earnings, we separately observe earnings and employer identity for the top three jobs (by income) of an individual in that year. We use the Employer Identification Number (EIN) of the employer associated with the highest annual earnings for the individual to assign workers to firms. In selecting the sample for year t, we require individuals to have strictly positive earnings from this employer in year $t+1$ to make sure that the employment relationship is still active by the end of year t. For workers for whom we observe a complete earnings history between years $t-5$ and t, we construct indicators for employment tenure by counting the number of consecutive years that the worker has received income from the current main employer.

A key focus of our analysis is on heterogeneity in the effects of risk premium and productivity shocks across the income distribution. We rank workers by their prior earnings relative to their peers. In particular, we sort workers by their last three years of total age-adjusted wage earnings, $w_{i,t-2,t}$, and compute the income rank of workers within their own firm. To compute these earnings ranks, we require observing at least 50 workers in the sample for a firm–year. We focus on quartiles of the initial earnings distribution, where we further separate out the top 5% from the remainder of the top quartile.

We use an internal Census table for mapping EIN to GVKEY identifiers to link firm information from Compustat to the worker earnings data. For most of our analysis, we focus on employees of publicly traded companies, for whom we have better measures of risk premium exposures and productivity shocks. We build our sample by first collecting data for all U.S. workers in the LEHD who are linked to Compustat firms in the base year t and constructing the yearly income ranks for this full sample. We exclude workers employed by firms with missing industry codes or who work in the utilities sector (NAICS codes starting with 22) or financial sector (NAICS codes starting with 52 or 53) from the sample.

After constructing all relevant variables, we draw a stratified random subsample of workers within each firm–year to keep the analysis computationally feasible. Let $N_{f,t}$ denote the total number of incumbent workers in firm f in year t. Workers at or below the 95th percentile of the within-firm earnings distribution are sampled with probability $p_{i,t} = \min\big(500 \,/\, (0.95 \cdot N_{f,t}),\, 1\big)$, and workers above the 95th percentile with probability $p_{i,t} = \min\big(250 \,/\, (0.05 \cdot N_{f,t}),\, 1\big)$. For large firms, this selects on average 500 workers from the bottom 95% and 250 from the top 5%; for smaller

firms, all workers are included. The top 5% is oversampled relative to its population share to ensure adequate coverage of the top earnings bin used in our analysis.

In all pass-through regressions, observations are weighted by $w_{i,t} = 1/(N_{f,t} \cdot p_{i,t})$, where $p_{i,t}$ is the worker's sampling probability. The inverse sampling probability corrects for the stratified design, and the additional scaling by $1/N_{f,t}$ weights firms equally regardless of size, as e.g. in Green et al. (2025).

An additional benefit of the LEHD is that it contains total earnings for each quarter in addition to the annual information. We use this information to construct a nonemployment indicator that takes the value of one if an individual has a quarter of zero earnings over a particular period. We also use worker earnings data split out per employer in future years to classify workers as stayers versus movers with respect to their initial job.

B.2 Productivity Shocks

We use the approach from İmrohoroğlu and Tüzel (2014) to estimate a revenue-based measure of total factor productivity (TFP) growth at the firm level based on the production function

$$y_{jt} = \beta_0 + \beta_k k_{jt} + \beta_l l_{jt} + \omega_{jt} + \eta_{jt}, \tag{A.29}$$

where y_{jt} is the log of value added for firm j in year t, k_{jt} and l_{jt} are log capital and labor, respectively, ω_{jt} is log firm TFP, and η_{jt} is an error term. We estimate the parameters β_k and β_l by implementing the semiparametric methodology of Olley and Pakes (1996). From these estimates, we then compute firm-level TFP growth as

$$\Delta\omega_{jt} = \Delta y_{jt} - \hat{\beta}_k \Delta k_{jt} - \hat{\beta}_l \Delta l_{jt}. \tag{A.30}$$

In their estimation of β_k and β_l, İmrohoroğlu and Tüzel (2014) use industry–time fixed effects to separate firm productivity from industry or aggregate effects. To obtain estimates of firm-level TFP growth that are suitable for aggregation, we re-estimate firm TFP growth based on their methodology but replace the industry–year fixed effects with industry fixed effects at the 3-digit SIC level.

We apply this methodology using data from Compustat, complemented by output and investment deflators from the Bureau of Economic Analysis and wage data from the Social Security Administration. We estimate the production function parameters for every year between 1964 and 2020 using all data up until that year to avoid using any forward-looking information. We winsorize the resulting firm-level growth series at the 1% and 99% levels.

We use this series rather than the TFP series from the Bureau of Labor Statistics (BLS) for several reasons. First, the İmrohoroğlu and Tüzel (2014) series is a direct estimate of revenue-based total factor productivity (TFPR) at the firm level, which Guiso et al. (2005) show has some pass-through to worker wages. By contrast, the TFP series from the BLS are defined as the difference between real output and a shares-weighted combination of factor inputs at the sector or industry level. Second, the BLS series are available only at a granular level for manufacturing industries. Third, for some industries, there are some salient differences between private and public firms; our

analysis is based on public firms, and the İmrohoroğlu and Tüzel (2014) measure of productivity directly applies to these firms.

B.3 Risk Premium Shocks

To measure variation in risk premia, we follow Meeuwis et al. (2025) and construct an index of risk premium shocks that captures fluctuations in either the level of risk or investors' risk-bearing capacity. We draw on nine existing monthly series from the literature: the excess bond premium of Gilchrist and Zakrajšek (2012); Shiller's CAPE ratio; the Chicago Fed's National Financial Conditions Index (NFCI); the financial uncertainty index of Jurado et al. (2015); the risk appetite index of Bauer et al. (2023); the risk aversion index of Bekaert, Engstrom, and Xu (2022); the variance risk premium of Bekaert and Hoerova (2014); the CBOE VIX; and the SVIX of Martin (2016). All series are signed so that an increase indicates elevated risk premia; innovations to all series are negatively correlated with stock market returns in the same month. Since the majority of the series are available from the 1980s and we link them to worker data starting from 1990, we collect data from December 1984.

Because each series is a noisy proxy, we focus on their common component. For each, we measure innovations by estimating AR(1) residuals. We then extract the first principal component across these residuals, following Bauer et al. (2023) in dealing with missing observations. We denote the resulting series of risk premium shocks as ϵ^{rp}. This common factor explains 60% of total variance, with correlations with individual series residuals ranging from 51% to 75%. Since earnings are annual, we construct annual risk premium shocks ϵ^{rp}_{t+1} by cumulating monthly shocks from mid-year t to mid-year $t+1$.

B.4 Pass-Through Regression Specification

In the pass-through regressions (31) and (33), the outcome variable $g_{i,t:t+H}$ is cumulative, age-adjusted earnings growth defined in equation (30), $f(i,t)$ denotes the employer of worker i at time t, and $\mathbb{1}_{\tau(i,t)=s}$ is an indicator for the worker's prior earnings bin $\tau(i,t)$, defined as the worker's rank by prior earnings relative to other workers within the same firm.

The controls $\mathbf{Z}_{i,t}$ include: a third-order polynomial in the log of average earnings over the past three years; the lagged risk premium index interacted with labor income group dummies; fixed effects for the worker's industry (defined at the 2-digit NAICS level) interacted with her labor income bin; and worker industry $\times$ age $\times$ gender fixed effects. Standard errors are clustered by worker and year. In the time-variation specification (33), we also consider alternative specifications that interact the firm-specific shocks by earnings group with aggregate output growth or the fraction of the year spent in NBER recessions.

B.5 CPS Data on Worker Flows

We measure gross flows between worker employment states using microdata from the Current Population Survey (CPS) between January 1978 and December 2019. The flows are calculated by making use of the rotating-panel sampling procedure, where households are included in the sample

for four months, rotated out for eight months, and then rotated back in for another four months. We follow the algorithm of Elsby et al. (2015); Krusell et al. (2017) in estimating worker flows for all respondents and the associated monthly transition flow probabilities between employment, unemployment, and nonparticipation.

It is well known that survey-based measures of gross flows between recorded employment states are sensitive to classification errors, especially between the states of unemployment and nonparticipation. We implement the Abowd-Zellner correction for classification errors that adjusts transition probabilities for the estimates of misclassification probabilities from Abowd and Zellner (1985), which are based on resolved labor force status from follow-up CPS interviews. The literature has found that all labor market states become more persistent after correction than what is implied by the unadjusted flows. Following the prior literature, we also implement a margin-error adjustment that restricts the estimates of worker flows to be consistent with the published aggregate labor market stocks of workers in employment, unemployment, and nonparticipation.

B.6 SIPP Data on Worker Flows

Given our focus on heterogeneity in labor market dynamics across workers with different income levels, we also want to measure worker flows conditional on wage earnings in the data. Since it is not possible to compute a time series of transition rates by income in the CPS, we turn to data from the Survey of Income and Program Participation (SIPP) of the U.S. Census Bureau to assess the relation between gross worker flows and earnings.

The SIPP is a longitudinal national household survey where participants are repeatedly interviewed on their labor market participation, income, demographic characteristics, and other economically relevant dynamics over a multiyear period. The SIPP consists of multiple panels that each last for several years. The SIPP had major redesigns in 1996 and 2014. Respondents are interviewed every four months (before 2014) or year (from 2014) about monthly outcomes over the past months.

We use data from the 1990–2019 panels of the SIPP, which cover the period from November 1989 to December 2019 with some gaps. We measure monthly employment status from reports in the last week of each month. Analogous to the CPS, we classify individuals as employed if they have a job and are working, absent without pay, or on paid leave. Individuals are classified as unemployed if they have no job and are either looking for work or on layoff. We also track workers who are not participating in the labor market.

