Time-Varying 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
MIT Sloan School of Management

February 2026

Abstract

We develop a directed search model with optimal dynamic wage contracts, two-sided limited commitment, and time-varying risk premia to study how financial conditions shape firm insurance of worker earnings risk. Two forces govern the optimal contract: an insurance motive that smooths wages and a retention motive that deters poaching. When risk premia rise, limited commitment constraints tighten and insurance erodes—especially for low-wage workers near the separation margin—generating state-dependent labor income risk. Using U.S. administrative data, we document consistent evidence: the pass-through of firm productivity shocks to worker earnings rises when risk premia are elevated, concentrated among lower-paid workers and operating through job destruction. The calibrated model matches labor market dynamics, asset prices, and the heterogeneous pass-through of firm shocks. Beyond these targeted moments, it generates realistic variation in non-Gaussian earnings dynamics across workers and over time, sizable welfare losses from idiosyncratic risk, depressed human capital valuations, and large welfare gains from state-contingent policies.

*We thank Brett Green, Virguliu Midrigan, and Giuseppe Moscarini for extremely helpful comments related to this paper. 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).

What are the key drivers of earnings risk across groups of workers and over the business cycle? Individual income growth rates are highly volatile and far from Gaussian—featuring 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). The aggregate unemployment rate moves considerably more than many benchmark search-theoretic models predict (Shimer, 2005), and job separations increase sharply in recessions. These patterns are hard to square with the large potential gains from firms insuring workers against idiosyncratic risk (see, e.g., Knight, 1921; Baily, 1974; Azariadis, 1975). Firms do insure workers, albeit imperfectly: they absorb transitory shocks to their own productivity but pass through permanent ones (Guiso, Pistaferri, and Schivardi, 2005), the degree of pass-through rises with worker income (Friedrich, Laun, Meghir, and Pistaferri, 2019; Chan, Xu, and Salgado, 2019), and many workers lose this insurance entirely upon job loss. If firms can insure workers, why do so many workers get fired and experience large income declines rather than accepting pay cuts and staying with their existing employers, especially in recessions?

We develop a directed search model that jointly explains these patterns with two key features: long-term contracting and time-varying risk premia. Workers are heterogeneous in their productivity and experience levels, where both evolve stochastically. Firms partially insure risk-averse, hand-to-mouth workers through state-contingent wage payments, but cannot commit to preserving matches with negative present values. Likewise, workers cannot commit to staying in matches that pay less than their outside option. The model generates realistic unemployment fluctuations with procyclical job finding and countercyclical separations, endogenous and state-dependent pass-through of firm shocks to worker earnings, and a U-shaped cross-sectional pattern in pass-through—high at the bottom of the earnings distribution (through job destruction) and high at the top (through wage adjustment)—with a countercyclical time-series pattern concentrated among low-wage workers. Further, it generates many other realistic patterns of earnings risk across the income distribution, both overall and conditional on job transitions, and allows us to evaluate the potential implications of state-contingent policies for labor market outcomes and welfare.

On the intensive margin, two forces shape the optimal wage contract for workers who remain employed: an *insurance motive*—firms smooth wages to shield workers from idiosyncratic consumption risk—and a *retention motive*—firms backload compensation to deter on-the-job poaching. When limited commitment constraints are slack and on-the-job search relatively inactive, the insurance motive dominates and wages are smooth. When either constraint binds, wages must adjust to satisfy it, eliminating the scope for insurance. When risk premia rise, match surplus falls, the firm's constraint binds for more workers, and insurance breaks down more often—especially for lower-paid workers near the separation threshold. At the top of the earnings distribution, the retention motive is more important:

stayers' wages respond to shocks because firms use backloaded incentives to deter outside offers. The degree of pass-through is therefore endogenous, heterogeneous across workers, and state-dependent.

On the extensive margin, the model builds on the duration-based mechanism of Kehoe, Lopez, Midrigan, and Pastorino (2023) and Meeuwis, Papanikolaou, Rothbaum, and Schmidt (2025). Nonemployment spells have long-lived consequences because productivity grows more rapidly when employed than when nonemployed and nonemployment spells entail a loss of job-specific experience. Workers with low productivity and experience therefore have backloaded benefits of employment that are more sensitive to discount rates than their outside option of nonemployment. The result is higher rates of endogenous job destruction when risk premia are elevated that are concentrated among the lowest-paid workers. Since investing in posting a vacancy involves paying upfront costs in exchange for future benefits, when risk premia rise, they post fewer vacancies and in turn job finding rates decline. These features also imply that transitory fluctuations in risk premia generate persistent declines in aggregate employment and output, because displaced workers lose experience and human capital during nonemployment and reenter the labor market at lower productivity levels.

Consistent with our model's predictions, we document a new empirical finding using U.S. administrative data: the pass-through of firm-specific TFP shocks to worker earnings rises with aggregate risk premia. This amplification is concentrated among lower-paid workers and operating primarily through higher rates of job destruction rather than wage cuts among job stayers. This pattern is robust to controlling for aggregate output growth and NBER recessions, holds across horizons of one to five years, and is confirmed by a movers-versus-stayers decomposition showing that the amplification is substantially larger for workers who leave their initial employer.

We calibrate the model to jointly match moments from asset markets, labor market dynamics, and the pass-through of firm productivity and risk premium shocks to worker earnings. The model reproduces realistic unemployment fluctuations driven by procyclical job finding and countercyclical separations, the heterogeneity in job-finding and separation rates across the earnings distribution, and the U-shaped pass-through of firm shocks to worker earnings. Feeding in our estimated series of risk premium shocks, the model generates time series for unemployment, job-finding rates, and labor market tightness that track the data closely, with correlations ranging from 50 to 80 percent.

Beyond these targeted moments, the model replicates many realistic features of income risk that are not targeted by our calibration. The optimal contract features a "kinked" response: wages are very stable when limited commitment constraints are slack but adjust sharply when they bind or when the match dissolves. This nonlinearity transforms Gaussian productivity shocks into non-Gaussian earnings dynamics. The resulting distribution of individual earnings growth matches the fat tails, excess kurtosis, and negative skewness documented by Guvenen et al. (2021), with a model kurtosis of 11.3 versus 14.9 in the data. The model also replicates how earnings risk varies across the income distribution—declining volatility from bottom to top, higher volatility for job movers

than stayers—and matches the large dispersion in cumulative lifetime earnings growth and lifetime employment rates. A model-generated series tracks the realized time path of countercyclical left-tail risk documented by Guvenen et al. (2014), with a correlation of 59 percent. It generates large and persistent earnings losses from job displacement that are amplified in recessions, consistent with Davis and Von Wachter (2011), and persistent earnings shortfalls for cohorts entering the labor market during downturns, consistent with Kahn (2010) and Schwandt and von Wachter (2019). Hence, while our calibration targets heterogeneity over the business cycle and passthrough of firm shocks, our framework also explains other patterns of heterogeneity in earnings risk.

The model's realistic earnings dynamics imply large welfare costs from idiosyncratic risk. We measure these costs by comparing certainty-equivalent consumption under the baseline calibration ($\gamma = 0.48$) with a counterfactual in which workers are risk-neutral toward idiosyncratic shocks, holding equilibrium dynamics fixed. Despite moderate risk aversion, workers would require a permanent increase in consumption of roughly 30 percent to be compensated for bearing undiversifiable earnings risk. These losses are sizable and fairly uniform across the initial productivity distribution, reflecting the substantial lifetime variation in consumption generated by the combination of non-Gaussian earnings dynamics, persistent displacement costs, and the hand-to-mouth assumption.

Idiosyncratic risk also substantially depresses the valuation of human capital. The ratio of human capital to current consumption—an analog of the price-earnings ratio for equity—averages roughly 15 in the baseline model, compared to nearly 30 if workers are risk neutral with respect to idiosyncratic shocks. The gap is largest for low-wage workers, who face the highest endogenous separation risk and the most volatile earnings paths; they therefore apply the largest subjective risk premium for idiosyncratic risk. We also calculate subjective discount rates on human capital as the internal rate of return consistent with the valuations discussed above, both with and without idiosyncratic risk. These subjective discount rates are considerably higher with idiosyncratic risk than without, especially for low-wage workers. These discount rates also fluctuate over the business cycle, peaking when risk premia are elevated.

Finally, we use the model to evaluate state-contingent labor market and financial policies. We consider two balanced-budget labor market policies activated in recessions: a countercyclical unemployment benefit subsidy, which provides additional insurance when job-finding rates are low and workers' marginal utility is high, and an employment subsidy that lowers the cost of retaining workers during downturns. Both raise ex-ante welfare, though through different channels: the unemployment subsidy improves insurance but weakens employment incentives, while the employment subsidy strengthens job creation and retention but is less well-targeted to the workers who need it most. Intriguingly, more is not better; welfare benefits from policies that only activate in the 10% of states with the highest risk premia are larger than those from policies that activate in the 20% of states with the highest risk premia. We also consider a financial policy that caps

risk premia—a "Fed Put." Since it is unclear to evaluate the cost of such a policy, we only restrict attention to quantifying its benefits. Even when the cap binds only in the worst 0.1 to 0.5 percent of periods, anticipated intervention raises GDP by 6 to 12 percent and welfare by 8 to 22 percent. The bulk of these gains come from anticipation rather than direct intervention: when agents know that extreme risk premia will not materialize, they value jobs and human capital more, firms create more vacancies, and fewer matches are destroyed. Unanticipated intervention, by contrast, has modest effects—underscoring that the credibility of the backstop matters more than its use.

Our paper connects three literatures that have developed largely in parallel. Most closely related are models of risk premia and labor markets: Kehoe, Midrigan, and Pastorino (2019); Kehoe et al. (2023) and Meeuwis et al. (2025) show that time-varying risk premia drive unemployment fluctuations through endogenous separations in directed search frameworks, but assume complete markets or exogenous wage rules and do not study income risk. We build directly on Meeuwis et al. (2025), replacing their exogenous wage rule with an optimal dynamic contract, adding risk-averse workers and on-the-job search, so that the extent of firm insurance is endogenous and state-dependent. These features create genuine demand for firm-provided insurance and make the welfare consequences of earnings risk first-order. In the implicit contracts literature, Rudanko (2011) embeds limited commitment contracts in a frictional labor market but without time-varying risk premia or the retention motive created by on-the-job search. Balke and Lamadon (2022) study dynamic contracts in a search model with incomplete markets, but their model features no time variation in risk premia and firms can commit to retaining ex-post unproductive workers, so that the degree of insurance does not fluctuate with aggregate conditions.

Our implications for earnings dynamics relate to several papers that study different, often complementary mechanisms which replicate some of the patterns we consider. Hubmer (2018) abstracts from the business cycle but shows that a search model generates non-Gaussian earnings growth through frictional transitions on a job ladder, but the mechanism operates entirely through job mobility: stayers' earnings follow a smooth unit root and exhibit no excess kurtosis. In our model, the kinked contract generates excess kurtosis even for stayers, and the frequency of these nonlinear events varies with aggregate financial conditions. Huckfeldt (2022) and Acabbi, Alati, and Mazzone (2026) explain the scarring effects of recessions through selective hiring and the collapse of human-capital ladders, complementing our discount rate-based mechanism. Ai and Bhandari (2021) develop a model in which optimal dynamic contracts between workers subject to idiosyncratic shocks and diversified business owners can generate a time-varying price of aggregate risk. Whereas in their model, countercyclical tail risk in idiosyncratic productivity is an exogenous input into their analysis, in our model countercylical tail risk in earnings growth arises endogenously when risk premia rise. Our empirical findings on the pass-through of firm shocks to worker earnings extend

the large literature on firm insurance (Guiso et al., 2005; Friedrich et al., 2019; Chan et al., 2019) along the time-series dimension.

The remainder of the paper is organized as follows. Section 1 presents the model, and Section 2 discusses its key mechanisms. Section 3 maps the model to the data: we document the empirical evidence on heterogeneous and time-varying pass-through and describe the calibration. Section 4 evaluates the model's fit. Section 5 explores the model's broader implications for earnings risk, welfare, human capital, and policy.

1 Model

We now formalize the environment outlined in the introduction. The model is a frictional labor market (Diamond, 1982; Mortensen, 1982; Pissarides, 1985) with heterogeneous risk-averse workers subject to shocks to their productivity. Firms search for workers with different productivity levels through a directed search process (Montgomery, 1991; Moen, 1997). Worker productivity is stochastic and persistent. Firms partially insure workers against these shocks, subject to a limited commitment friction: they cannot commit to preserving employment matches that ex-post yield a negative value to the firm. If that were to occur, firms adjust wages to the point that the surplus to the firm is weakly positive. If that is not feasible, the match is terminated. Similar to Kehoe et al. (2019, 2023), worker productivity grows faster, on average, during employment than during nonemployment. The model also features on-the-job search (Postel-Vinay and Robin, 2002), which creates the retention motive discussed above.

We model risk premium shocks as shocks to the effective discount rate that agents (both workers and firms) use to value risky future cashflows, in the spirit of Lettau and Wachter (2007) and Meeuwis et al. (2025). A positive risk premium shock leads to a lower valuation of a stream of risky future cashflows. For simplicity, we assume that discount rate shocks are perfectly (negatively) correlated with aggregate productivity. Since the decisions to hire a worker and to maintain an existing worker–firm match involve calculating the present value of the relative benefits of keeping the worker in the job or not and these benefits are uncertain, fluctuations in discount rates directly affect labor allocations. In addition, discount rates also affect the surplus value of the match to firms, and hence indirectly affect the degree of insurance firms can provide to workers.

1.1 Environment

The model is set in discrete time. There is a unit measure of ex ante identical workers who can be employed by a large number of ex ante identical firms. The workers are indexed by i, have heterogeneous productivity, and are either employed by a firm, unemployed and searching for a job, or nonparticipants in labor markets. Firms employ workers to produce output and can post vacancies to attract new workers, targeting workers with a specific productivity level.

Timing

Each period in the model consists of five stages:

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

2. A random share s of existing matches are destroyed for exogenous reasons, and all firms and workers in existing matches decide whether to continue the match or to endogenously terminate the match.

3. Firms randomize the contract values that are promised to continuing workers over a two-point lottery. This lottery serves to ensure that expected firm profits are concave.

4. Search and matching stage: workers who are still employed get a chance to search for a new job with probability χ, and all workers in the unemployment pool search for a new job. Firms post vacancies to attract new workers, and new matches are formed. Poaching offers cannot be countered, so existing matches are terminated if on-the-job search is successful.

5. For continuing and new matches, firms collect output and pay wages to the workers. Workers who are out of a job decide whether to enter the unemployment pool and search for a job in the next period or to remain out of the labor force, and receive the corresponding flow benefits.

Production

The output produced by an employed worker is equal to

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

Here, A_t is aggregate productivity, while $h_{i,t}$, $z_{i,t}$, and $\lambda_{i,t}$ are different components of worker productivity; the worker's current level productivity is summarized by the vector $\Omega = (h, z, \lambda)$.

Aggregate productivity A follows a random walk with a 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-specific productivity h is permanent and evolves according to the following process:

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

We can think of h as capturing general human capital that grows with experience. Importantly, the growth rate of human capital depends on the worker's current employment status: $g_{i,t} \in \{g_E, g_O\}$. As in Ljungqvist and Sargent (1998), human capital grows with work experience, and workers experience long-term costs from being out of a job; therefore, $g_E > g_O$.

The second component z evolves according to a mean-reverting process:

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

The key difference between z and h is that z refers to employment-specific productivity whereas h refers to overall human capital; we model this distinction by assuming that the flow benefits of unemployment and nonparticipation scale perfectly with h but imperfectly with z. As a result, shocks to h do not affect allocations of workers to jobs, but shocks to z do.

The third component of worker productivity $\lambda \in \{\bar{\lambda}_L, \bar{\lambda}_H\}$ is a two-point Markov process; while the worker is employed, λ evolves according to the following transition probability matrix,

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

The component λ is meant to capture an element of job-specific experience. Hence, as long as the worker remains employed, there is a probability f each period that an inexperienced worker becomes experienced at her job. An experienced worker retains her experience level as long as she remains employed. However, if she goes through a nonemployment spell then she loses that job-specific experience—that is, λ is reset to $\bar{\lambda}_L$. Further, a key difference between λ and h or z is that it does not affect the flow benefits of the worker's outside option (unemployment or nonparticipation).

Last, newly-born workers enter the economy at time $t_0(i)$ without a job, no job-specific experience $\lambda_{i,t_0(i)} = \bar{\lambda}_L$, and with idiosyncratic productivity components $h_{i,t_0(i)} = \bar{h}$ and $\log z_{i,t_0(i)} \sim N(\log \bar{z}, \sigma_{z0}^2)$.

Worker Preferences and Payoffs

Workers derive utility from their consumption. Their preferences are defined recursively as

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

Equation (6) defines a worker's utility V, in consumption units, over sequences of consumption $c_{i,t}$. The recursive formulation is of the Epstein and Zin (1989) form, subject to one modification: workers' attitudes towards aggregate risk are modified by the process

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

where the effective degree of worker risk aversion towards aggregate shocks is time-varying,

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

Examining our specification for worker utility (6)–(8), we note that at each date t, workers evaluate their continuation utility under a distorted conditional distribution for $\varepsilon_{A,t+1}$. This probability distortion is captured by the Radon-Nikodym derivative (7), where x_t is an $\mathcal{F}_t$ measurable

process. This specification can be microfounded under the assumption that there exists ambiguity about the distribution of the aggregate shock ε_A, in the spirit of Hansen and Sargent (2001). Here, we allow the degree of the agent's local desire for robustness or the degree of ambiguity aversion, captured by (8), to vary over time as in Epstein and Schneider (2003).[1] By contrast, workers have no ambiguity about their own productivity: their effective degree of risk aversion over idiosyncratic shocks is γ. Hence, shareholders do not care about idiosyncratic shocks while workers evaluate their idiosyncratic shocks under the physical probability measure subject to a level of risk aversion given by γ. Skiadas (2013) shows that, under these assumptions, agents exhibit source-dependent risk aversion: workers exhibit time-varying risk aversion over aggregate shocks that is a function of x_t but constant risk aversion towards their idiosyncratic shocks equal to γ. Hence aggregate shocks are effectively concavified relative to idiosyncratic shocks in worker preferences. Last, given this specification, the workers' elasticity of intertemporal substitution (EIS) is constant and equal to $1/\gamma$.

Workers are hand-to-mouth: they do not have access to financial markets. Hence, employed workers consume their flow wage: $c_{i,t} = w_{i,t}$. Workers that are out of a job receive benefits and engage in home production; they receive and consume a payoff that depends on their labor market participation status. Workers that decide to enter the unemployment pool and actively search for a new job consume

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

Following Hall (2017) and Kehoe et al. (2023), the opportunity cost of employment has a unit elasticity to aggregate productivity A and to the permanent component of worker productivity h, which is consistent with Chodorow-Reich and Karabarbounis (2016). As in Kehoe et al. (2019), we also allow for the worker opportunity cost to depend on worker productivity z to match the dynamics of labor market flows across the earnings distribution.

Workers that decide not to participate in the labor market, that is, they decide to not search for a new job in the upcoming search-and-matching stage, consume a benefit given by

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

The main difference between the flow benefits of unemployment (9) and nonparticipation (10) is that the latter does not depend on z. This assumption is motivated by the fact that unemployment benefits (imperfectly) scale with workers' prior earnings. Importantly, this assumption helps dampen the cyclicality of labor force participation in the model, consistent with the data (Mukoyama, Patterson, and Şahin, 2018).

[1] An alternative interpretation of our preferences is that agents have distorted (pessimistic) beliefs regarding the distribution of future aggregate shocks $\varepsilon_{A,t+1}$, with the Radon-Nikodym derivative (7) capturing the extent of this behavioral distortion.

Firms and Financial Markets

Firms are owned by infinitely-lived shareholders who collect the output from their existing employees, pay out wages, and post vacancies to attract new workers. The objective of a firm is to maximize the net present value of profits to the shareholders. Shareholders have preferences of the same form as workers (6), but with two differences. First, shareholders are infinitely lived ($\zeta = 0$). Second, shareholders are well-diversified across firms and hence bear no idiosyncratic risk; we therefore set $\gamma = 0$, which, given our preference specification, also implies an infinite elasticity of intertemporal substitution.

Given the above, the present value of a claim to a stream of future cashflows 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 Λ_{t+s} is the one-period shareholders' stochastic discount factor (SDF) between periods $t+s-1$ and $t+s$ given by

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

That is, the stochastic discount factor has the same form as in Lettau and Wachter (2007), with the risk-free rate determined by the shareholders' time discount factor β.

Directed Search and Matching

Workers out of employment decide each period whether to search for a job. Employed workers can also search for a job with a probability that depends on their current state $\chi(\Omega)$. The probability of on-the-job search for a worker with productivity type Ω is

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

Hence, the search pool consists of all workers in unemployment, as well as a random share of employed workers who get an opportunity to search on the job.

Firms can post vacancies directed at workers of a particular productivity type. Labor markets are competitive—all firms can freely enter the labor market for any type of worker in each period. The per-period cost to post a vacancy directed at a worker of productivity Ω is given by

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

As in Meeuwis et al. (2025); Afrouzi, Blanco, Drenik, and Hurst (2024), the cost of posting a vacancy targeting a specific type of worker is increasing in the worker's productivity z, with the parameter $\bar{\kappa}_1 > 0$ determining the elasticity of the vacancy posting cost with respect to z. The assumption that vacancy costs are proportional to A and h ensures that the limiting employment distribution is

not degenerate, while the assumption that they increase with z ensures that job-finding rates are fairly similar across workers with different prior earnings levels, as is the case in the data.

Vacancies are tailored to specific worker types Ω and are also characterized by an expected total lifetime utility value V that the firm promises to the worker. Thus, labor markets are organized in a set of submarkets (Ω, V). Workers that are actively searching for a job choose to direct their search to a submarket with an offered value V for their current productivity type Ω.

The likelihood of a vacancy being filled is a function of market tightness $\theta \equiv \nu/u$, where u is the number of workers searching for a job and ν is the number of posted vacancies by firms in a given submarket. Following den Haan, Ramey, and Watson (2000), the number of matches in a submarket with unemployment rate u and vacancies v is given by

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

This matching function ensures that job-finding and vacancy-filling rates are bounded between zero and one. Specifically, equation (15) implies that the probability that a vacancy is filled in a market with tightness θ is $q(\theta) = (1 + \theta^\alpha)^{-1/\alpha}$ and the probability that a job searcher obtains a new match is $p(\theta) = \theta(1 + \theta^\alpha)^{-1/\alpha}$.

1.2 Labor Market Search and Contracting

Firms partially insure workers by offering dynamic contracts. These contracts define the flow wage for an employed worker given the history of worker productivity which is common knowledge and hence contractible. While the full history of shocks is observable and contractible, worker search decisions on the job are private information and are unobserved by the firm. Workers and firms have limited commitment: neither side can commit ex-ante to continuing an employment match that delivers negative value to either side.

The contract is fully flexible in how wages can respond to aggregate and idiosyncratic shocks, and therefore the pass-through of shocks to wages is endogenously determined by the solution to the optimal wage contracting problem. Following Spear and Srivastava (1987), the state space can be expressed in terms of current promised utility to avoid having to keep track of the full history of shocks. Hence, the optimal labor contract solves a dynamic optimization problem with forward-looking constraints. Thus, each contract delivers a promised utility V that the firm promises to a worker of type Ω.

We next elaborate on the problem faced by the worker and the firm.

Worker Problem

Employed workers choose whether to keep their current match. Job seekers (unemployed or on the job) choose a submarket to search in, while nonparticipants decide whether to begin searching next period.

Continuation decision. First, consider a worker who enters the second stage in an existing match with a firm, and suppose this match is not terminated for exogenous reasons. Let $V_t^O(\Omega)$ be the worker value in nonemployment. Limited worker commitment implies that a worker with continuation value V_t^C of staying in the current match decides to break up the match and leaves for nonemployment if and only if V_t^C is below $V_t^O(\Omega)$.

Worker directed search. Next, consider a worker who is actively searching for a new job during the search-and-matching stage of the current period. Given her current outside option value V, which can represent either the value of her current job or the value of unemployment, the worker directs her search to the submarket with the best offer for her productivity type Ω. That is, she chooses the submarket that maximizes the expected gain R in lifetime value:

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

Thus, the worker targets the job posting that offers the best trade-off between the probability of finding a match and the expected lifetime utility conditional on finding a match. We denote the solution to this search problem by $\mathcal{V}_t^*(\Omega, V)$, with associated job-finding probability $p_t^*(\Omega, V) \equiv p(\theta_t(\Omega, \mathcal{V}_t^*(\Omega, V)))$. If this worker is in an existing match with a firm, 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. Finally, consider a worker of productivity type Ω who ends the current period in nonemployment. This worker faces the choice of whether to enter the next period as a nonparticipant (which yields a continuation value $V_t^N(\Omega)$) or to enter the search pool for the next period as an unemployed worker (obtaining a continuation value $V_t^U(\Omega)$). Thus, her continuation value equals

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

A worker that decides to be out of the labor force collects the non-participation benefit (10) and, conditional on surviving, will be a nonemployed worker in the next period. Her continuation value is therefore equal to

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

An unemployed worker collects unemployment benefits (9) and, conditional on surviving, actively searches for a job during the search-and-matching stage of the next period. Hence, her continuation value is

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

Firm Problem

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

Continuation decision. Let $J_t^{FC}(\Omega, V^C)$ be the net present value of firm profits conditional on keeping the match alive in the second stage, when V^C is the utility that has been promised to the worker. The firm cannot commit to keeping a match intact and therefore terminates the match when $J_t^{FC}(\Omega, V^C)$ falls below zero (the outside option for the firm). Combined with limited worker commitment, we define the following continuation indicator that captures the endogenous firm and worker separation decisions in the second stage:

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

Randomization. In the third stage, the firm chooses the value of lifetime utility V^S that is promised to current workers conditional on them staying with the firm rather than transitioning to a new firm via on-the-job search. The promise-keeping constraint imposes that the total ex-ante value derived by the worker is consistent with what was previously promised. As in Thomas and Worrall (1990), firms randomize over worker continuation utilities in order to implement incentive-compatible allocations while preserving concavity of the firm value function J. In particular, firms choose the randomization probabilities π_j and promised utility V over a two-point lottery,

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

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

Vacancy posting. Next, we consider the creation of new vacancies by firms. The value to a firm of posting a vacancy directed at workers of type Ω and with a promised utility of V is

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

Firms create such a vacancy if and only if the benefit (24) of doing so is (weakly) positive. Due to free entry, the number of vacancies keeps increasing until market tightness is such that the expected profit from each type of vacancy is non-positive:

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

with equality when there is a positive amount of new vacancies created, $\theta_t(\Omega, V) > 0$.

Optimal dynamic contract. Finally, consider a worker who is employed by a firm in the final stage of the period. The worker's current continuation value is V^S. This continuation value reflects the fact that the worker will receives a wage w, which she consumes, and is promised a future state-

contingent lifetime utility $V^{C\prime}$ conditional on continuation of the match until the end of the second stage of the next period. To satisfy promise-keeping, the wage and future continuation values to the worker need to deliver at least

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

The firm's chooses the flow wage w and the state-contingent promised utility $V^{C\prime}$ to maximize the net present value of profits derived from this 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 as an 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). The equilibrium consists of a set of firm and worker value functions, a market tightness function, optimal contract policy functions, a job retention probability function, a search policy function, a distribution of workers over individual states conditional on the aggregate state, and a distribution of vacancies across submarkets. The model admits a block recursive equilibrium.

2 Model Mechanisms

We next discuss the mechanisms that drive the model's key predictions.

2.1 Job Finding and Separation Rates

As in Kehoe et al. (2023); Meeuwis et al. (2025), higher risk premia x reduce vacancy creation and lower job finding rates. Both the separation threshold $z^*(x)$ and the participation threshold $\underline{z}(x)$ increase in x, raising job destruction and lowering labor force participation (Figure A.3). The mechanism is the same as in Meeuwis et al. (2025): employment payoffs are more backloaded than those of nonemployment, so higher discount rates erode the value of a match for the marginal worker. This sensitivity is strongest among low-experience(low λ) workers, who face a lower cost of separation.

The novel element here is worker risk aversion. Figure A.3 compares the baseline thresholds with a $\gamma = 0$ counterfactual in which workers are risk neutral (with perfect household risk sharing). Risk aversion affects thresholds through two channels. First, it raises the *level* of both thresholds: because risk-averse workers with incomplete markets value the partial insurance that employment provides less than risk-neutral workers value employment, fewer workers participate and more matches are dissolved. Second, and more importantly, risk aversion *attenuates* the sensitivity of the thresholds to risk premia. Separation destroys the partial insurance the match provides against both aggregate

and idiosyncratic risk. When x is high, nonemployment is also riskier—expected unemployment durations rise, durations are more uncertain (with long spells inducing faster depreciation of z), and firms can insure newly-hired workers less, especially those near the separation threshold. Together, these forces make separation and nonparticipation more costly for low-productivity workers precisely when risk premia are elevated, dampening the response of z^* and $\underline{z}$ to x relative to the $\gamma = 0$ case.

2.2 Firm Insurance and Wage Dynamics

A central element of our model is that firms partially insure risk-averse workers through optimal dynamic contracts. Two forces shape the resulting wage path: an *insurance motive*—firms smooth wages to shield hand-to-mouth workers from idiosyncratic risk—and a *retention motive*—firms backload compensation to deter on-the-job poaching. We characterize each in turn and then illustrate how wages respond to shocks.

Optimal wage path. Along any path for which the limited commitment constraints are not binding—i.e., $V^{C\prime}_{t+1} > V^{O}_{t+1}$ and $J^{FC}_t(\Omega', V^C) > 0$—the first-order conditions of the firm's problem (27) together with the promise-keeping constraint (26) imply the following first order condition that links spot wages to the worker's continuation value:

$$-w_t^{-\gamma} = \frac{[V_t^S]^{-\gamma}}{\left[\frac{\partial J^{FS}(\Omega_t, V_t^S)}{\partial V_t^S}\right]} = \underbrace{\beta\Gamma_{t+1}[V_{t+1}^C]^{-\gamma}}_{\substack{\text{marginal effect of} \\ V_{t+1}^C \text{ on lifetime utility} \\ \text{(per unit probability)}}} \Bigg/ \underbrace{\left[\Lambda_{t+1}\frac{\partial J^{FC}(\Omega_{t+1}, V_{t+1}^C)}{\partial V_{t+1}^C}\right]}_{\substack{\text{marginal cost to} \\ \text{firm of providing } V_{t+1}^C \\ \text{(per unit probability)}}} = \frac{[V_{t+1}^C]^{-\gamma}}{\left[\frac{\partial J^{FC}(\Omega_{t+1}, V_{t+1}^C)}{\partial V_{t+1}^C}\right]}, \quad (28)$$

which can be rearranged to yield:

$$\frac{w_{t+1}}{w_t} = \left[1 + \left(\frac{V_{t+1}^S}{w_t}\right)^{\gamma} J^{FS}(\Omega_{t+1}, V_{t+1}^S)\frac{\partial \log \tilde{p}_{t+1}(\Omega_{t+1}, V_{t+1}^S)}{\partial V_{t+1}^S}\right]^{1/\gamma}. \quad (29)$$

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 all states in which neither limited commitment constraint binds. The first equality follows from the envelope theorem, and the last equality uses the fact that $\beta\Gamma_{t+1}/\Lambda_{t+1} = 1$. This condition is the hallmark of efficient risk sharing between the firm and the worker, given the constraints of the problem.

Equation (29) characterizes workers' wage growth along the equilibrium path. The second term in the 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 lower the likelihood that the worker receives a better outside offer. To deter poaching, firms optimally backload wages. Without on-the-job search this elasticity is zero and wages are 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* (keep wages stable) and

a *retention motive* (backload wages to deter outside offers). When neither limited commitment constraint binds, the firm fully insures the worker by providing a stable, nondecreasing wage path.

Behavior at the bounds and separation. At the limited commitment bounds, equation (29) no longer holds and wages adjust to satisfy the promise-keeping constraint (26). When the worker's outside option binds ($V^{C\prime}_{t+1} = V^{O}_{t+1}$), the firm raises wages; when the firm's participation constraint binds ($J^{FC}_t(\Omega', V^C) = 0$), the firm lowers wages. If $J^{FC}_t(\Omega, V^O) < 0$, no feasible contract satisfies both constraints and the match is dissolved following (21). Since higher risk premia x reduce match surplus, they raise the separation threshold and bring more workers—particularly low-productivity ones—into the region where the firm's constraint binds. This erodes the scope for insurance precisely when workers are most vulnerable, linking the firm's insurance provision to the aggregate state.

Illustration: response to shocks. Figure 7 illustrates these forces by plotting how (log)continuation values (top panels) and wages (middle panels) respond to shocks.

The left panels show the response to human capital shocks h. Since separation risk and on-the-job search do not depend on h, the insurance motive dominates. For small shocks, the firm fully insures the worker: wages stay flat while continuation values adjust within the non-binding region. Once h' is sufficiently far from h, one of the constraints binds and wages respond one-for-one thereafter. At higher levels of z where on-the-job search rates are higher, the retention motive also plays a role: wages respond to h even inside the bounds to deter outside offers (via the last term in (29)).