In our calibration, we separately target the dynamics of separation and job-finding rates by worker earnings levels. For separation rates, we restrict attention to incumbent workers with positive wage earnings who report having a job in all weeks of the initial month. We sort these employed workers into income groups based on their wage earnings in the current month and compute the share of workers that become unemployed in the next month by earnings quartile bin. For job-finding rates, we sort unemployed workers into income groups based on their last reported (full-month) monthly wage income during the prior 12 months, if any. We then compute the share of workers that report having a job in the next month by prior earnings quartile bin.

It is well established that there is a significant level difference in flow rates computed using the CPS versus the SIPP (Fujita, Nekarda, and Ramey, 2007). Since we calibrate the model to conventional moments of aggregate flows based on the CPS, we adjust the flow rates from the SIPP by removing the level effect. Specifically, we scale the monthly transition probabilities for each earnings group by the respective unconditional average flow rate. That is, we only use the SIPP to estimate relative differences in flows across the earnings distribution.

B.7 Cyclical Dynamics

Both in the data and in the model, we average all monthly labor market stocks and flows at the quarterly frequency. Following Shimer (2005), we apply a low-frequency HP filter with smoothing parameter 10^5 to these series to capture business-cycle fluctuations.

References

Abowd, J. M. and A. Zellner (1985). Estimating gross labor-force flows. *Journal of Business & Economic Statistics 3*(3), 254–283.

Bauer, M. D., B. S. Bernanke, and E. Milstein (2023). Risk appetite and the risk-taking channel of monetary policy. *The Journal of Economic Perspectives 37*(1), pp. 77–100.

Bekaert, G., E. C. Engstrom, and N. R. Xu (2022). The time variation in risk appetite and uncertainty. *Management Science 68*(6), 3975–4004.

Bekaert, G. and M. Hoerova (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics 183*(2), 181–192. Analysis of Financial Data.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Elsby, M. W., B. Hobijn, and A. Şahin (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics 72*, 64–82.

Fujita, S., C. J. Nekarda, and G. Ramey (2007). The cyclicality of worker flows: New evidence from the sipp. Working paper.

Gilchrist, S. and E. Zakrajšek (2012, June). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

Green, B., L. Kogan, D. Papanikolaou, and L. D. W. Schmidt (2025). Winners and losers: Competition, creative destruction, and labor income risk. Working paper, MIT Sloan School of Management.

Guiso, L., L. Pistaferri, and F. Schivardi (2005). Insurance within the firm. *Journal of Political Economy 113*(5), 1054–1087.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Hansen, L. and T. J. Sargent (2001, May). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

İmrohoroğlu, A. and Ş. Tüzel (2014). Firm-level productivity, risk, and return. *Management Science 60*(8), 2073–2090.

Jurado, K., S. C. Ludvigson, and S. Ng (2015, March). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216.

Kehoe, P. J., P. Lopez, V. Midrigan, and E. Pastorino (2023). Asset Prices and Unemployment Fluctuations: A Resolution of the Unemployment Volatility Puzzle. *Review of Economic Studies 90*(3), 1304–1357.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin (2017, November). Gross worker flows over the business cycle. *American Economic Review 107*(11), 3447–76.

Lettau, M. and J. A. Wachter (2007). Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium. *The Journal of Finance 62*(1), 55–92.

Maenhout, P. J. (2004, 04). Robust portfolio rules and asset pricing. *The Review of Financial Studies 17*(4), 951–983.

Martin, I. (2016, 10). What is the Expected Return on the Market? *The Quarterly Journal of Economics 132*(1), 367–433.

Meeuwis, M., D. Papanikolaou, J. L. Rothbaum, and L. D. Schmidt (2025). Time-varying risk premia and heterogeneous labor market dynamics. Working paper.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*(6), 1263–1297.

Shimer, R. (2005, March). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review 95*(1), 25–49.

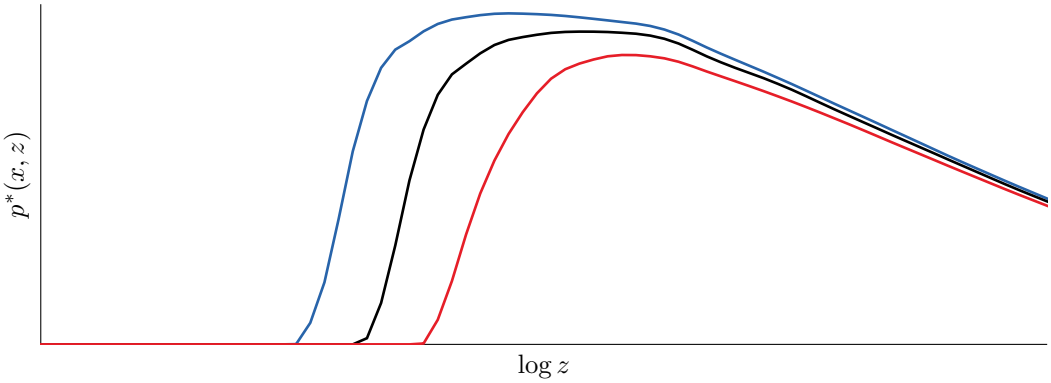
Skiadas, C. (2013a). Scale-invariant uncertainty-averse preferences and source-dependent constant relative risk aversion. *Theoretical Economics 8*(1), 59–93.

Skiadas, C. (2013b). Smooth ambiguity aversion toward small risks and continuous-time recursive utility. *Journal of Political Economy 121*(4), 775–792.

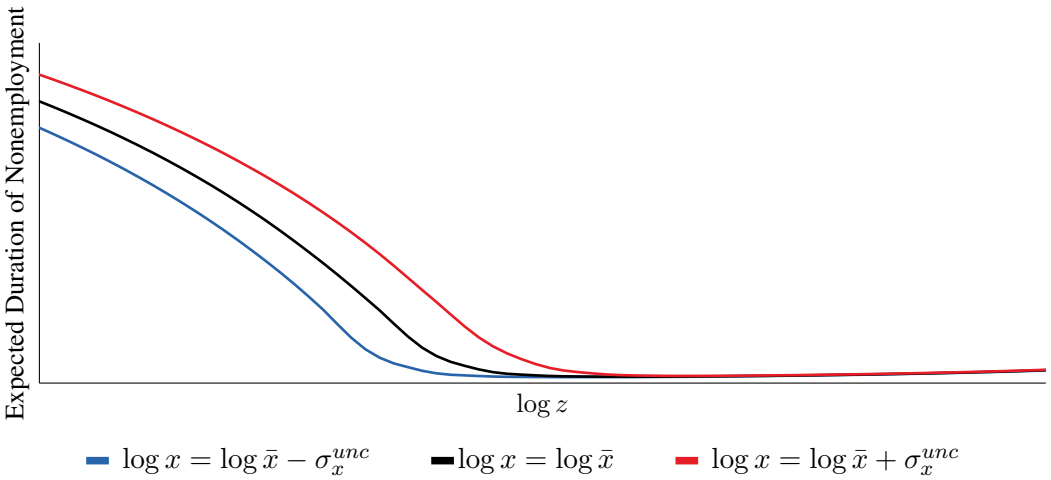
Thomas, J. and T. Worrall (1990, August). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory 51*(2), 367–390.

Figure A.1: Job-Finding Rate and Nonemployment Duration

(a) Probability of Finding a Job for Unemployed Worker: $p^*(x, z)$



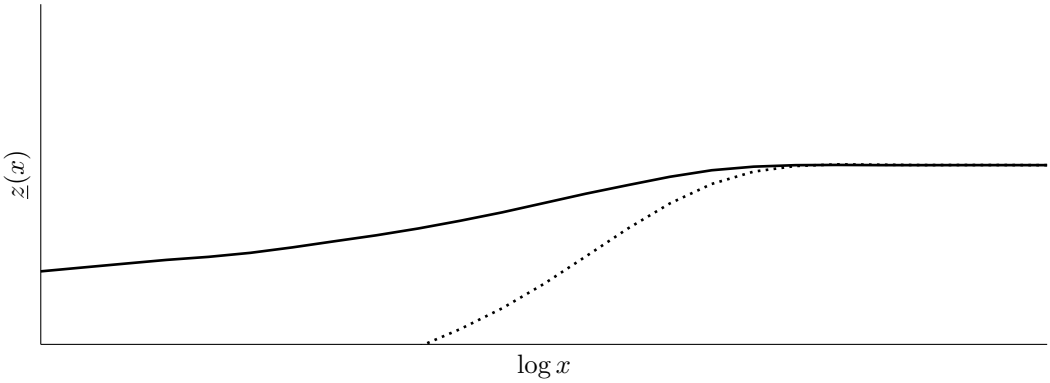
(b) Expected Duration of Nonemployment for Workers Without Job



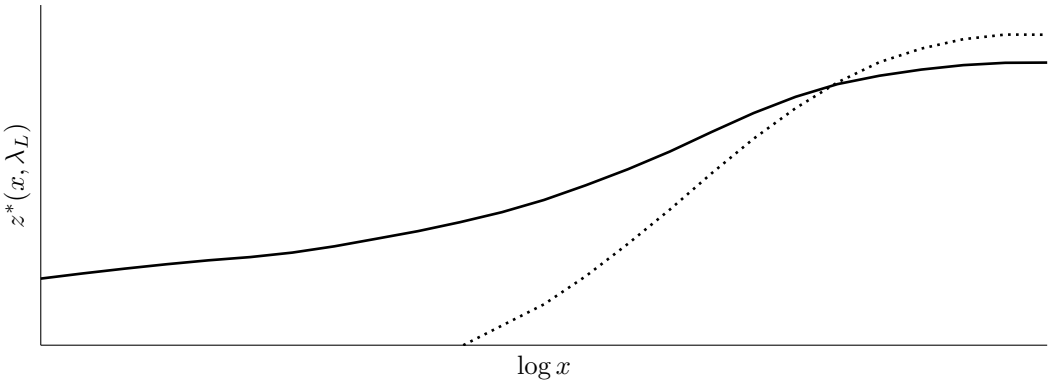
Panel (a) plots the monthly job-finding rate $p_t^*(\Omega, V_t^O(\Omega))$ for unemployed workers against worker productivity z. Unemployed workers enter nonemployment with $\lambda = \bar{\lambda}_L$, so the job-finding rate depends only on x and z. Panel (b) plots the expected duration of nonemployment for workers without a job against z. Each curve corresponds to a different level of aggregate risk premia: one unconditional standard deviation below, at, and above the unconditional mean of $\log x$.

Figure A.2: Risk Premia and Thresholds for Worker Search and Separation

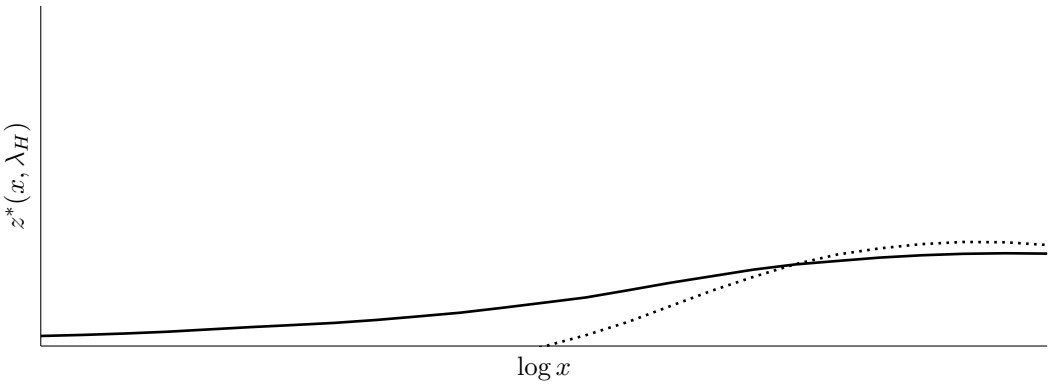
(a) Worker Search Threshold: $\underline{z}(x)$