The right panels show the response to productivity shocks z, which additionally affect separation and the possibility of an outside-offer. On the downside, small negative shocks are absorbed through lower continuation values; wages fall only once the firm's limited commitment constraint binds. If z crosses the separation threshold, there is no feasible wage contract $J^{FC}_t(\Omega, V^O) < 0$ and the match dissolves. On the upside, the retention motive is now active compared to the case of h shocks: since the outside-offer probability rises with z (see (13)), the firm preemptively raises wages and continuation values to deter poaching. The reason is that, for high z workers, the more relevant outside option is successful on-the-job search rather than separation into nonemployment. Hence, the firm keeps continuation values near—but not at—the upper bound (Figure 7b). This is because all new matches start with positive surplus, placing an upper limit on the worker's gains from on-the-job search. Once compensation is high enough that poaching is ineffective, only the insurance motive remains, and keeping the constraint slack preserves the firm's ability to insure the worker against future adverse shocks.

3 Mapping the Model to the Data

We now turn to the question of how we take the model to the data. Since all decisions in the model are bilateral—between a single firm and a single worker—the boundary of the firm is indeterminate

without further assumptions. However, we can interpret firm-level shocks as a common component in the idiosyncratic productivity shocks (z and h) experienced by coworkers at the same firm. Under this interpretation, the model's predictions about how the pass-through of worker productivity shocks to earnings varies across workers and with aggregate conditions (Section 2.2) translate directly into testable predictions about the pass-through of firm-level shocks to worker earnings.

Two predictions are worth highlighting. First, the model implies that the overall pass-through of firm shocks to worker earnings is shaped by two competing forces that operate at opposite ends of the earnings distribution. At the top, high-productivity workers have better on-the-job search opportunities (equation (13)), so the optimal contract tilts toward retention incentives rather than insurance, making stayers' earnings more responsive to shocks. At the bottom, workers near the separation threshold face high rates of endogenous job destruction, which makes their total earnings—including nonemployment spells—highly sensitive to firm shocks through the extensive margin. For workers in the middle, neither force is dominant: the retention motive is weak and limited commitment constraints rarely bind. The model therefore predicts a U-shaped pattern in overall pass-through, with the sources differing across the distribution—driven by the contract at the top and by separations at the bottom. Second, for lower-paid workers, the firm optimally insulates wages from shocks as long as limited commitment constraints are slack. However, these constraints bind more frequently when risk premia are elevated, so the model predicts that the pass-through of firm shocks to low-wage workers' earnings should be countercyclical—rising when discount rates are high, primarily through increased separation risk.

We document empirical evidence consistent with both predictions. The cross-sectional patterns—established in Meeuwis et al. (2025) and Friedrich et al. (2019)—display a U-shaped pass-through profile that matches the model's prediction, with high exposure at both tails driven by distinct mechanisms. In addition, we provide new evidence that the pass-through of firm-specific TFP shocks to lower-paid workers' earnings is indeed countercyclical. We then describe the calibration of the model's parameters.

3.1 Data and Methodology

Our empirical analysis builds on Meeuwis et al. (2025). We use a 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 tracks the subsequent earnings of incumbent U.S. workers in public firms; Appendix A.1 provides full details. The main outcome variable 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; the denominator adjusts for the life-cycle earnings profile. Workers are included if employed by a Compustat firm in year t and are followed regardless of subsequent employment status, so that the growth measure captures both wage changes and nonemployment spells.

We link worker earnings to two sources of shocks. Firm productivity growth $\epsilon^{tfp}_{f,t+1}$ is measured as annual revenue-based TFP following İmrohoroğlu and Tüzel (2014) (Appendix A.2). To measure variation in risk premia, we follow Meeuwis et al. (2025) and extract 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)—that capture fluctuations in risk or investors' risk-bearing capacity. We denote the resulting series of risk premium shocks as ϵ^{rp}; this common factor explains 60% of total variance across the nine series. Since earnings are annual, we cumulate monthly shocks from mid-year t to mid-year $t+1$ to construct annual risk premium shocks ϵ^{rp}_{t+1}. Appendix A.3 provides further details.

3.2 Cross-Sectional Heterogeneity in Pass-Through to Workers

Following Meeuwis et al. (2025), we estimate the pass-through of firm productivity and risk premium shocks to worker earnings, allowing coefficients to vary by the worker's prior earnings rank within her 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 the employer of worker i at time t and $\mathbb{1}_{\tau(i,t)=s}$ is an indicator for the worker's prior within-firm earnings bin. Controls include a polynomial in lagged earnings, industry $\times$ earnings group fixed effects, and worker industry $\times$ age $\times$ gender fixed effects; standard errors are clustered by worker and year. Appendix A.4 provides further details on the specification.

We briefly summarize the key cross-sectional patterns, which closely mirror Meeuwis et al. (2025) and Friedrich et al. (2019). 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 rises sharply to 15–20 percent for workers in the top few percentiles (Figure 1a). Adverse firm shocks also raise nonemployment risk, particularly among low-rank workers, for whom firm downturns translate primarily into separations rather than wage cuts (Figure 1b). Taken together, the total exposure of worker earnings to firm TFP shocks—combining wage adjustments and job destruction—displays the U-shaped pattern predicted by the model, with the sources differing across the distribution: the high pass-through at the top reflects the retention motive in the optimal contract, while the high exposure at the bottom operates primarily through the extensive margin. The pass-through of risk premium shocks is U-shaped at short horizons—high for both low- and

high-wage workers, low in the middle—and nearly monotonically decreasing at longer horizons as the cumulative exposure of low-wage workers rises further (Figure 2). The strong exposure at the bottom of the earnings distribution primarily reflects job destruction rather than wage declines (Figure 2b).

3.3 Time Variation in Pass-Through to Workers

The cross-sectional patterns above are well established (Meeuwis et al., 2025). We now turn to the model's second prediction: that the pass-through of firm shocks to lower-paid workers' earnings should be countercyclical (see also Figure A.4b). A new finding in this paper is that this prediction is borne out in the data: the pass-through of *firm-specific* shocks varies systematically over time, and this time variation is linked to aggregate risk premia. To document this, we first estimate 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 $\times$ earnings group $\times$ year fixed effects. Figure 3 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.

To explore this relationship more formally, we 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. These interaction terms imply that higher risk premia are associated with a greater overall pass-through of firm-specific shocks to worker earnings, 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).

To understand the source of these earnings losses, we examine the role of job separations. 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.

Finally, Table A.1 estimates 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, confirming that the amplification operates primarily through the extensive margin.

In sum, the degree of firm insurance—as measured by the pass-through of firm-specific shocks to worker earnings—varies systematically across both time and workers. Higher risk premia weaken insurance, especially for lower-paid workers, primarily through higher rates of job destruction. This countercyclical pattern in firm insurance confirms the model's prediction that limited commitment constraints tighten when discount rates rise, and provides the key empirical target for our calibration.

3.4 Calibration

Having established these empirical patterns, we now calibrate the model. We first formalize the mapping from the model's worker-level shocks to firm-level shocks previewed above, and then calibrate the remaining parameters in two steps. Table 3 summarizes our baseline calibration.

Parameters Calibrated a Priori

We first calibrate a subset of parameters using a priori information, listed in Panel A of Table 3. 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, consistent with 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 Asset and Labor Markets

The second step of our calibration selects the remaining parameters to jointly match moments from asset and labor markets. We target 91 empirical moments using 20 parameters, whose values are reported in Panel B of Table 3. The target moments and associated parameters are grouped into four

categories, discussed below. While equilibrium outcomes reflect the joint interaction of all parameters, this structure helps clarify which parameters are most directly tied to each set of moments.

Asset markets. Following Meeuwis et al. (2025), 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 $\beta = 0.999$ 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.

Labor market dynamics. We discipline the 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 Sahin (2017) (see Appendix A.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 A.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, $g_O = 0.03\%$ per year, shapes time variation in endogenous separation and job-finding rates by affecting the duration of the match surplus. 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. We discuss the labor market dynamics implied by the calibrated model and their fit to the targeted moments in Section 4.1.

Volatility of firm TFP growth and worker earnings growth. Since production is linear, the value of each worker–firm match can be evaluated in isolation. To connect the model to our empirical analysis—which focuses on the pass-through of firm-level TFP shocks to worker earnings—we introduce a model-based notion of firm-level shocks by assuming that a component of individual productivity shocks is common to all workers currently employed by the same firm. This assumption

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 size of firm-level shocks 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) and set $\sigma_h = 2.7\%$.

Pass-through regressions. Finally, we discipline the remaining parameters by targeting the empirical pass-through regression coefficients. Specifically, we match the model-implied regression coefficients of worker earnings growth and the probability of nonemployment on 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 $\lambda_H = 2.4$ 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 such as the model matches the pass-through of risk premium shocks to worker earnings. Section 4.2 discusses the resulting model's fit to the empirical pass-through regression estimates.

3.5 Quantitative Model Properties

We now characterize the quantitative properties of the calibrated model.

Pass-through of model shocks. The model features several sources of worker heterogeneity: earnings are driven by idiosyncratic shocks to z, h, and λ. Figure A.4 summarizes how each feeds into worker earnings. Shocks to z have the biggest effect (panel a), both because they have the largest volatility and because they directly affect the probability of separation (panel c). The earnings pass-through of z-shocks exhibits a U shape across the earnings distribution, while the impact on separation risk is concentrated among lower-paid workers. Pass-through is also state dependent: when risk premia x rise, the sensitivity of earnings to z also rises (panel b), especially for low-wage workers. Conditional on staying employed, however, the pass-through of z shocks increases with workers' earnings levels—a pattern driven by the retention motive for high-productivity stayers. The second most important source of earnings variation is permanent human capital shocks, which mainly affect higher-paid workers but do not influence separations. Aggregate shocks also pass through to earnings, but their volatility is small and they account for only a limited share of overall earnings variation.

Value of firm insurance. How valuable is the insurance the firm provides? Figure 7e compares the present value of worker's compensation under the optimal contract (valued at the firm's SDF) with a piece-rate alternative that keeps allocations unchanged but restricts wages to be proportional to productivity. Negative values mean the optimal contract is cheaper to the firm. Indeed, this is the case across most of the z distribution: insuring workers reduces the NPV of compensation to the firm by more than 10 pp for the least productive workers and around 5 pp at $z = 1$. These gains reverse in the right tail, where the retention motive outweighs the insurance motive: to retain its most productive workers, the firm provides strong, nonlinear, backloaded incentives, making stayers' earnings respond sharply to z shocks at high levels of z (see Figure A.4e). The value of insurance is largely insensitive to discount rates, because the gains from trade between worker and firm stem from idiosyncratic rather than aggregate risk. For very high-z workers, the value is lower when x is high: low job-finding rates reduce the risk of on-the-job poaching, so the retention premium shrinks.

Impulse responses to aggregate shock. Figure 8 shows the impulse response of key model variables to the aggregate shock ε_A, which drives both risk premia (Figure 8a) and aggregate TFP (Figure 8b). Following a negative shock, output falls (Figure 8c) and discount rates rise, which lowers employment and raises unemployment (Figure 8d). As a consequence, worker earnings decline, especially for low earners (Figure 8e). This decline in earnings is driven both by the increased

probability of separation and by wage declines for stayers (Figure 8f). Left-tail income risk rises persistently (Figure 8g), while right-tail income risk slightly falls in the short run and increases afterwards (Figure 8h).

4 Model Fit

We now evaluate the model's ability to replicate key features of the data.

4.1 Heterogeneous Labor Market Dynamics

Panel A of Table 4 shows that the model replicates the volatility, persistence, and cyclicality of key labor market indicators. The unemployment rate mean of 6.5% and volatility of 1.4% match that in the data, and the model reproduces the high degree of persistence observed empirically. Labor market tightness is volatile and strongly procyclical, and the employment-to-population ratio moves closely with the cycle. The model somewhat overstates the volatility of labor force participation, reflecting the absence of non-economic motives for non-participation.

Panel B shows that job-finding and separation rates in the model vary over the business cycle with realistic magnitudes: the job-finding rate falls in recessions, whereas the separation rate into unemployment rises. Consequently, fluctuations in the aggregate unemployment rate are driven both by a procyclical job-finding margin and a countercyclical separation margin.

Panel C decomposes unemployment fluctuations following Shimer (2005, 2012) by computing counterfactual unemployment rate series that hold either separations or job finding constant over the cycle. In both the data and the model, most of the variation in the unemployment rate arises from movements in job finding rather than separations. Hence, the model captures the dominant role of job creation in aggregate unemployment dynamics, while also producing realistic and strongly countercyclical fluctuations in separations.

Figure A.1 shows that the model also captures the heterogeneity in job-finding and separation rates across workers. In the data, job-finding rates are as good as homogeneous across the earnings distribution. In the model, average job-finding rates are somewhat increasing in prior income, but their cyclicality is very similar across groups. Both in the data and in the model, the level and cyclicality of separation rates are substantially higher for low-wage workers.

4.2 Pass-Through of Productivity and Risk Premium Shocks

The calibrated model also replicates the cross-sectional heterogeneity in worker exposures to firm productivity and financial shocks documented in Section 3. We estimate the same regression specification (33) using model-simulated data and compare the resulting coefficients to their empirical counterparts.

Figure 4 compares the estimated coefficients of worker earnings growth across different horizons

on firm TFP growth, risk premium shocks, and their interaction, separately for the five income groups. Panel (a) shows that, as in the data, the unconditional pass-through of firm TFP growth to earnings increases with worker income, consistent with the retention motive driving higher wage sensitivity among top earners. Panel (b) shows that the model also reproduces workers' exposures to risk premium shocks closely, as in Meeuwis et al. (2025). Panel (c) illustrates that the interaction between firm TFP growth and risk premium shocks is positive for low-wage workers but close to zero for high-wage workers. Even for lower-paid workers, the interaction effects are modest at short horizons but strengthen and are consistent with the data at medium-term horizons.

Figure 5 presents the same set of coefficients, but with the outcome variable replaced by an indicator for having a zero-earnings quarter as a measure of job destruction. Adverse TFP shocks raise the probability of job loss, particularly for lower-paid workers (panel a), reproducing the other side of the U-shaped total exposure documented in Section 3.2: while the earnings pass-through rises with income, the separation channel is concentrated at the bottom. Increases in risk premia also significantly raise the probability of nonemployment, disproportionately among low-income workers—a pattern that the model matches quantitatively (panel b). Finally, panel (c) shows that, as for earnings growth, the amplification of firm TFP shocks when risk premia rise is modest at short horizons but aligns with the data over longer horizons.

In Appendix Figure A.2, we replicate the analysis from Table A.1, estimating exposures separately for stayers and movers. The model overstates the unconditional response to firm productivity shocks for low-wage movers, but it closely matches the empirical pattern that wage earnings respond much more strongly to risk premium shocks and to the interaction between firm productivity growth and risk premium shocks for movers than for stayers, especially among lower-paid workers.

4.3 Model-Implied Fluctuations

We next examine whether the model can replicate realized fluctuations when fed with actual shocks. Specifically, we take our empirical measure of risk premium shocks ϵ^{rp}_{t+1} from Section 3.1 as proxies for the model's aggregate shocks $\varepsilon_{A,t+1}$, which jointly drive aggregate productivity and financial conditions. We then compute the implied time series for labor market variables. Because the model is scale invariant, these variables do not depend on the realized path for aggregate TFP A.

Figure 6 plots the model-implied series. Fed with the empirical risk premium shock series, the model produces aggregate labor market dynamics that closely resemble those observed in the data. The model reproduces the cyclical dynamics of the unemployment rate, with a correlation of 67% between the model-implied and observed series. Periods of elevated risk premia—such as during the early 2000s and the 2008–09 financial crisis—coincide with pronounced increases in unemployment and slow recovery thereafter.

The simulated job-finding and separation rates also exhibit realistic magnitudes and cyclical

properties. Job finding and labor market tightness (V/U) are strongly procyclical, while separations are countercyclical, consistent with the evidence in Table 4. Importantly, the path of labor market tightness (V/U) in the model closely tracks that in the data, accounting for the high sensitivity of vacancy creation to aggregate conditions emphasized by Shimer (2005). While model-implied total employment also strongly comoves with the data, the model overstates the volatility of the employment-to-population ratio, as noted in Section 4.1. Overall, the realized paths indicate that fluctuations in risk premia account for a large share of observed labor market variation.

We also estimate the model analog of equation (32) on the realized paths, yielding a measure of the pass-through of firm-level productivity shocks for each period. We run these regressions of three-year earnings growth separately for workers in the bottom and top income groups and compare the results to their empirical counterparts in Figures 6g (low income) and 6h (high income). The series are demeaned to remove the baseline level of the coefficients. For low-income workers, the model-implied pass-through coefficients closely track both risk premia and the empirical pass-through estimates. The correlation between the two series is high (53%), although the fluctuations over time are somewhat larger in the data. For high-income workers, by contrast, both the data and the model display little systematic cyclicality in pass-through.