(b) Worker Termination Threshold, Low Experience: $z^*(x, \bar{\lambda}_L)$

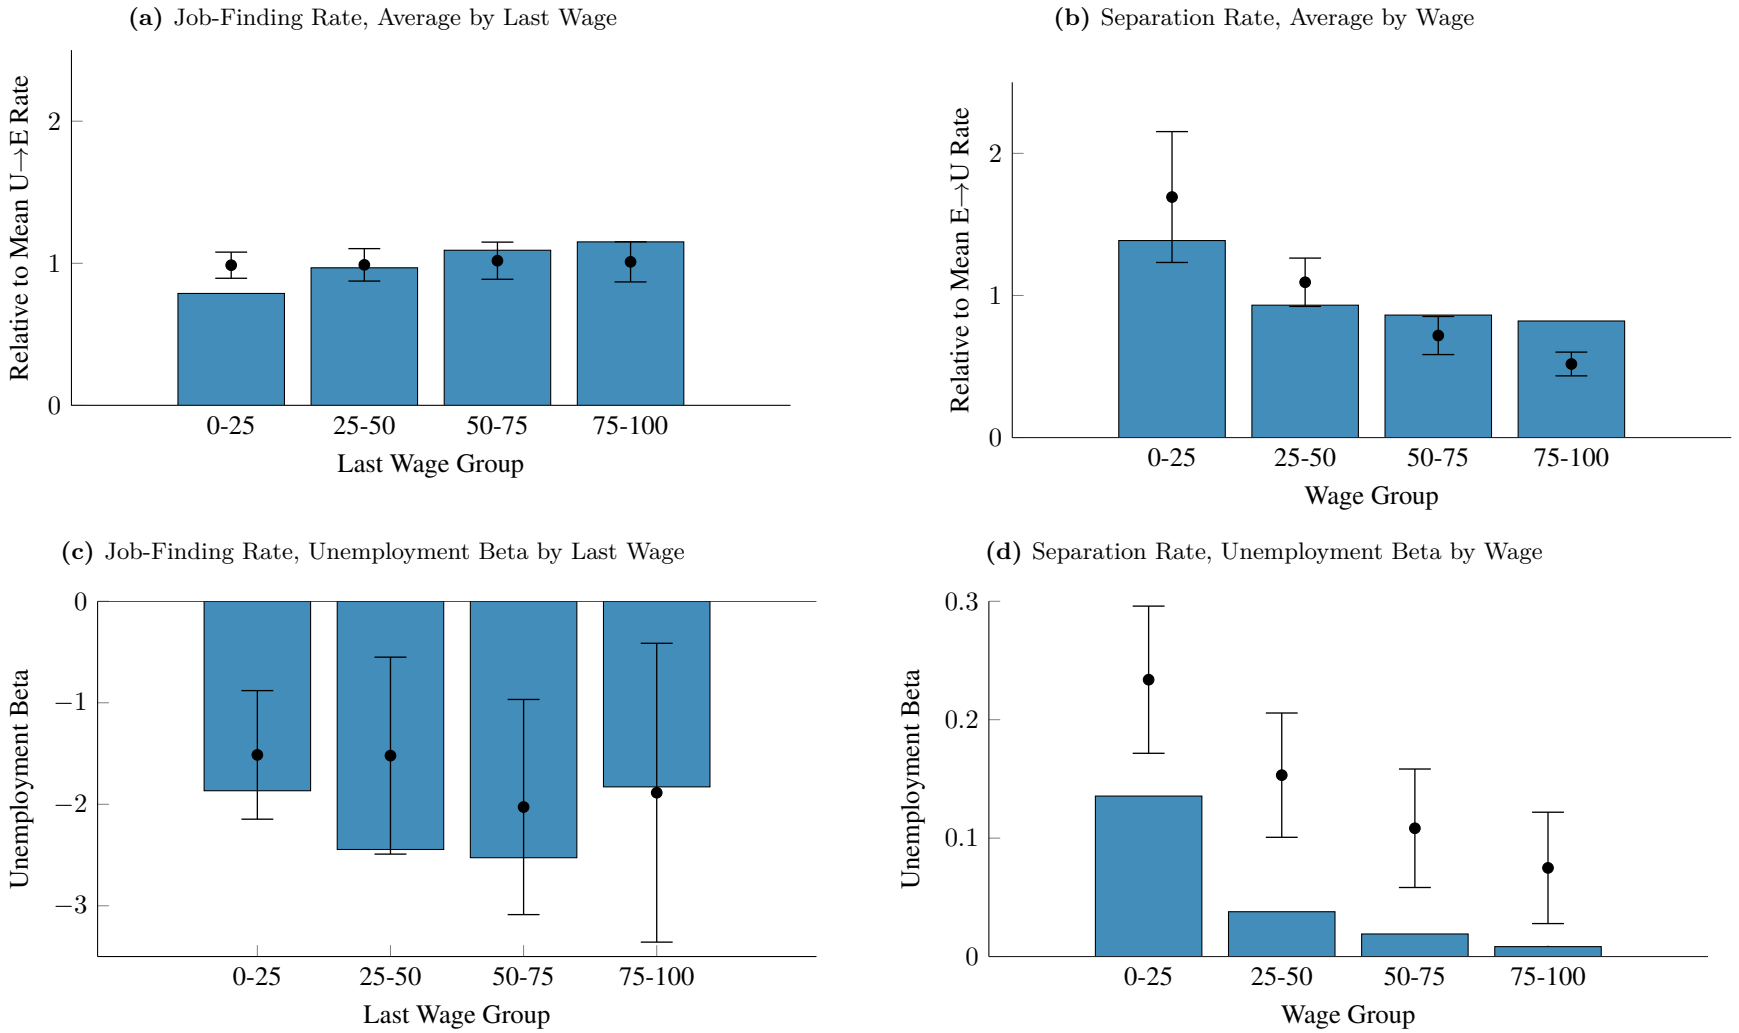


(c) Worker Termination Threshold, High Experience: $z^*(x, \bar{\lambda}_H)$



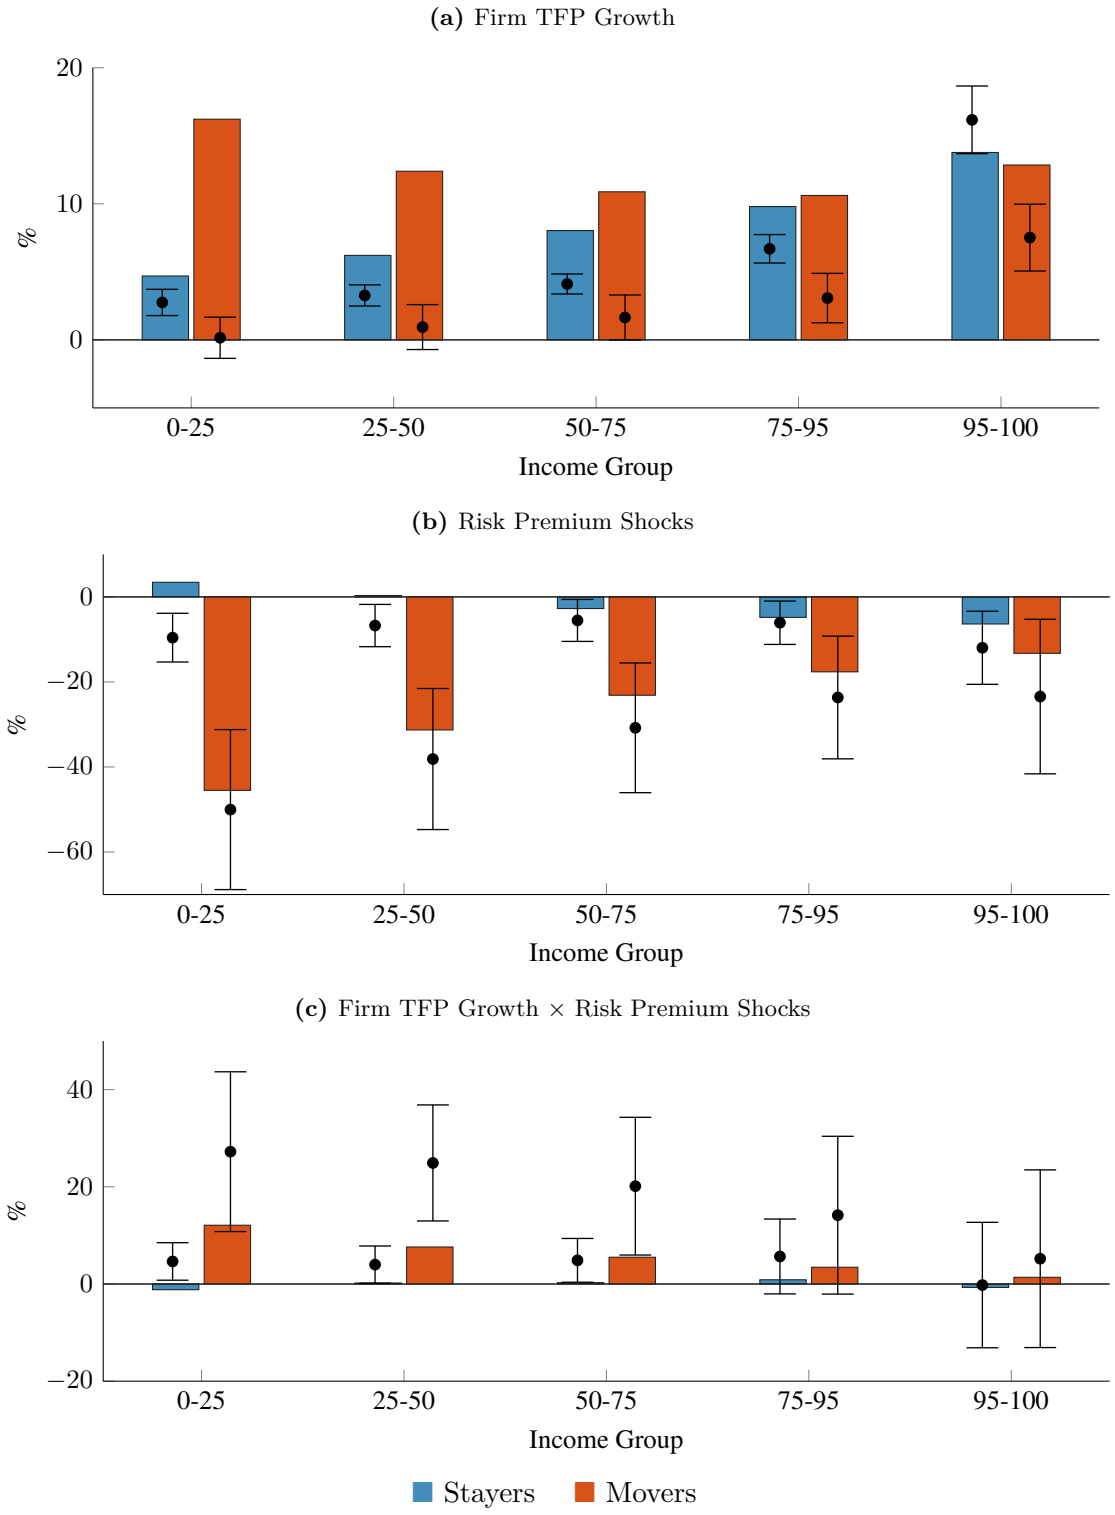
Panel (a) plots the search threshold $\underline{z}(x)$ as a function of risk premia x. A nonemployed worker enters the search pool when $z \geq \underline{z}(x)$, i.e., when $V_t^U(\Omega) \geq V_t^N(\Omega)$, and exits the labor force otherwise. Panels (b) and (c) plot the separation threshold $z^*(x, \lambda)$ as a function of x for low-experience ($\bar{\lambda}_L$) and high-experience ($\bar{\lambda}_H$) workers, respectively. An existing match is continued when $z \geq z^*(x, \lambda)$, i.e., when $J_t(\Omega, V_t^O(\Omega)) \geq 0$. Solid lines: baseline calibration with risk-averse workers ($\gamma > 0$). Dotted lines: counterfactual with risk-neutral workers ($\gamma = 0$).

Figure A.3: Separation and Job-Finding Rates by Worker Income: Model vs. Data



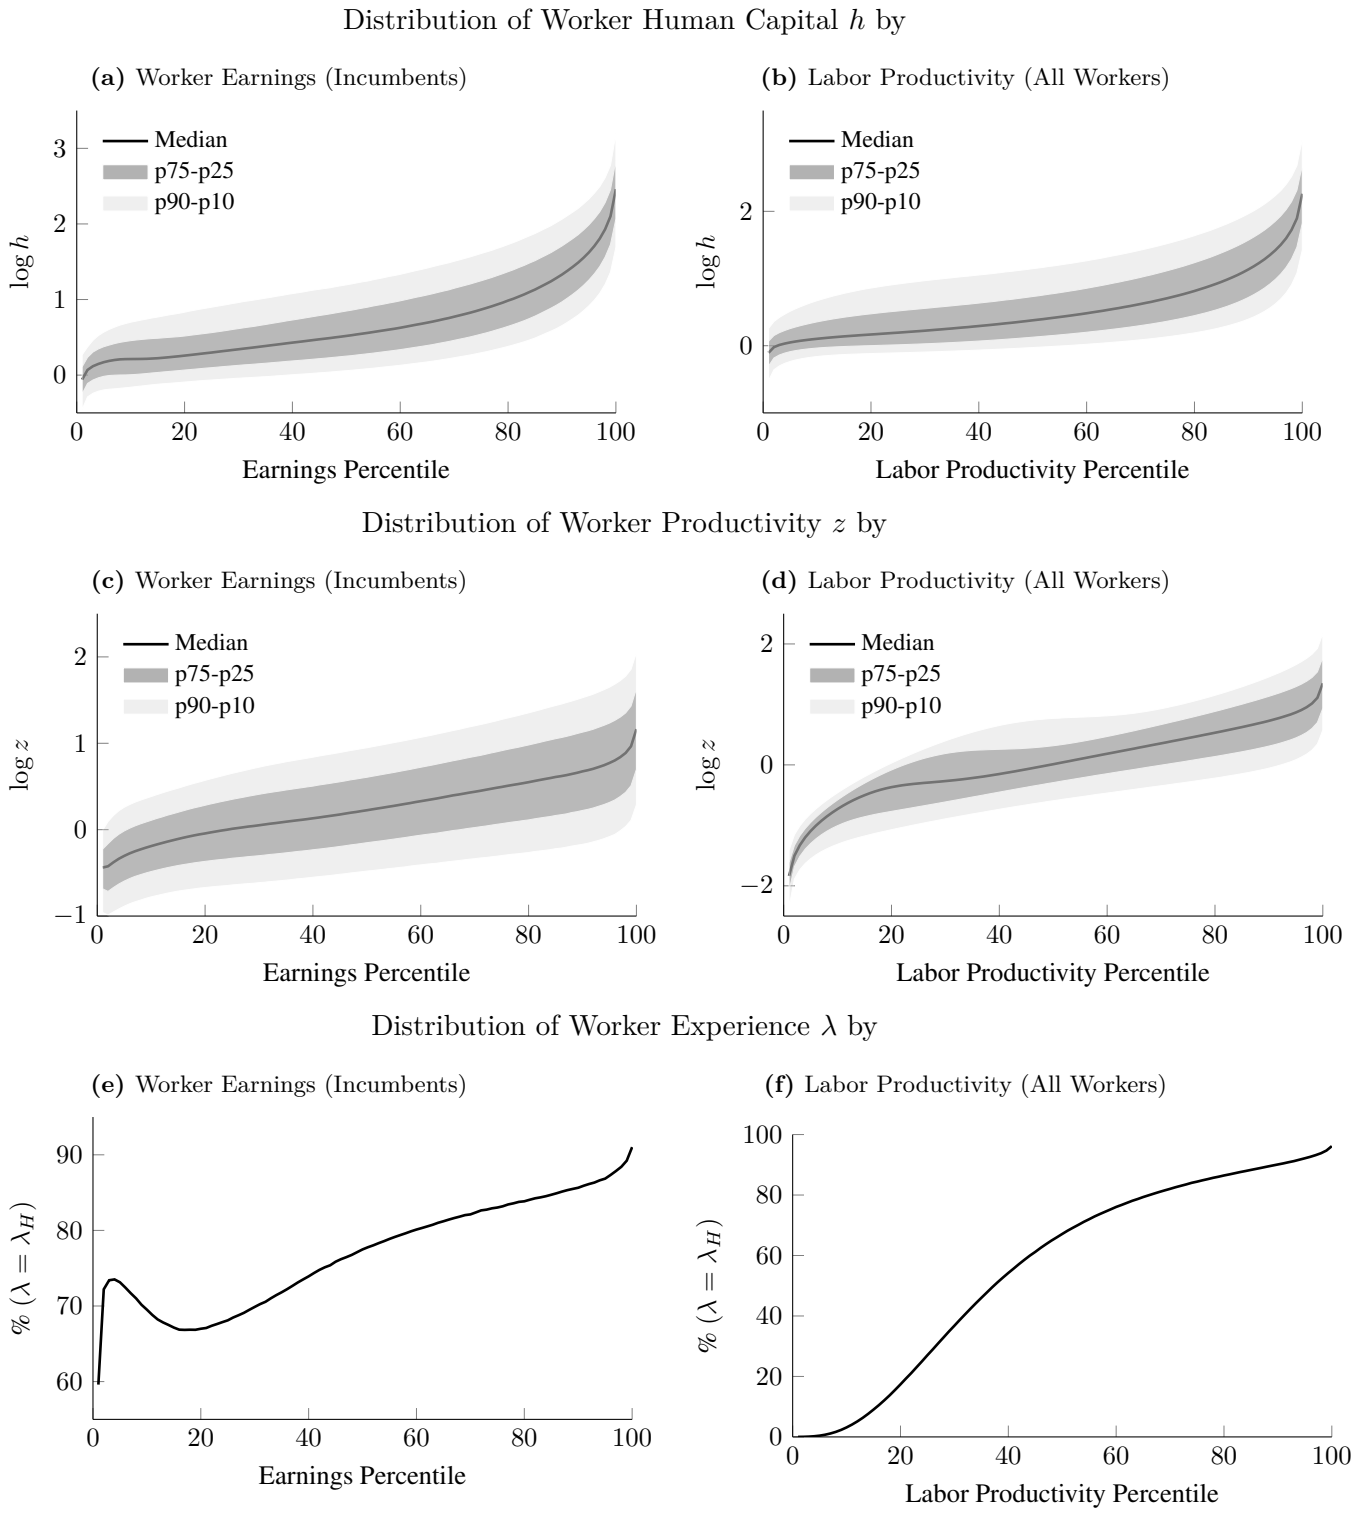
This figure compares the average and cyclicality (unemployment beta) of the job-finding rate (U→E) and the separation rate into unemployment (E→U) by income group in the model and in the data. The empirical counterparts are computed from the SIPP, adjusted for flow level differences from the CPS. Unemployed workers in panels (a) and (c) are binned into groups based on their earnings the last time they were employed in the prior twelve months (if any). Incumbent workers in Panels (b) and (d) are binned into groups based on their current wage earnings.

Figure A.4: Response of Wage Earnings, Stayers vs. Movers: Model vs. Data



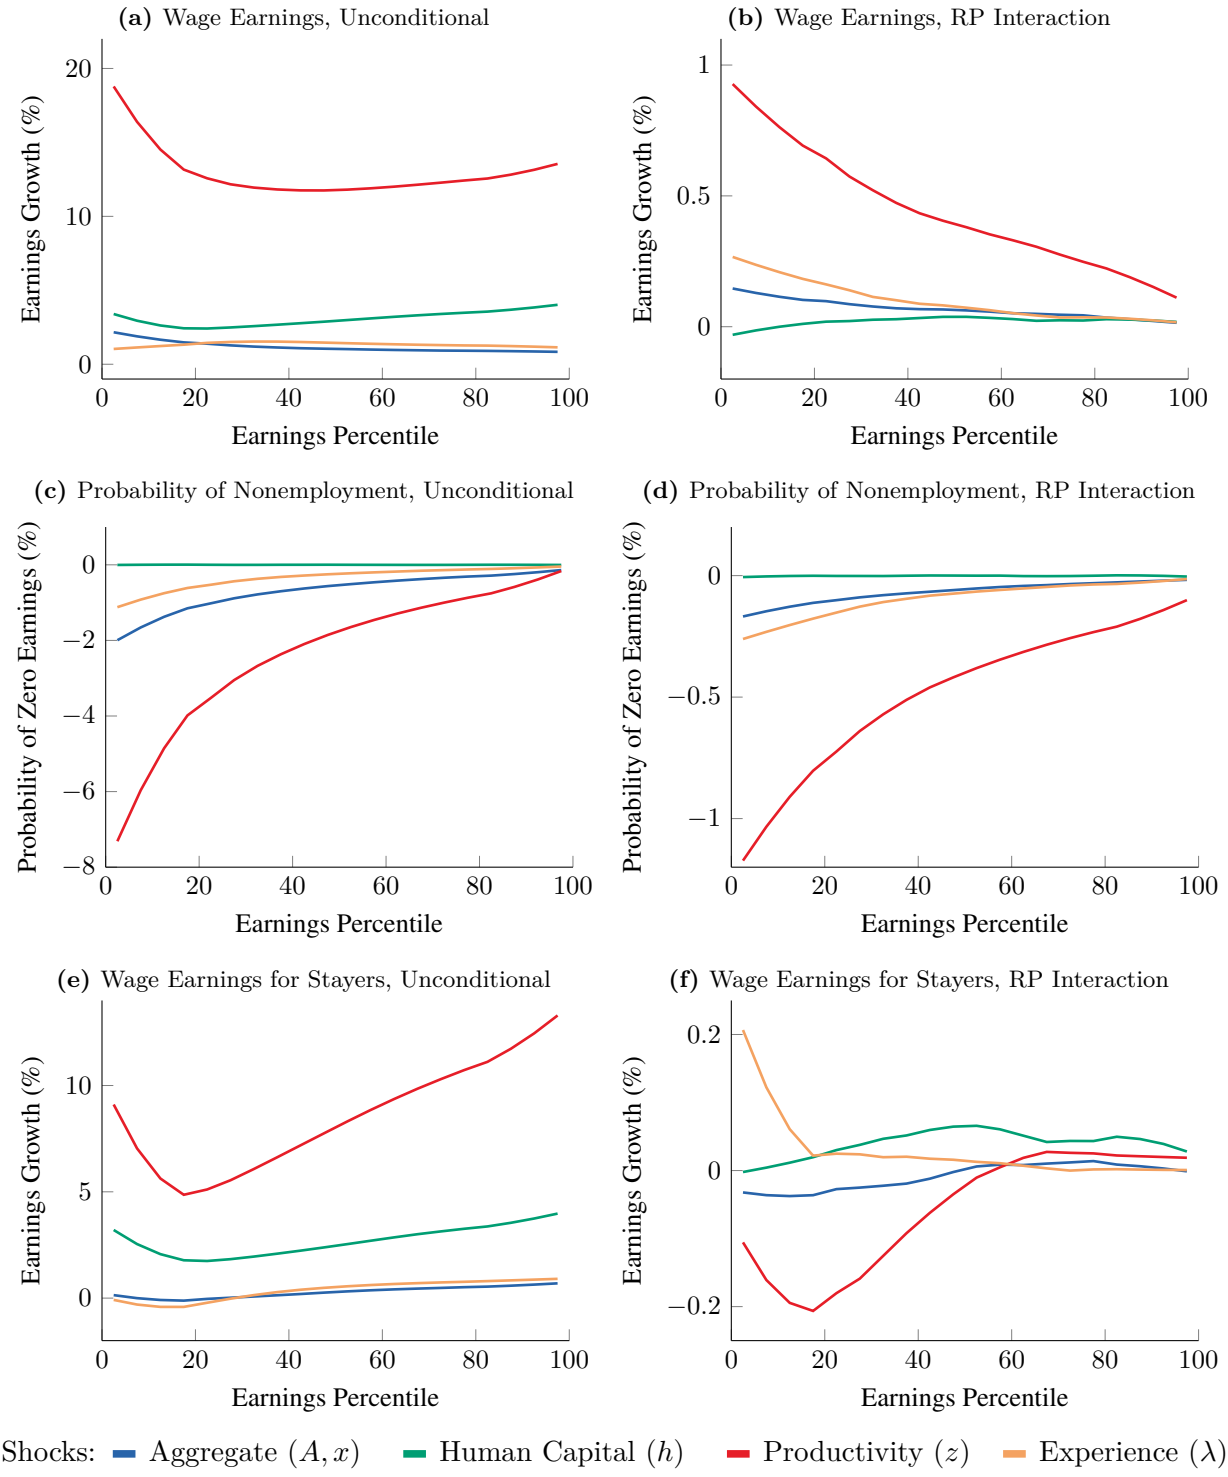
This figure compares estimated coefficients from regression (33) in the data (dots with 95% confidence intervals) and in model-simulated data (bars), by earnings group at a three-year horizon, estimated separately for stayers (workers who remain with their initial employer) and movers (workers who leave). The outcome variable is cumulative age-adjusted earnings growth. Panel (a): pass-through of firm TFP growth ($\beta_{0,s}$). Panel (b): pass-through of risk premium shocks (γ_s). Panel (c): interaction of firm TFP growth with risk premium shocks ($\beta_{1,s}$).

Figure A.5: Distribution of Worker Productivity Components



This figure plots the ergodic distribution of the three worker productivity components—permanent human capital h (panels a–b), persistent productivity z (panels c–d), and experience λ (panels e–f)—in the calibrated model. Left panels sort incumbent workers by earnings percentile; right panels sort all workers (including nonemployed) by labor productivity percentile.