5 Model Implications

Having shown that the calibrated model fits the moments used in calibration—including labor market dynamics, asset prices, and the heterogeneous pass-through of firm shocks to worker earnings—we now explore its broader implications. We first examine whether the model can replicate salient features of the distribution of individual earnings growth that were not targeted in the calibration. We then quantify the welfare costs of idiosyncratic risk, the value of human capital, and the effects of state-contingent labor market and financial policies.

5.1 Worker Earnings Risk

We start with the model's implications regarding the distribution of labor earnings growth.

Overall Distribution of Earnings Growth

Using administrative data from the U.S. Social Security Administration, Guvenen et al. (2021) show that the distribution of individual earnings growth is far from normal. It is highly leptokurtic, with a sharp peak at zero and very fat tails. Roughly one-third of workers experience almost no earnings change in a given year, yet large increases and declines occur far more often than under a Gaussian distribution. The log density declines approximately linearly in both tails, consistent with a Pareto distribution at both ends. Since the left tail is thicker, the earnings growth distribution is negatively skewed.

Examining Figure 9a, we see that our model matches these key features even though individual productivity growth is normally distributed.[2] Our model replicates the distribution of earnings growth remarkably well, including the heavy tails that deviate sharply from a Gaussian distribution: the kurtosis of annual earnings growth is 11.3 in the model versus 14.9 in the data.

The excess tail risk in worker earnings arises endogenously from the joint dynamics of labor markets and wage contracts. The model generates non-Gaussian earnings growth because earnings are a nonlinear function of underlying shocks. When wage contracts can be smoothly adjusted within the limited-commitment bounds, earnings barely move; when a constraint binds or the match is destroyed, wages are reset and earnings can change sharply, including through job loss. This endogenous "kinked" response produces a large mass near zero and fat tails even with Gaussian productivity growth. Because labor market transition rates and the degree of wage smoothing differ across aggregate and worker states, some shocks are well insulated while others lead to very large changes in earnings. This logic is related to the job-ladder explanation in Hubmer (2018): infrequent transitions between employment states generate large and asymmetric earnings changes. Our framework adds firm insurance through dynamic contracts and links the frequency of these nonlinear events to aggregate financial conditions.

Earnings Risk across Workers

A key finding of Guvenen et al. (2021) is that earnings risk varies systematically across the worker earnings distribution: lower-income workers face a more symmetric distribution of earnings growth with substantially higher variance, while higher-income workers experience smaller but more asymmetric and leptokurtic growth rates. Figure 9b assesses the model's ability to replicate this cross-sectional pattern. The model quantitatively matches the declining pattern of earnings volatility from the bottom up to about the 90th percentile, but it does not reproduce the sharp rise in volatility at the very top. It also matches the fact that earnings growth for job switchers is much more volatile than for job stayers, particularly for low-income workers (Figure 9c). Appendix Figures A.5b–A.5e show how skewness and kurtosis vary across workers; the model qualitatively matches the hump-shaped pattern of both moments, though not with the same magnitudes as in the data.

In the model, these cross-sectional patterns reflect heterogeneity in exposure along both the extensive and intensive margins, as illustrated in Figure A.4. 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 endogenously separate 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

[2]For ease of visualizing the tails of the distribution, Figure 9a plots the log of the density (on such a scale, a normal density is a quadratic). Appendix Figure A.5a plots the density itself.

wage growth volatility among stayers is higher for high-wage workers, who benefit less from wage smoothing.

Guvenen et al. (2021) also study the persistence of earnings changes across workers. For low-income individuals, negative earnings shocks tend to be short-lived while positive shocks are highly persistent; among high-income workers, this pattern is reversed. Appendix Figure A.6 shows that in our model, the impulse responses of earnings closely resemble those for low-income workers in the data: there is strong mean reversion with respect to negative earnings growth, but not with respect to positive growth. Unlike in the data, the model generates little heterogeneity in this pattern across the earnings distribution—the shape of the impulse responses is broadly similar for all worker groups.

Lifetime Earnings and Employment across Workers

Figures 9d and 9e replicate the analysis of lifetime earnings growth and employment rates in Guvenen et al. (2021). Figure 9d plots cumulative earnings growth between ages 25 and 55 against workers' lifetime earnings percentiles. The upward slope is partly mechanical—faster growth implies higher overall earnings—but the degree of heterogeneity and similar to the data: average earnings grow by 154% at the median, 635% at the top 5%, and 1024% at the top 1%. Figure 9e plots the distribution of years employed over the life cycle. As in the data, the model generates substantial dispersion in lifetime employment rates, although it 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 key moment in this literature that the model matches reasonably well. Together, these results show how short-run tail events cumulate into large lifetime heterogeneity, without overstating the cross-sectional variance of earnings.

Earnings Risk over Time

Guvenen et al. (2014) document how the distribution of individual earnings growth evolves over the business cycle. They show that cyclical variation arises primarily from changes in higher-order moments rather than the second moment: the variance of idiosyncratic shocks is only weakly countercyclical, whereas skewness is strongly so. During recessions, the left tail of the distribution expands and the right tail contracts—large negative earnings shocks become much more frequent, while large positive shocks become rarer—producing strongly countercyclical skewness even as the overall dispersion of earnings changes moves little.

We next assess whether the model can replicate these cyclical dynamics of labor income risk by comparing the realized paths of earnings risk in the model and in the data. We use our empirical series of risk premium shocks as a direct proxy for risk premium shocks in the model, as described in Section 4.3. The top panel of Figure 10 plots the difference between the median and the 10th percentile of earnings growth, capturing the evolution of left-tail risk; the bottom panel plots the difference between the 90th percentile and the median, capturing the right tail. In both the model

and the data, periods of economic downturn coincide with a widening left tail and a narrowing right tail of the earnings growth distribution. Quantitatively, the model tracks the cyclical fluctuations in these moments well, with correlations of 59% for the left tail and 37% for the right tail.

Costs of Job Loss and Entering the Labor Market in a Recession

A large literature documents that job displacement is associated with substantial and persistent earnings losses. Davis and Von Wachter (2011) estimate that men displaced in mass layoffs lose, in present value, roughly 1.6 years of pre-displacement earnings when displacement occurs during expansions and about 2.5 years when it occurs during recessions—highlighting that the cost of job loss is strongly cyclical. We ask whether our model can replicate these findings.

Following the empirical design of Davis and Von Wachter (2011), we restrict attention to incumbent workers below age 50 with at least three years of tenure at their current employer. Given their findings that lower-tenure and lower-wage workers are much more likely to be displaced in mass layoffs (see also Farber, 2017; Margolis and Montana, 2024, for related evidence), we also restrict the sample to workers with pre-displacement earnings below the median in these comparisons. In the model, a worker is classified as 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 compute average earnings around the displacement event and compare them to the average earnings of otherwise similar workers who are not displaced.

Figure 11 shows that displacement induces large and highly persistent earnings losses, with a pronounced deterioration in the first few years after job loss and only partial recovery even at long horizons. These losses are driven by 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 when it occurs in expansions.

The model also predicts that entering the labor market during a recession leads to persistent earnings losses. 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 12 shows that entering in times of high unemployment has a substantial initial effect on earnings; the effect fades gradually but persists for up to 10 years. The pattern is consistent with the evidence in Kahn (2010) and Schwandt and von Wachter (2019), who document sizable and long-lived earnings shortfalls for cohorts entering the labor market in recessions. Schwandt and von Wachter (2019) also document that these costs are particularly high for workers who enter with lower levels of education, a result that aligns with our model's prediction that these costs are larger for workers with lower levels of initial productivity z_0.

5.2 Welfare Cost of Idiosyncratic Risk

We evaluate welfare in our model by converting worker values into certainty-equivalent consumption units. Recall that the value function $V_{i,t}$ in (6) denotes the expected lifetime utility of worker i at time t. The certainty-equivalent consumption $CE_{i,t}$ is the constant consumption level that yields the same lifetime utility as the stochastic consumption stream implied by the model, given the current aggregate and individual state. This measure is obtained directly from the value function as

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

To analyze how idiosyncratic risk affects welfare, we compute certainty-equivalent consumption under two valuations of the same model outcomes. First, we evaluate utility under the baseline model with the calibrated risk aversion parameter $\gamma = 0.48$. Second, we re-evaluate welfare under the exact same consumption dynamics but with $\gamma = 0$ to compute counterfactual value functions. Thus, the two economies share the same labor and consumption dynamics, and only differ in how workers value risks. When $\gamma = 0$, workers are risk neutral with respect to idiosyncratic shocks. This lets us interpret the gap between the two certainty equivalents as a measure of the welfare cost of idiosyncratic risk.

Under the veil of ignorance, overall welfare is the expectation of welfare for newborn workers over the unconditional distribution of the aggregate state x and the idiosyncratic initial productivity level z_0. Certainty-equivalent consumption in the baseline valuation is 23 percent lower than in the counterfactual. This implies that workers would require a permanent increase in consumption of roughly 30 percent 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 realistic earnings dynamics and the assumption that workers are hand-to-mouth.

Figure 13a plots welfare in the baseline model relative to the counterfactual as a function of initial productivity z_0. For each z_0, welfare is integrated over the unconditional distribution of x. Welfare losses are sizable but fairly uniform across the distribution—20 to 25 percent—with the biggest losses for workers around the middle and at the top of the distribution. Since hand-to-mouth behavior is empirically plausible for lower-income households but less so at the top of the distribution, true heterogeneity in welfare losses from idiosyncratic risk is likely substantially larger once worker savings are taken into account, which is beyond the scope of this paper.

5.3 Value of Human Capital

We next quantify the value of human capital in the model. Unlike the welfare measure above—which aggregates the entire stream of flow utility including nonemployment payoffs—the value of human capital focuses exclusively on the present value of future wages earned while employed. For each worker, we compute this object under two valuations applied to the same equilibrium dynamics of the baseline model: a baseline valuation using the actual risk aversion parameter $\gamma > 0$, and

a counterfactual valuation with $\gamma = 0$. As in the welfare analysis, the difference reflects how risk aversion affects the valuation of labor income risk, not the underlying income dynamics.

In the bottom panels of Figure 13, we plot the annualized ratio of human capital to current consumption—an analog of the price-earnings ratio for equity—under these two different valuation methods. On average, the human capital to consumption ratio is around 15 in the baseline model. This value is substantially lower than estimates in the literature. For instance, Lustig, Van Nieuwerburgh, and Verdelhan (2013) report an average wealth–consumption ratio of 83 for U.S. households. The lower valuation ratios in our model reflect two forces. First, human capital wealth is exposed to sizable aggregate wage and employment risks, which are discounted by the aggregate SDF. Second, idiosyncratic earnings and employment risks increase households' effective discounting: valuation ratios in the counterfactual with $\gamma = 0$ are nearly twice as large as in the baseline around the middle of the distribution.

The distributional patterns across workers reveal substantial heterogeneity. In Figure 13b, low-productivity workers have low human capital values despite strong mean reversion in productivity. This arises because these workers face a high probability of nonemployment and remain so for extended periods, which depresses the present value of future wage income. Figure 13c shows a different pattern for incumbent workers sorted by current wage. Under $\gamma = 0$, the human capital to consumption ratio declines steeply with wage, reflecting expected mean reversion in wages. However, after taking into account the cost of idiosyncratic risk, these ratios compress markedly. Low-wage workers are exposed to the largest idiosyncratic risks, so their expected future wages are discounted most heavily. As a result, the model implies that total wealth behaves much more like equity than previously thought, even for low-wage workers.

We also compute human capital discount rates, defined as the internal rate of return: we first 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. Figure 14 plots annualized discount rates in the cross section and over time. The discount rate under the baseline valuation is substantially higher than in the counterfactual with $\gamma = 0$, reflecting idiosyncratic risk that workers cannot diversify. The additional discount is highest for low-wage incumbent workers, who face the largest idiosyncratic risk. Figure 14c shows that the aggregate human capital discount rate fluctuates substantially over the business cycle and peaks following the Great Recession.

5.4 Policies

Finally, we use the calibrated model to evaluate the positive and normative implications of state-contingent labor market and financial policies.

Labor market policies. We start by analyzing two state-contingent labor market policies that mimic policies that are often introduced in recessions. Each policy is triggered only when the economy

is in a recession, and 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 state-contingent policies in the model based on the level of risk premia x_t, the aggregate state variable that drives business cycles in our framework. Specifically, we define the recession indicator $\mathbb{1}_t^R$ that marks periods in which labor markets are slack as

$$\mathbb{1}_t^R = \begin{cases} 1 & \text{if } \log x_t > \log \bar{x} + \sigma_x^{unc}\, \Phi^{-1}(1-\delta) \\ 0 & \text{otherwise,} \end{cases} \tag{38}$$

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

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

$$b_t(\Omega) = \left(1 + \xi_b\, \mathbb{1}_t^R\right) A_t\, h\left(\bar{b}_0 + \bar{b}_1\, z\right), \tag{39}$$

so that in recessions the worker receives a fraction ξ_b more than the baseline benefit. We set ξ_b so that the flow payoff in unemployment rises by 20% when the policy is active.

The second policy 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 lower the cost of employment in downturns—for example, a temporary payroll tax cut or a wage subsidy. We set ξ_y to 6%, roughly corresponding to the employer share of Social Security contributions that might be paused in a payroll tax holiday.

Table 5 reports the required tax rate τ (1.2–3.5% across specifications) and the effects on ex-ante welfare and ex-post labor market outcomes for each policy in isolation and for the two policies combined, under two recession thresholds $\delta = 10\%$ and $\delta = 20\%$, which bracket the empirical frequency of NBER recession months (13% over 1960–2019). The unemployment benefit subsidy provides additional insurance in states where job-finding rates are low and workers' marginal utility is high, but it weakens employment incentives. As a result, it reduces aggregate employment and output. In contrast, the output subsidy raises the value of employment to the firm and therefore strengthens job creation and retention. It is less effective at providing insurance, however, because the incidence of the subsidy is skewed to more productive workers who are more likely to be employed.

Both policies raise ex-ante welfare despite the balanced-budget requirement, because workers value the state-contingent transfers more than firms discount the tax. Welfare gains are larger when intervention is rarer (lower δ): the unemployment subsidy yields a 1.1% welfare gain at $\delta = 10\%$ but only 0.7% at $\delta = 20\%$, while the employment subsidy yields 0.4% and 0.1% respectively. The effects on labor market quantities are larger for higher δ, as the policy is active in more periods. Combining the two policies yields the largest welfare gain (1.4% at $\delta = 10\%$, 0.8% at $\delta = 20\%$), but employment and output fall relative to the baseline because the disincentives from the unemployment subsidy dominate the incentives from the output subsidy.

Financial policies. We next consider a financial policy that caps the level of risk premia, in the spirit of a "Fed Put": the authority intervenes when risk premia exceed a threshold, so that the effective price of risk is $\tilde{x}_t = \min\{x_t, x^H\}$. The cap x^H is set so that the policy binds in the worst δ fraction of periods (when x_t is in the upper tail),

$$\tilde{x}_t = \min\{x_t, x^H\}, \qquad \log x^H = \log \bar{x} + \sigma_x^{unc}\, \Phi^{-1}(1-\delta). \tag{40}$$

We abstract from the (unknown) implementation costs of the policy and present the results as a counterfactual in which risk premia are capped at no direct resource cost; the welfare and quantity effects therefore only reflect the benefits of lower discount rates in bad states.

Table 6 distinguishes two scenarios. Columns (1)–(3) report outcomes when the policy is fully anticipated: agents know that x_t will always be capped at x^H in the future, so labor allocations and wage dynamics respond to the reduced likelihood of extreme risk premia. The threshold δ is set to 0.1%, 0.25%, and 0.5%—so the cap binds only in the worst 0.1–0.5% of periods—yet average GDP rises by 5.7–11.6%, employment increases with 5.0–10.3 ppt, and the unemployment rate falls by 2.0–3.2 ppt. These large effects indicate that extreme outcomes in risk premia are a first-order driver of equilibrium valuations and labor market outcomes. Since part of the welfare gain from capping risk premia is due to the direct effect of lower risk premia on workers' utility, we report welfare gains relative to the welfare that is obtained under the baseline model dynamics but where welfare is evaluated under the new Γ_t that is driven by $\tilde{x}_t$ instead of x_t. Even after this adjustment, the policy yields large welfare gains of 8.3%, 17.1%, and 22.4% respectively.