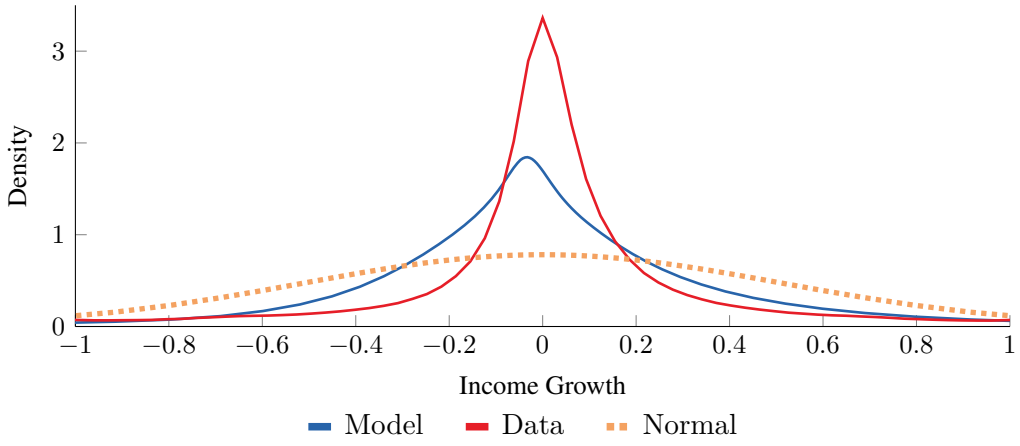
Figure A.6: Response of Worker Earnings and Nonemployment to Model Shocks



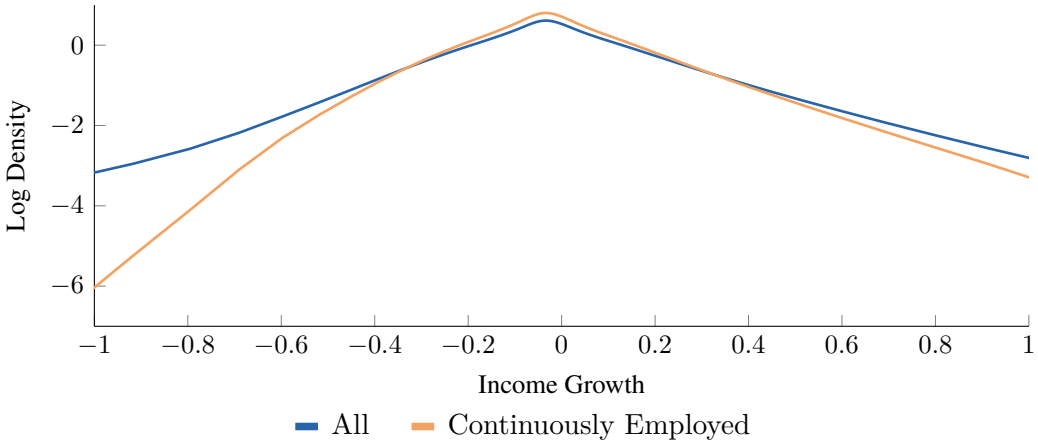
This figure shows the response of annual worker earnings and nonemployment to a one annual standard deviation shock to each model variable, by prior earnings percentile. Left panels plot unconditional responses; right panels plot the additional effect when risk premia x are one standard deviation above their mean.