For comparison, columns (4)–(6) report outcomes when the policy is unanticipated: we compute welfare and quantities under the baseline equilibrium, so that agents never anticipate the intervention, and ex post cap x_t at x^H in each period. Even with substantially higher values of δ, welfare gains are much smaller, and the effects on GDP, employment, and unemployment are modest. The comparison makes clear that the bulk of the Fed Put's benefit in this model comes from anticipation: when agents know that risk premia will not spike beyond x^H, they value jobs and human capital more and firms create more vacancies and retain more workers, with large effects even when the cap binds very rarely.

Conclusion

This paper develops a model of endogenous firm insurance in frictional labor markets and shows that the scope for insurance depends on aggregate financial conditions. The model features directed search, optimal dynamic wage contracts with two-sided limited commitment, risk-averse hand-to-mouth workers, on-the-job search, and time-varying risk premia. Two forces shape the optimal contract: an insurance motive that smooths wages to shield workers from idiosyncratic risk, and a retention motive that backloads compensation to deter poaching. The balance between these forces is state-dependent. When risk premia rise, the value of employment matches falls, limited

commitment constraints bind for more workers, and the scope for insurance shrinks—especially for lower-paid workers near the margin of job destruction. At the top of the earnings distribution, the retention motive dominates, making high earners' wages responsive to shocks even in normal times. The model therefore predicts a U-shaped cross-sectional pattern in the pass-through of firm shocks to worker earnings, with the sources differing across the distribution, and a countercyclical time-series pattern concentrated among low-wage workers.

We document new empirical evidence consistent with both predictions: using U.S. administrative data, we show that the pass-through of firm-specific TFP shocks to worker earnings rises when risk premia are elevated, with this amplification concentrated at the bottom of the earnings distribution and operating primarily through job destruction.

The calibrated model matches its targeted moments—asset prices, labor market flows, and the heterogeneous pass-through of firm shocks—and generates a rich set of untargeted implications. The distribution of individual earnings growth features the fat tails, excess kurtosis, and negative skewness documented by Guvenen et al. (2021), emerging endogenously from the kinked response of the optimal contract. The model reproduces the countercyclical left-tail risk in earnings growth (Guvenen et al., 2014), the large and persistent earnings losses from job displacement (Davis and Von Wachter, 2011), and the scarring effects of entering the labor market during recessions (Schwandt and von Wachter, 2019)—none of which are targeted in the calibration.

The framework highlights financial conditions as a central driver of worker earnings dynamics and labor income risk. The large, time-varying, and undiversifiable idiosyncratic risk generated by the model implies sizable welfare losses for workers—roughly 30 percent in consumption-equivalent terms—and substantially depresses the valuation of human capital, particularly for low-wage workers. State-contingent labor market policies that provide additional insurance in recessions are welfare-improving, and a financial policy that limits extreme risk premia has large positive effects on equilibrium employment and output, with the bulk of the gains arising from anticipation effects.

Our model abstracts from a job ladder in which workers climb toward better-paying firms or those offering job stability over time. Incorporating such a ladder, as in Jarosch (2023) and Hubmer (2018), would provide an additional channel through which financial conditions affect worker earnings: recessions would not only increase job destruction but also disrupt the process of moving up the job ladder, amplifying the costs of downturns for lower-paid workers. This interaction would be particularly strong if firms influence the rate of human capital accumulation as in Acabbi et al. (2026). We leave this extension for future work.

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.

Azariadis, C. (1975). Implicit contracts and underemployment equilibria. *Journal of Political Economy 83*(6), 1183–1202.

Baily, M. N. (1974). Wages and employment under uncertain demand. *Review of Economic Studies 41*(1), 37–50.

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.

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.

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.

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 M. Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory 113*(1), 1–31.

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.

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 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.

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.

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.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Sahin (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.

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.

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).

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.

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.

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

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.

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.

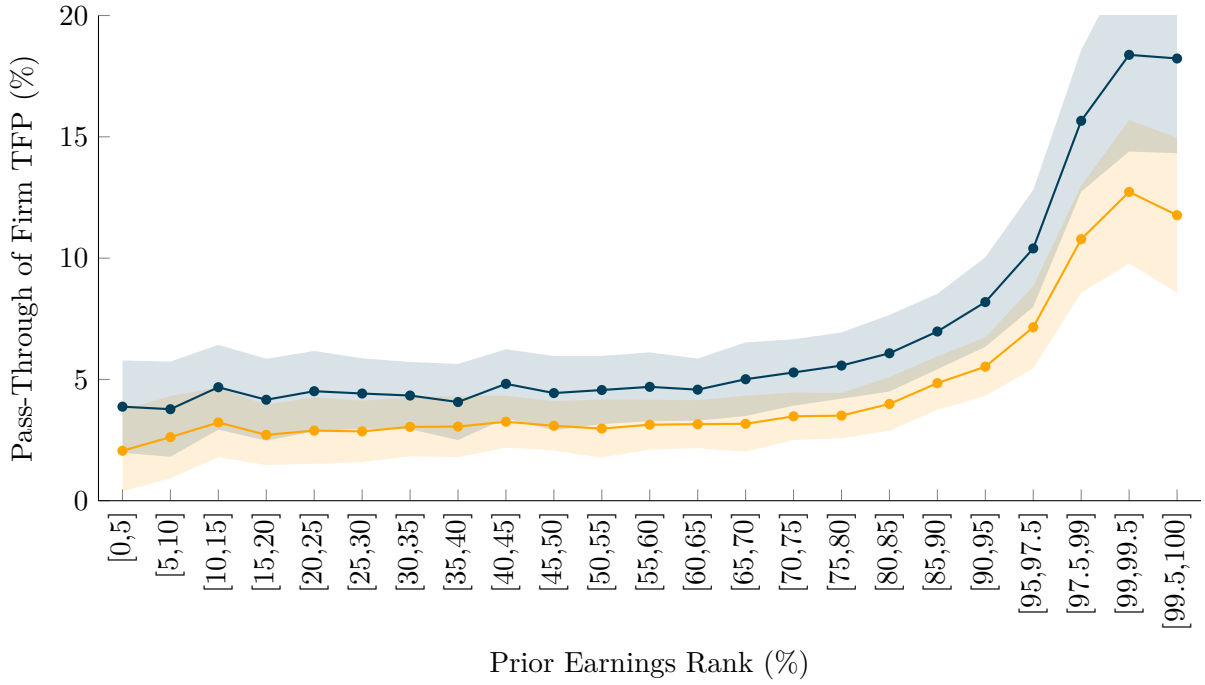
Skiadas, C. (2013). Scale-invariant uncertainty-averse preferences and source-dependent constant relative risk aversion. *Theoretical Economics 8*(1), 59–93.

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: Pass-Through of Firm TFP Growth to Workers

(a) Wage Earnings



(b) Extensive Margin (Probability of Zero-Earnings Quarter)

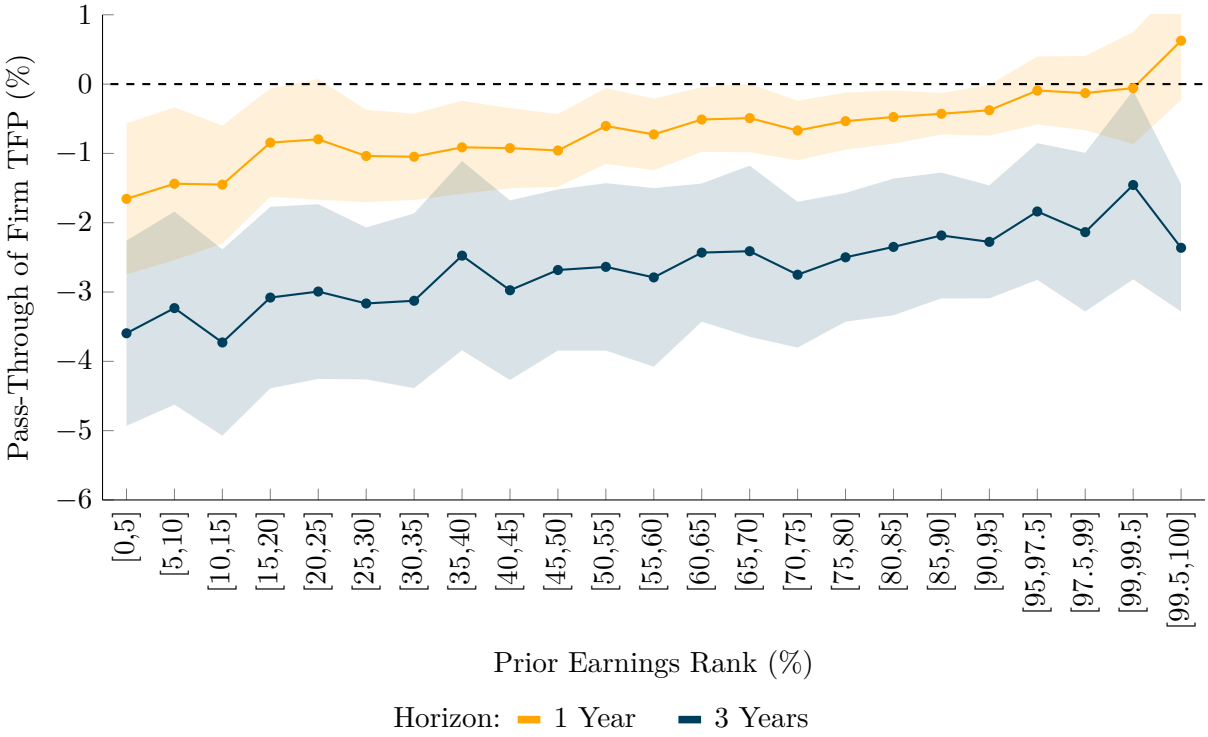
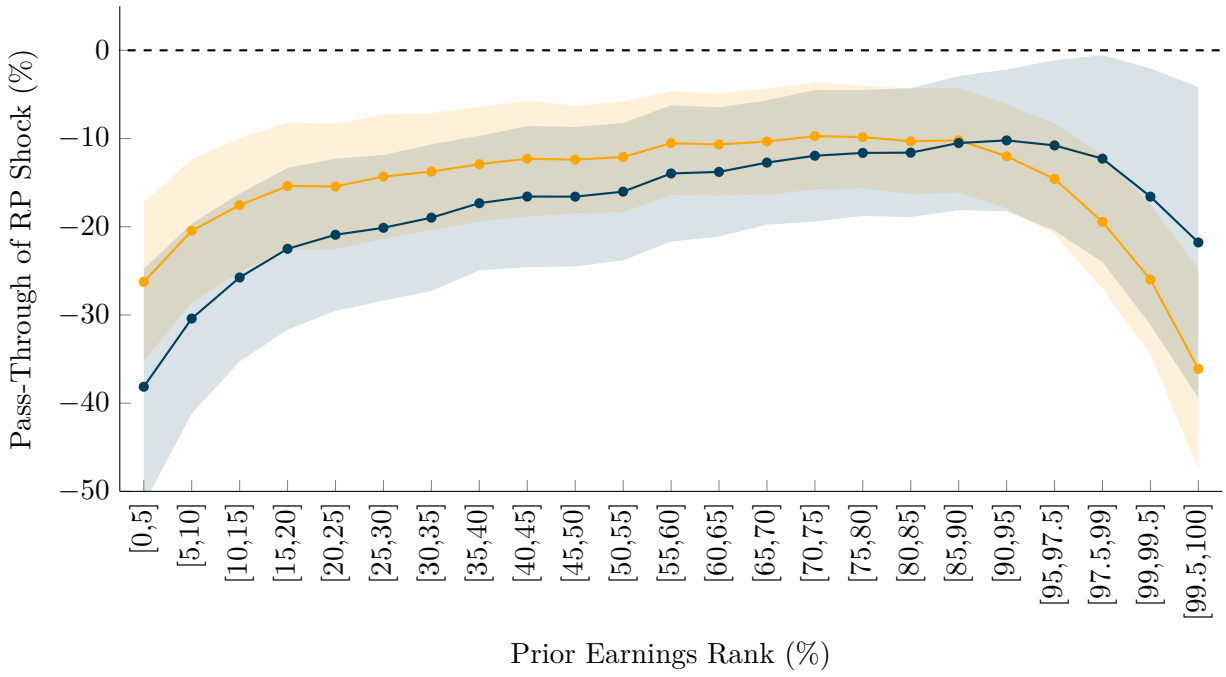


Figure 2: Pass-Through of Risk Premium Shocks to Workers

(a) Wage Earnings



(b) Extensive Margin (Probability of Zero-Earnings Quarter)