Figure A.7: Distribution of Worker Earnings Growth

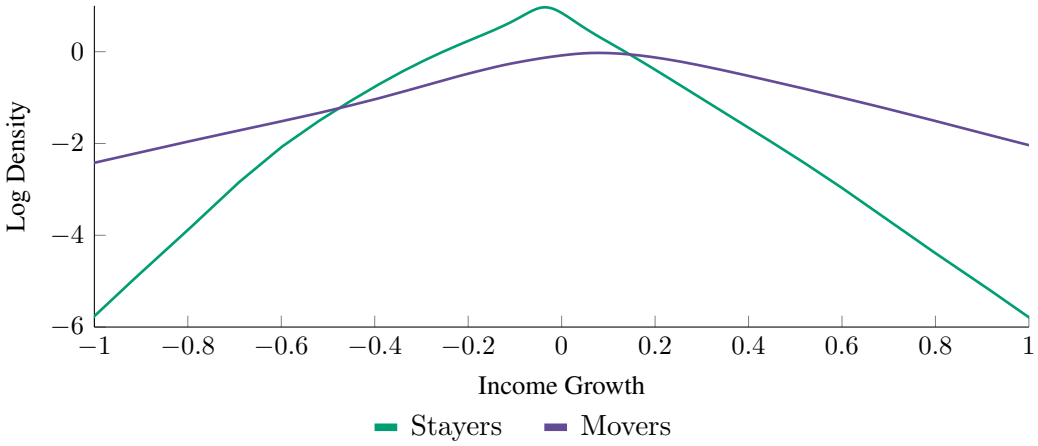
(a) Density of Annual Earnings Growth



(b) Log Density of Annual Earnings Growth: All vs. Continuously Employed

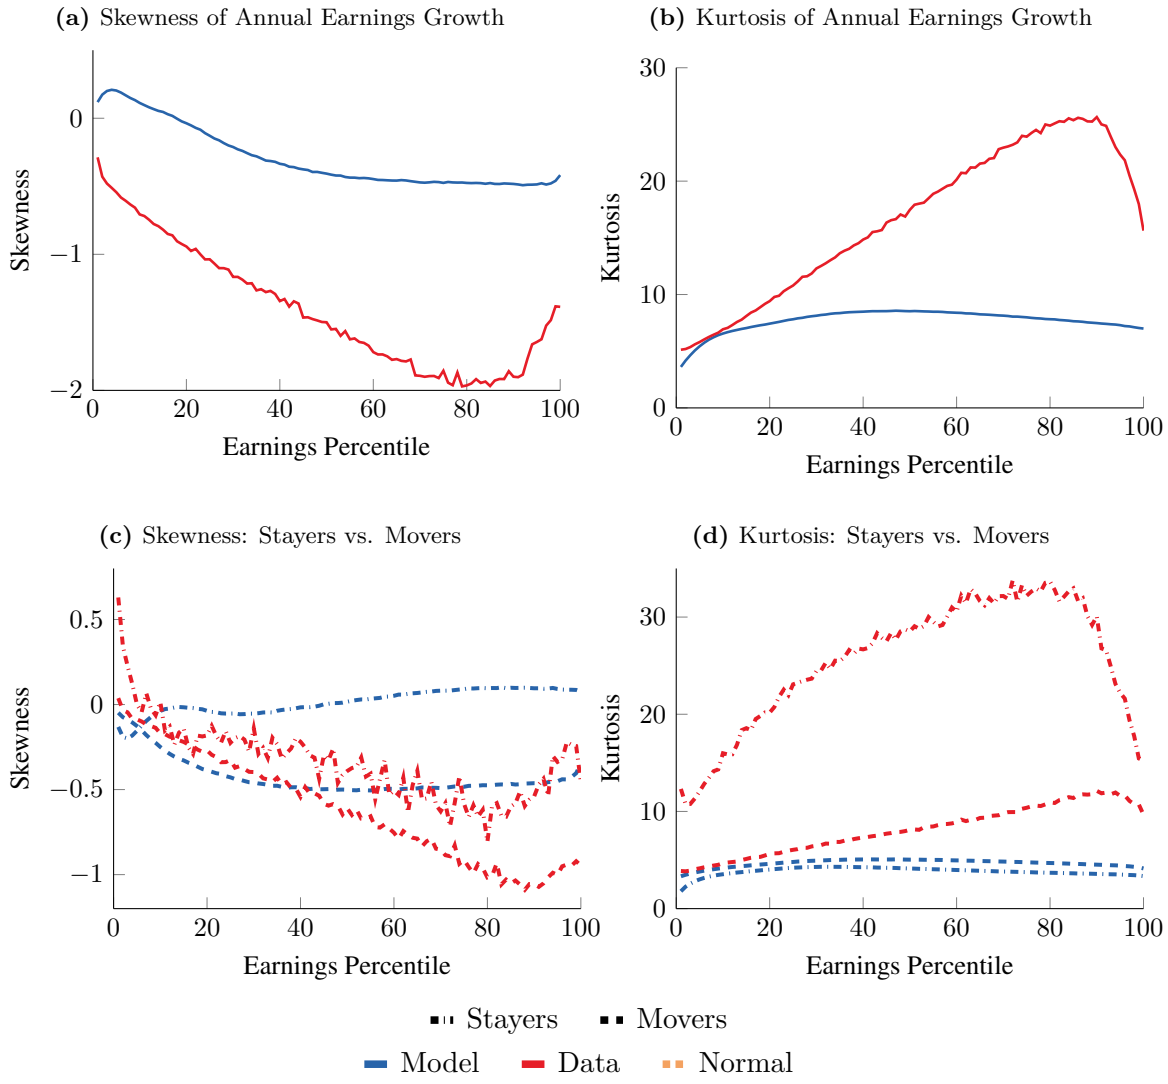


(c) Log Density of Annual Earnings Growth: Stayers vs. Movers



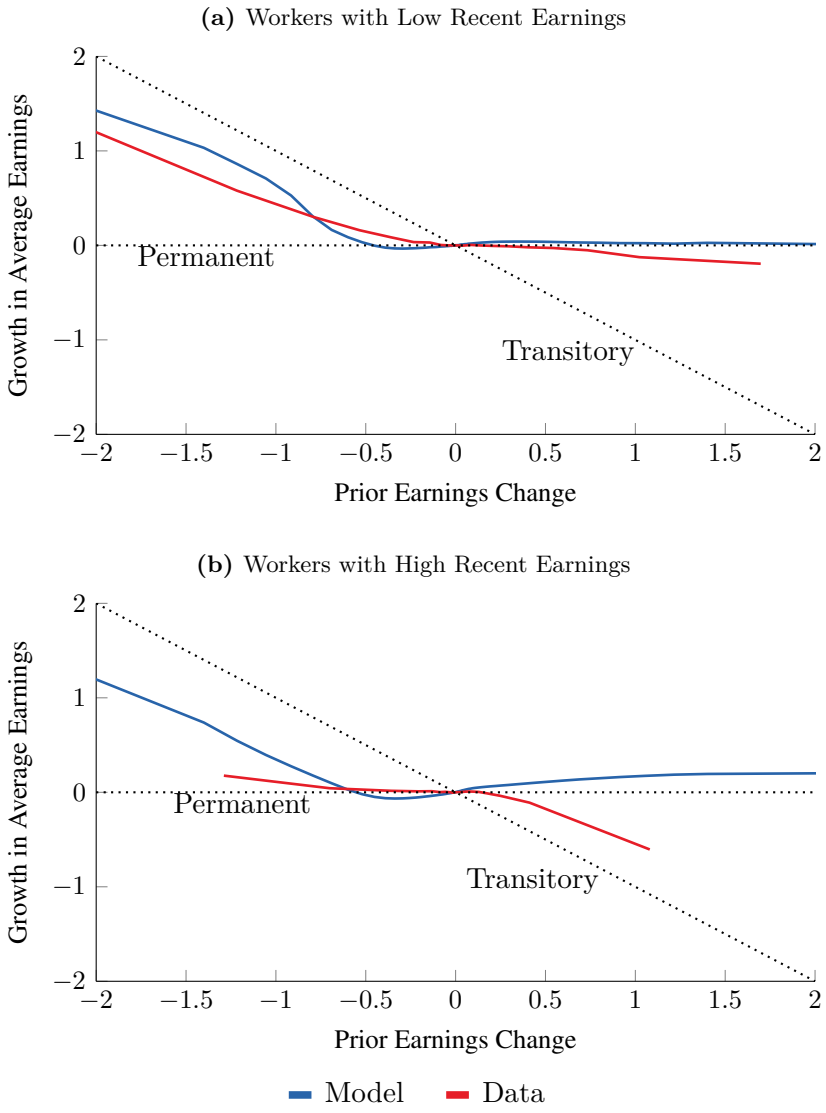
Panel (a) compares the density of annual earnings growth to the data (Guvenen et al., 2021). Panel (b) plots the log density for all workers versus continuously employed workers (employed throughout years t and $t+1$) in the model, and panel (c) plots the log density for stayers versus movers.

Figure A.8: Labor Income Risk, Higher Order Moments: Model vs. Data



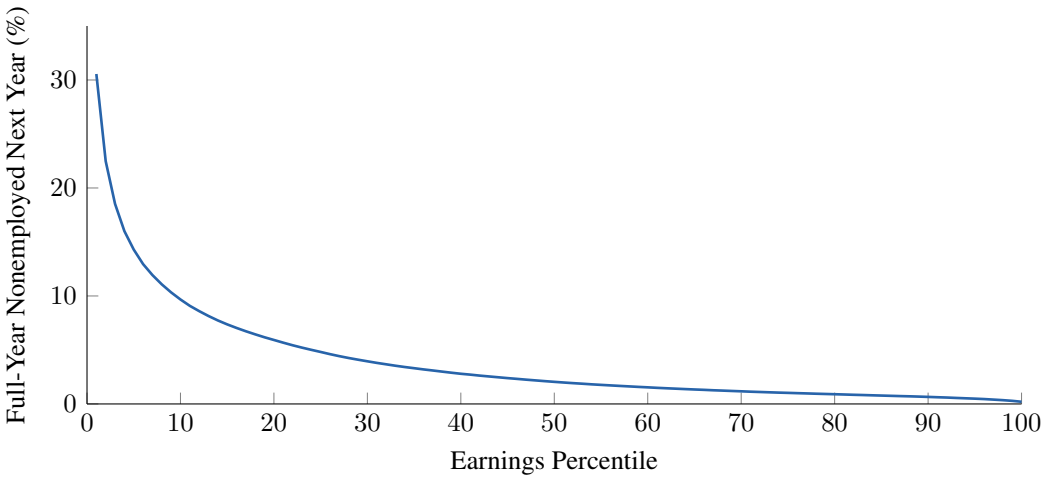
This figure compares higher-order moments of worker earnings risk in the model and in the data. The data series are from Guvenen et al. (2021).

Figure A.9: Impulse Responses of Earnings Growth



This figure plots impulse response functions of earnings, following Guvenen et al. (2021). Workers are sorted by their initial earnings change from $t-1$ to t into quantiles. For each quantile, the y-axis shows the change in log average earnings from t to $t+1$, plotted against the average initial change on the x-axis. Panel (a) shows workers with low recent earnings (P6–P10); panel (b) shows workers with high recent earnings (P91–P95). A flat line at zero indicates full persistence; a line with slope -1 indicates full mean reversion.

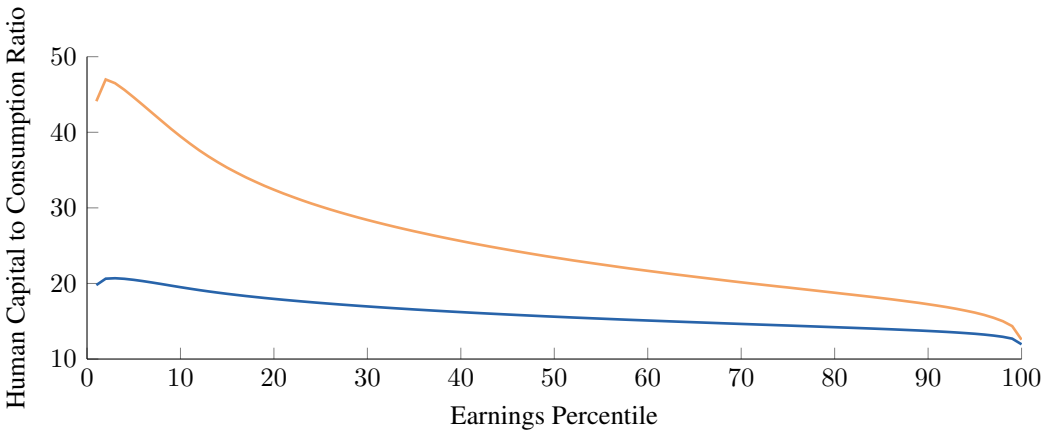
Figure A.10: Fraction of Nonemployed Workers



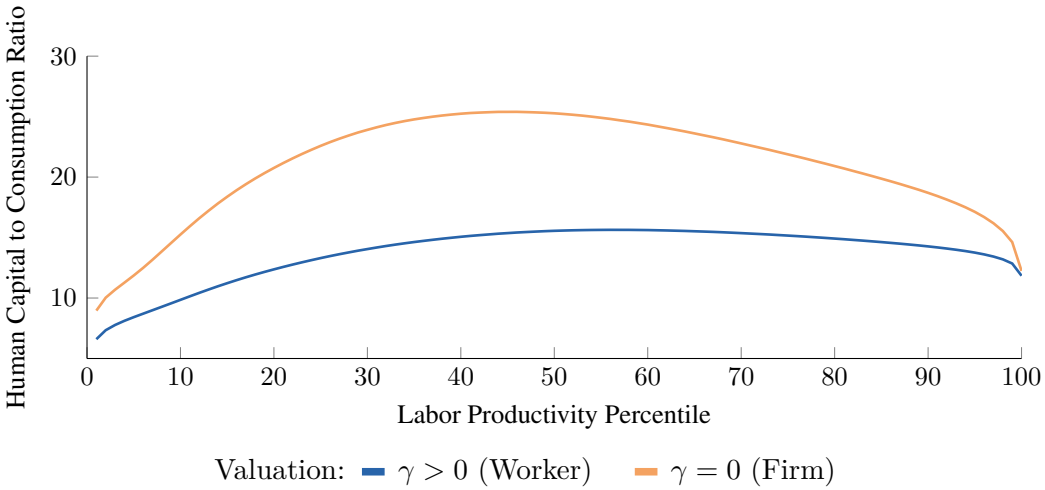
This figure plots the fraction of workers with positive earnings in year $t-1$ that become full-year nonemployed in year t, sorted by recent worker earnings percentile.

Figure A.11: Value of Human Capital in Model

(a) Average Human Capital to Consumption Ratio by Earnings (Incumbent Workers)