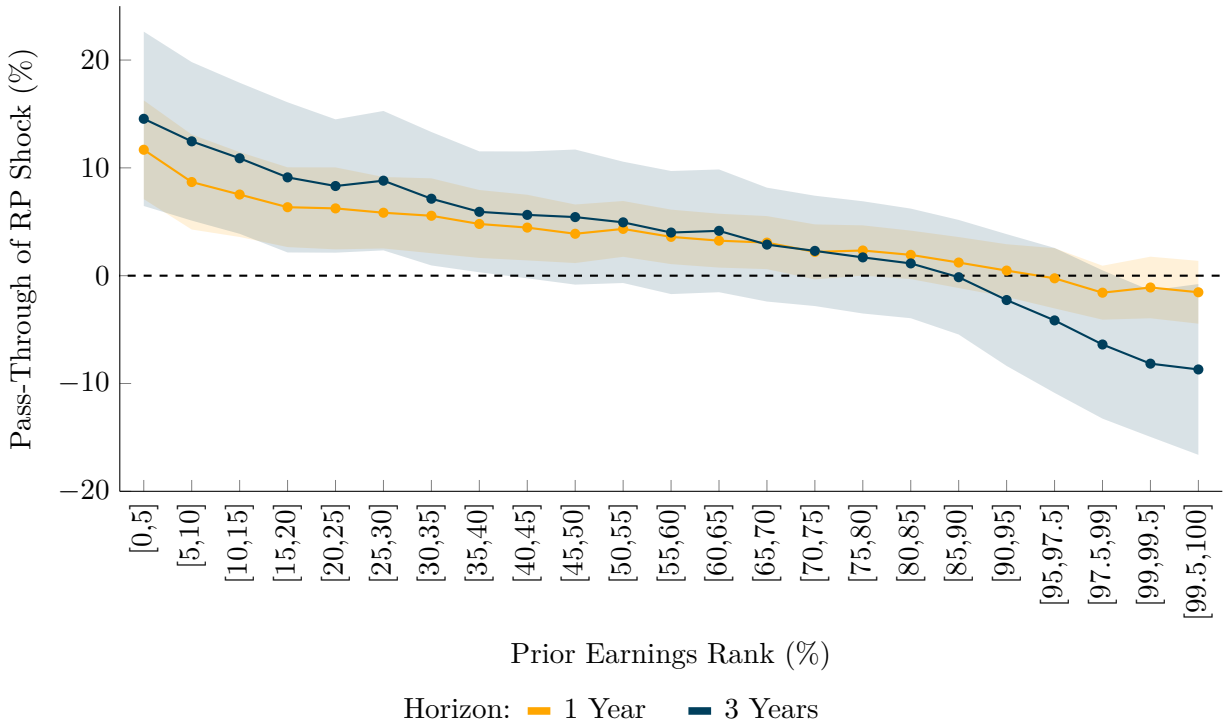


Figure 3: Time-Varying Pass-Through of Firm TFP Growth

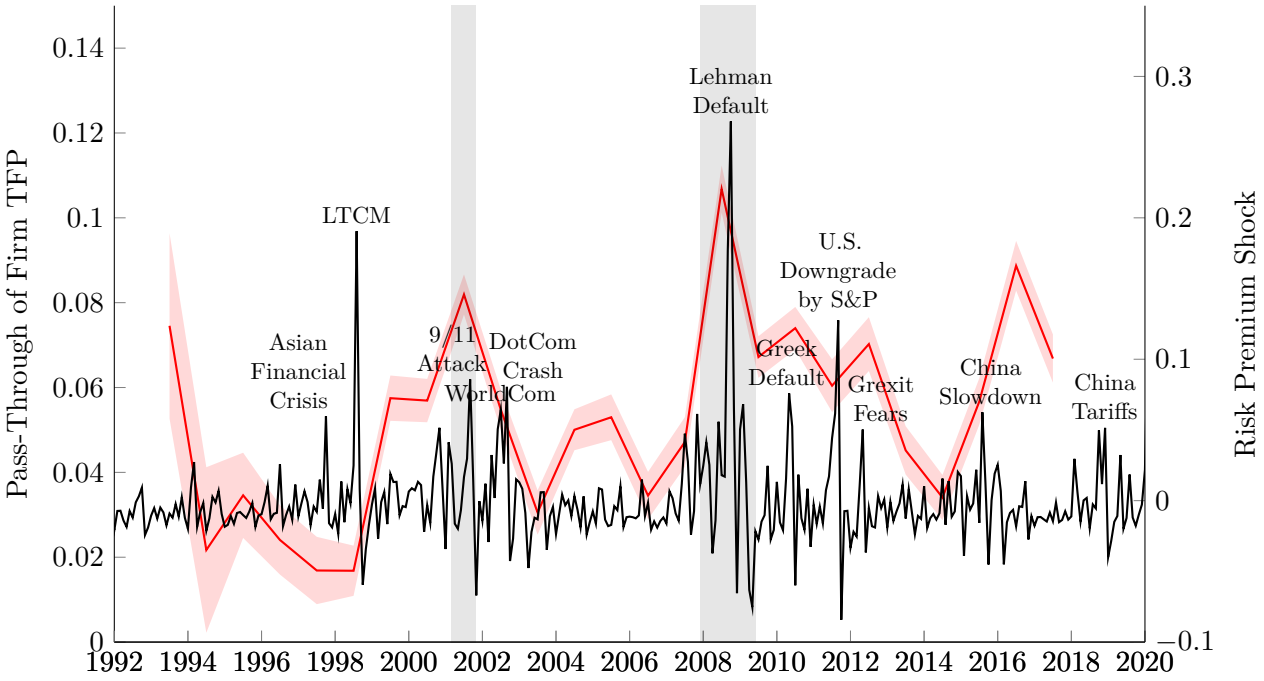
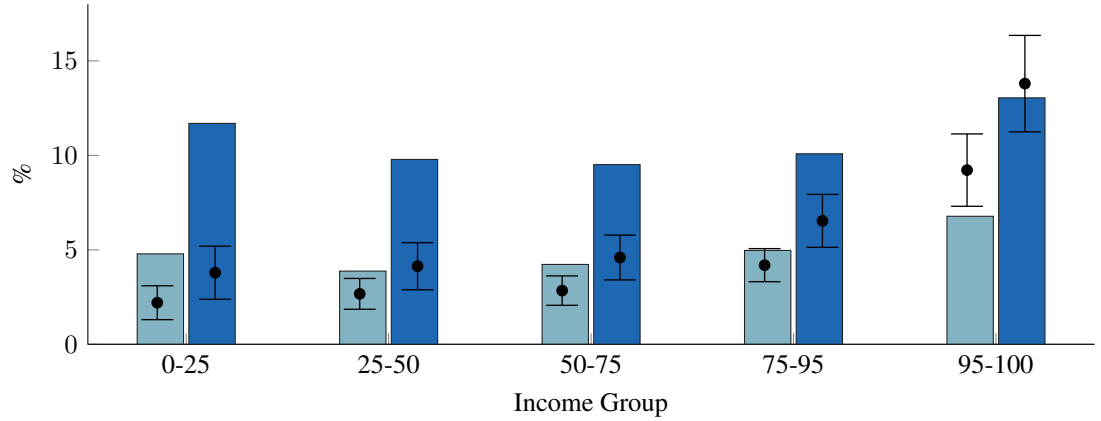
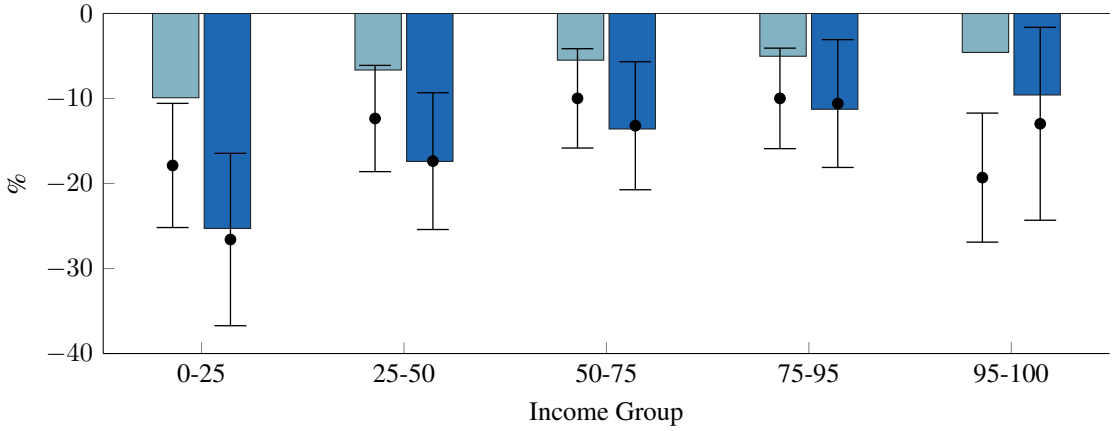


Figure 4: Response of Wage Earnings: Model vs. Data

(a) Firm TFP Growth



(b) Risk Premium Shocks



(c) Firm TFP Growth × Risk Premium Shocks

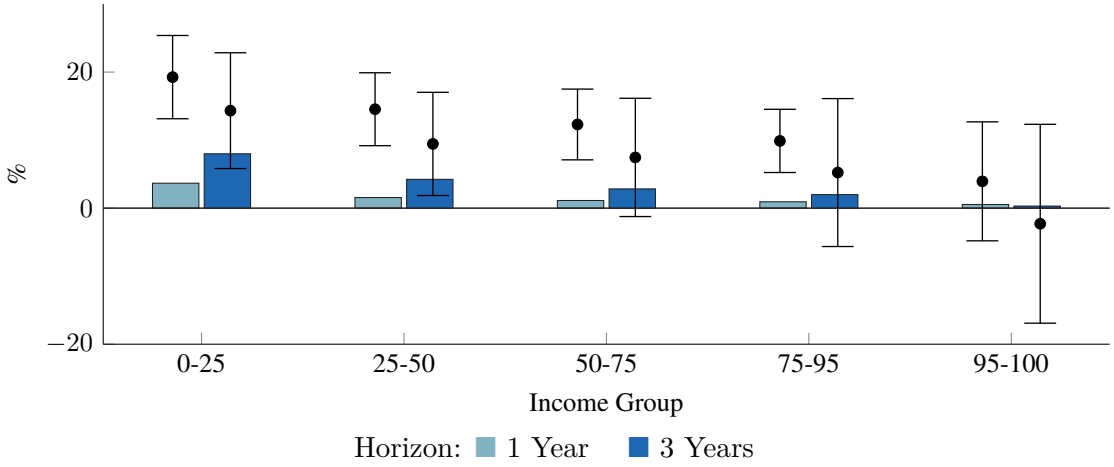
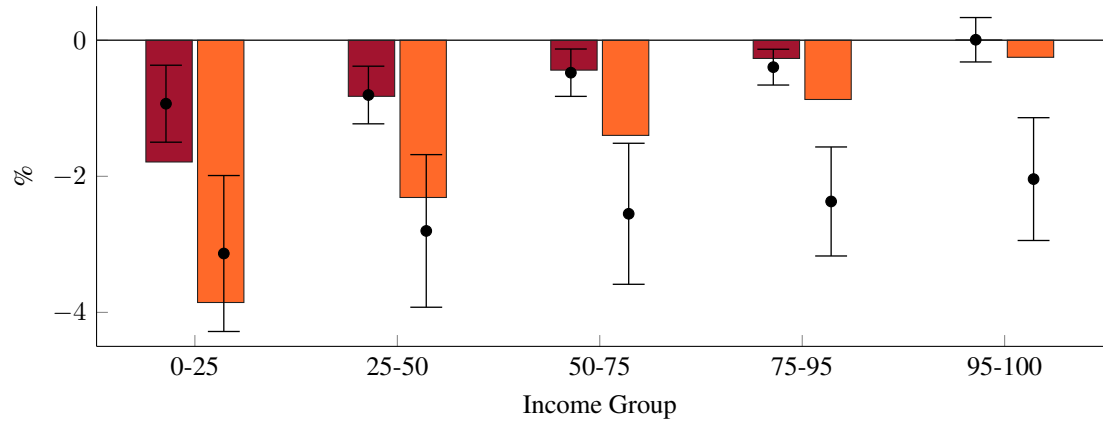
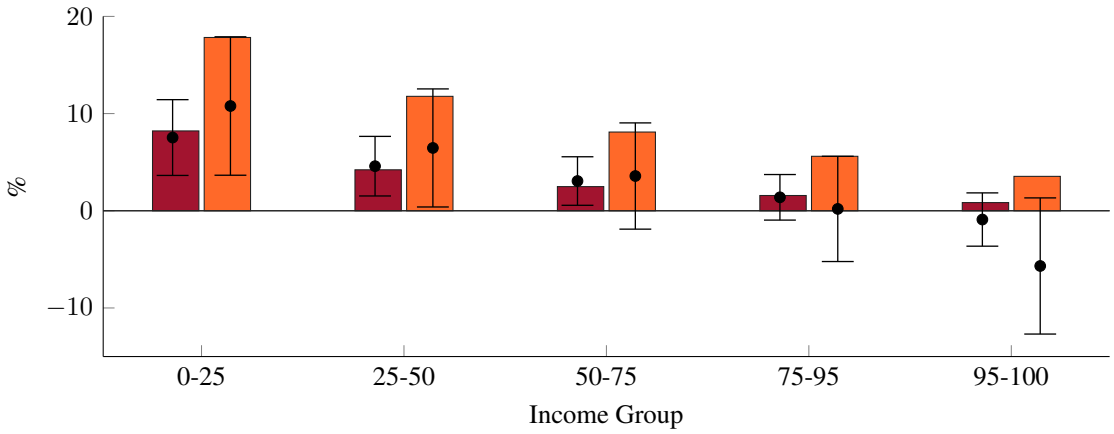


Figure 5: Response of Job Destruction: Model vs. Data

(a) Firm TFP Growth



(b) Risk Premium Shocks



(c) Firm TFP Growth × Risk Premium Shocks

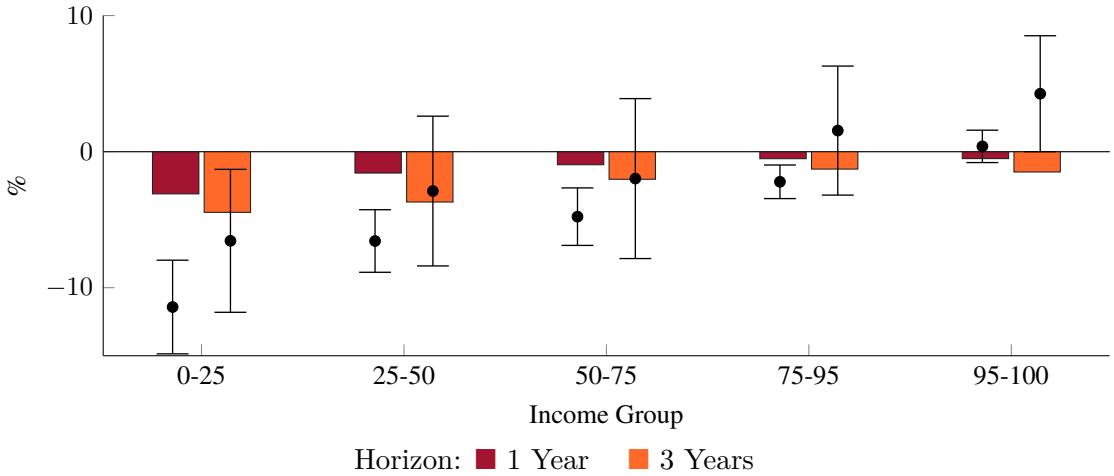
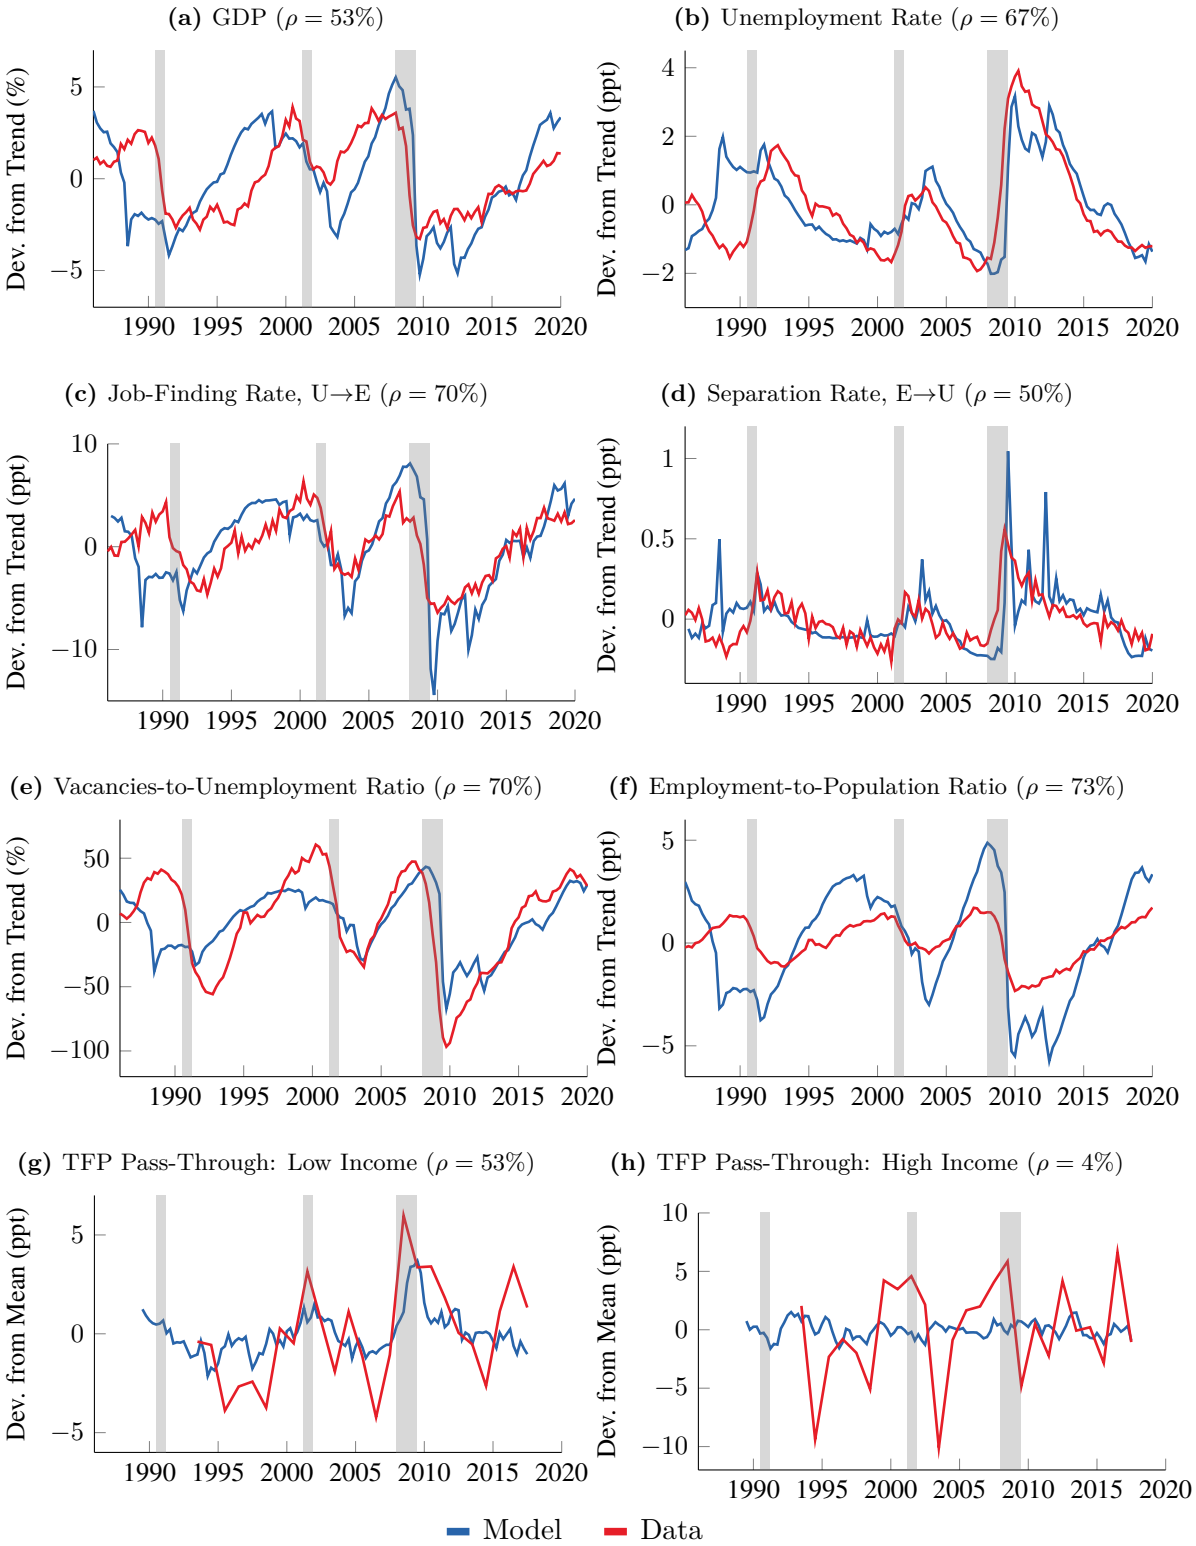
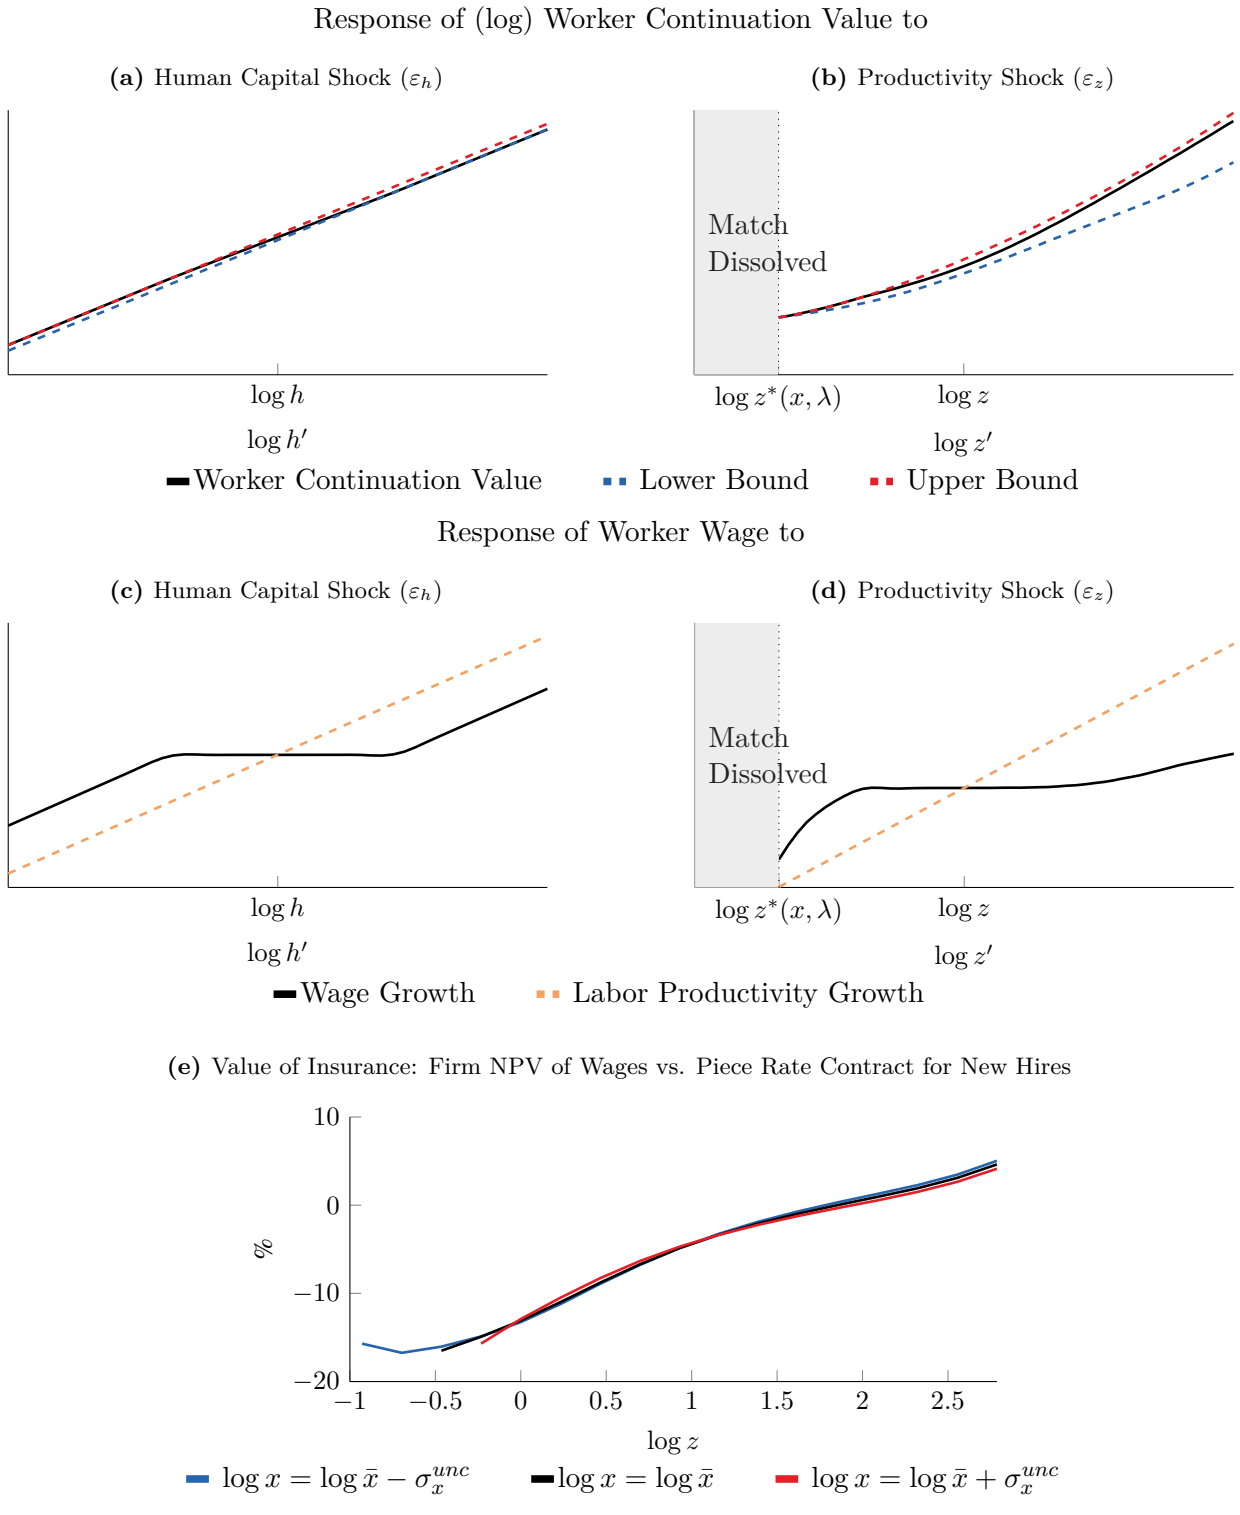


Figure 6: Realized Fluctuations in Labor Markets: Model vs. Data



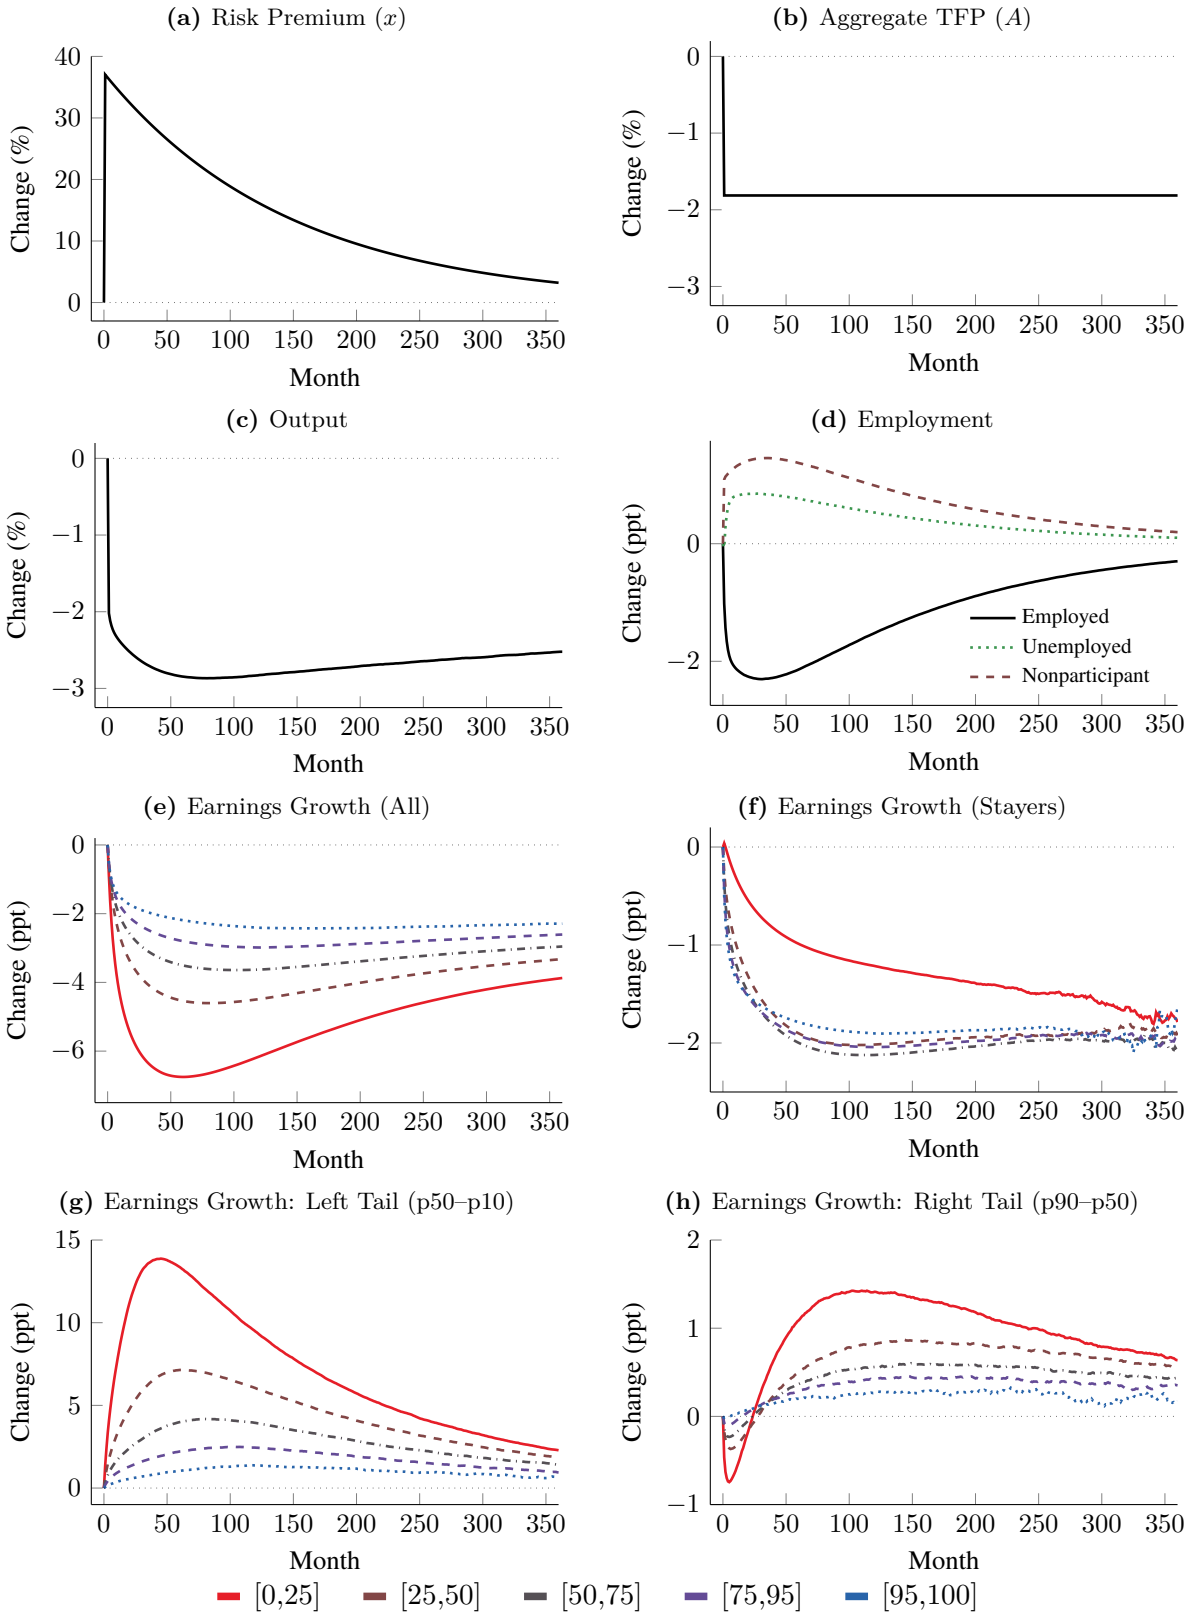
This figure compares the realized paths of key variables between the model and the data. We directly feed into the model our (scaled) empirical measures of risk premium shocks ϵ^{rp}. We detrend all series using an HP filter with quarterly smoothing parameter 10^5.

Figure 7: Wage Dynamics in Model



Top panels: worker continuation value (y-axis) against subsequent idiosyncratic productivity level (x-axis) with z in the left panel and h in the right panel, with limited commitment upper and lower bounds shown as dashed lines. Middle panels: wage growth (y-axis) against shock size (x-axis). Bottom panel: NPV of wages under the optimal contract relative to a piece-rate contract (y-axis) against worker productivity z (x-axis), for different levels of aggregate risk premium.

Figure 8: 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 9: Labor Income Risk: Model vs. Data