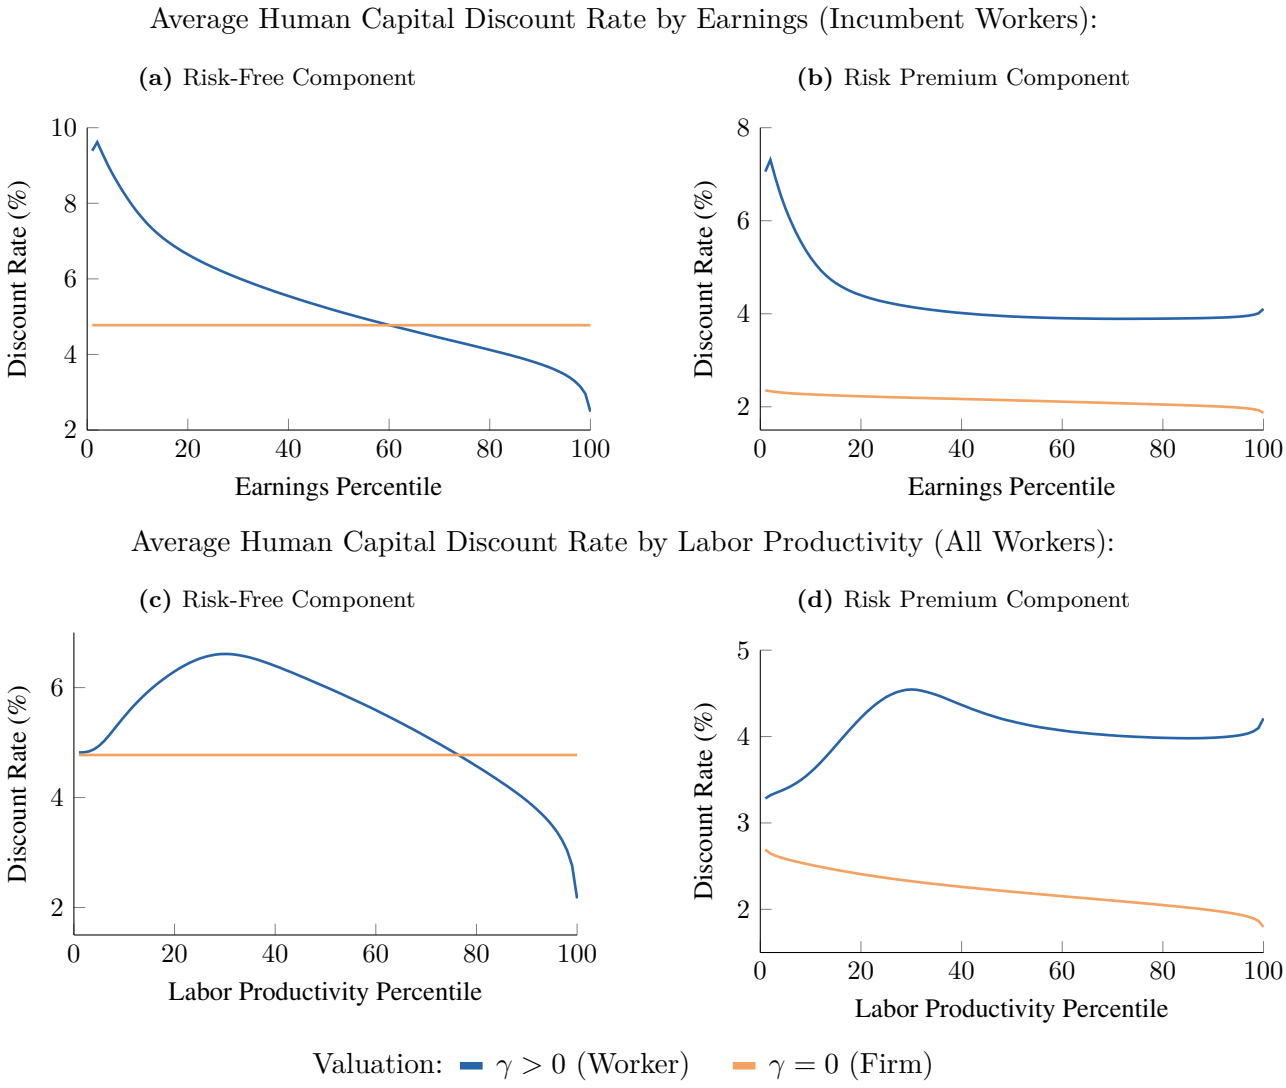


(b) Average Human Capital to Consumption Ratio by Labor Productivity (All Workers)



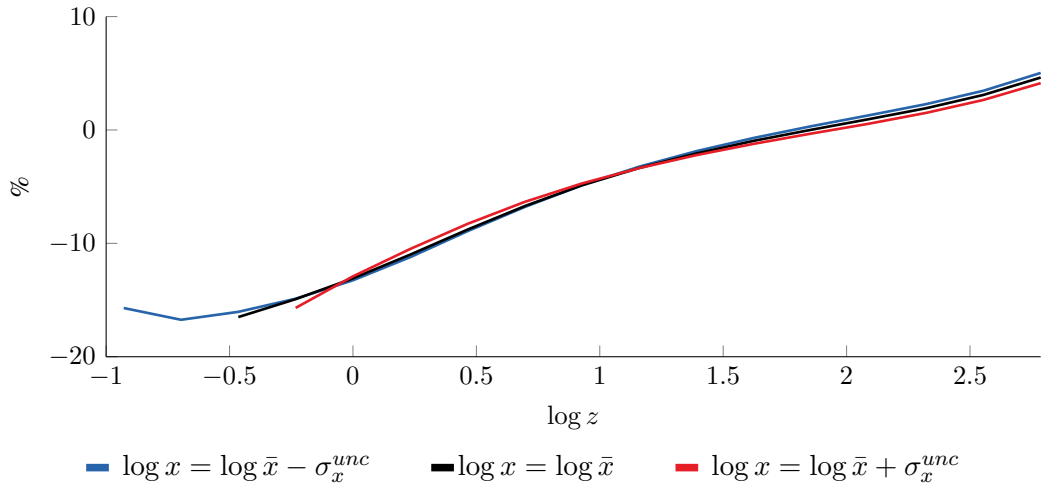
Panels (a) and (b) plot the annualized ratio of human capital to current consumption—under the worker ($\gamma > 0$) and firm ($\gamma = 0$) valuations—for incumbent workers sorted by prior earnings and for all workers sorted by labor productivity, respectively.

Figure A.12: Human Capital Discount Rate in Model: Decomposition



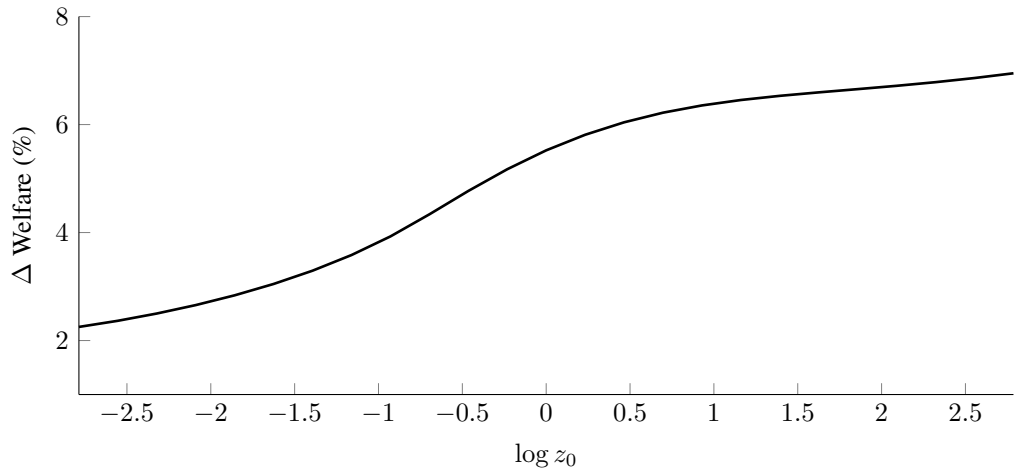
This figure plots a decomposition of annualized human capital discount rates—under the worker ($\gamma > 0$) and firm ($\gamma = 0$) valuations—for incumbent workers sorted by prior earnings and for all workers sorted by labor productivity, respectively. The human capital discount rate is the internal rate of return: for each worker, we compute expected payoffs by horizon, then find the per-period discount rate such that the resulting net present value equals the value of human capital. The risk-free component is the corresponding discount rate under risk-neutral valuation, and the risk premium component is the difference between the human capital discount rate and the risk-free component.

Figure A.13: Value of Insurance: Firm NPV of Wages vs. Piece Rate Contract for New Hires



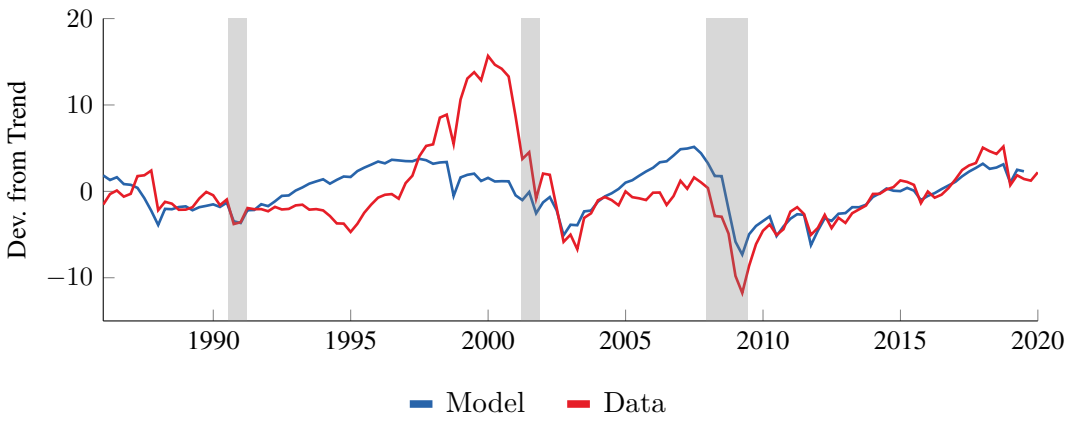
This figure plots the net present value of wages under the optimal contract relative to a piece-rate contract that sets wages proportional to productivity, for newly hired workers. The NPV is computed using the firm's stochastic discount factor. The horizontal axis is initial productivity z; each line corresponds to a different level of the aggregate risk premium x.

Figure A.14: Welfare Gains From Eliminating On-the-Job Search



This figure plots welfare gains from a counterfactual in which the on-the-job search probability $\chi(\Omega)$ is set to zero, as a function of initial productivity z_0, holding other parameters constant. Welfare gains are measured as the percentage change in certainty-equivalent consumption for newborn workers (integrated over aggregate states) relative to the baseline economy.

Figure A.15: Realized Fluctuations in Price–Earnings Ratio: Model vs. Data ($\rho = 53\%$)



This figure plots the price–earnings ratio in the model (computed by feeding the model the empirical series of risk premium shocks ϵ^{rp}) and in the data (Shiller's CAPE ratio).

Table A.1: Asset Market Moments: Model vs. Data

	Model	Data
Average P/E	17.6	18.2
Autocorrelation of $\log P/E$	0.9	0.9
Average risk-free return (%)	1.4	1.4
Average excess market return (%)	8.2	7.9
Volatility of excess market return (%)	20.3	20.0

This table reports targeted asset market moments in the model and in the data. The price–earnings ratio and equity returns in the model are computed from a levered claim on aggregate TFP.