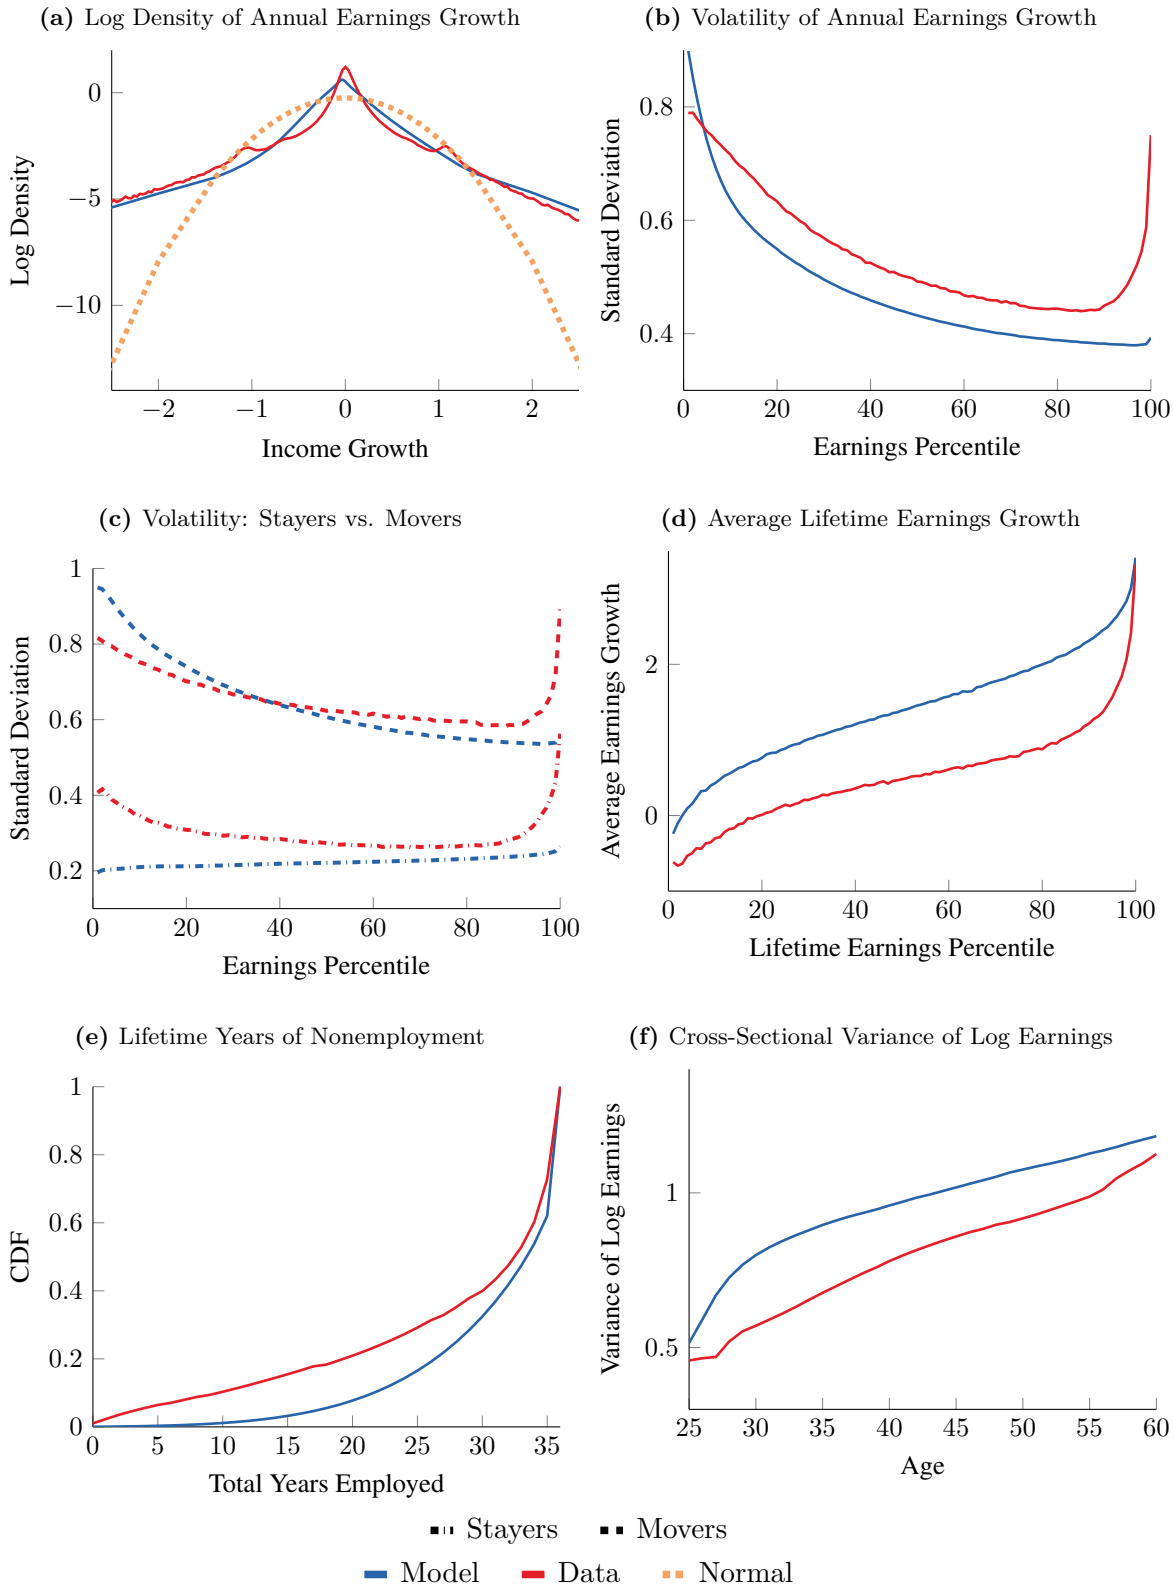
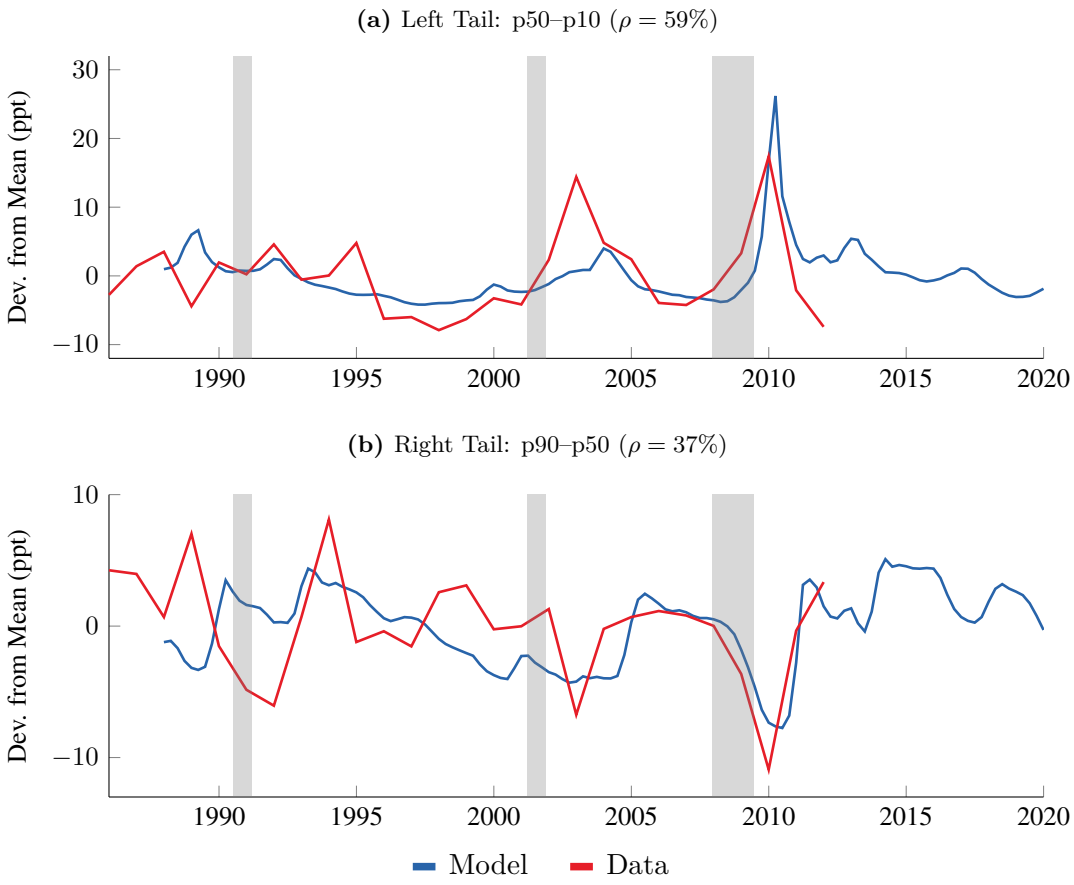


Figure 10: Realized Fluctuations in Earnings Risk: Model vs. Data



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}.

Figure 11: Earnings Losses Due to Job Displacement

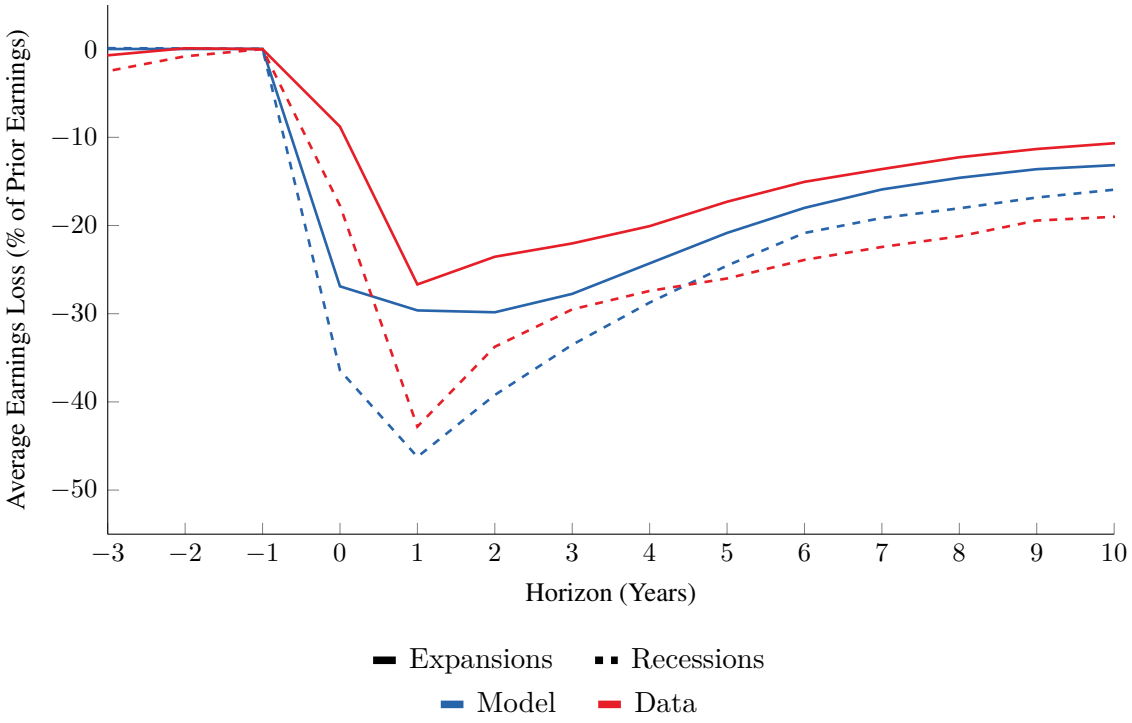


Figure 12: Earnings Losses from Entering the Labor Market in a Recession

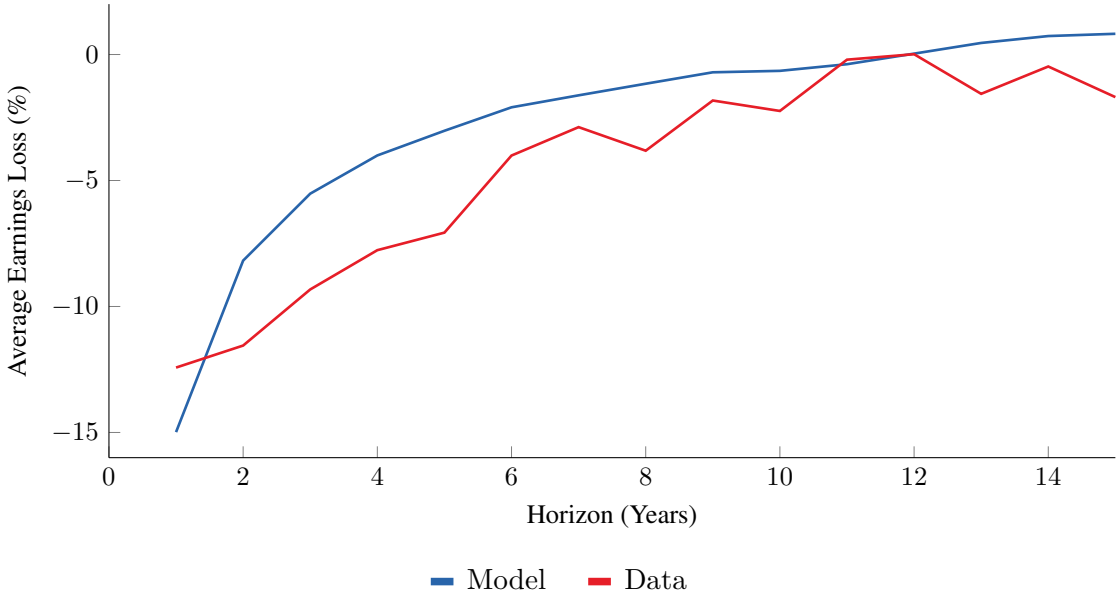
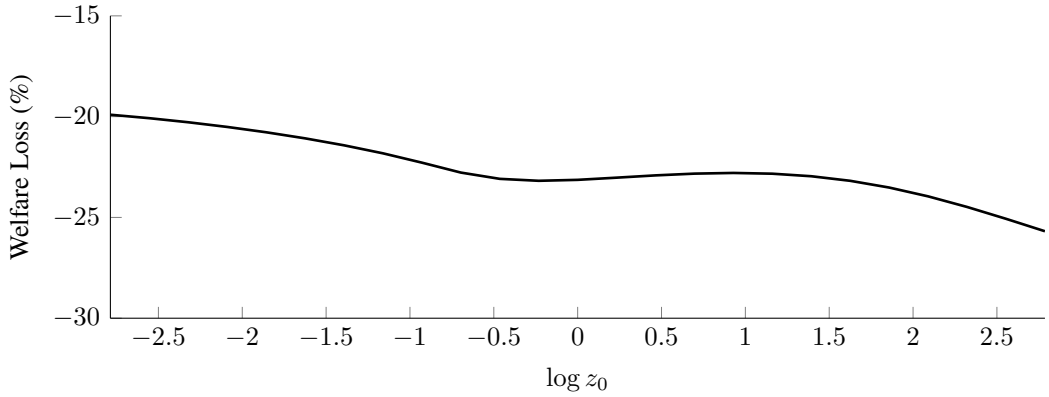
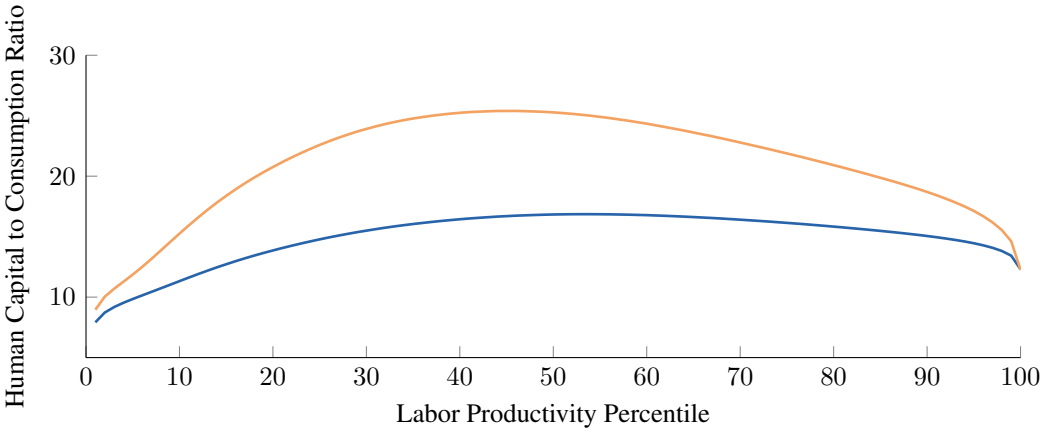


Figure 13: Welfare and Value of Human Capital in Model

(a) Welfare Loss Due to Idiosyncratic Risk ($\gamma > 0$ vs. $\gamma = 0$)



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



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

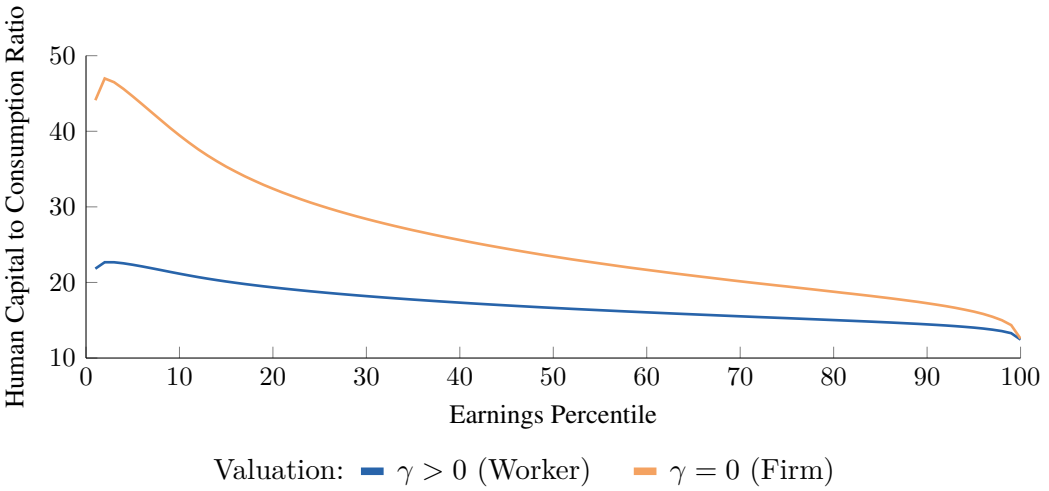
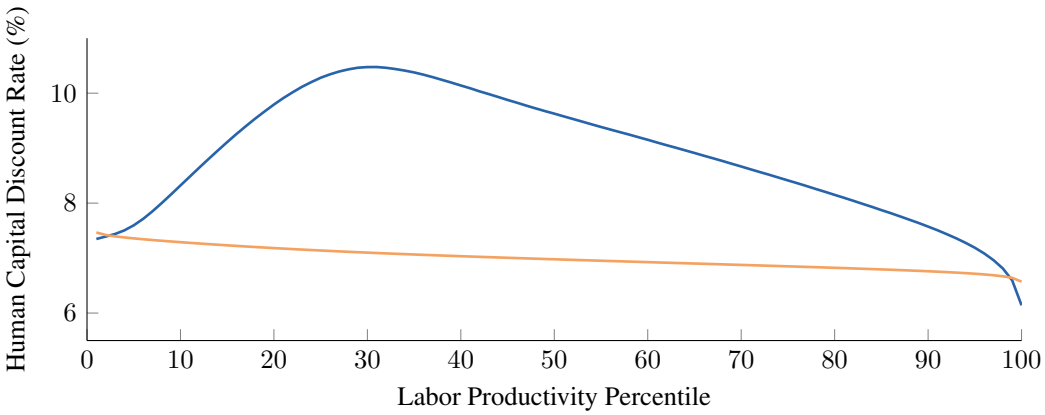
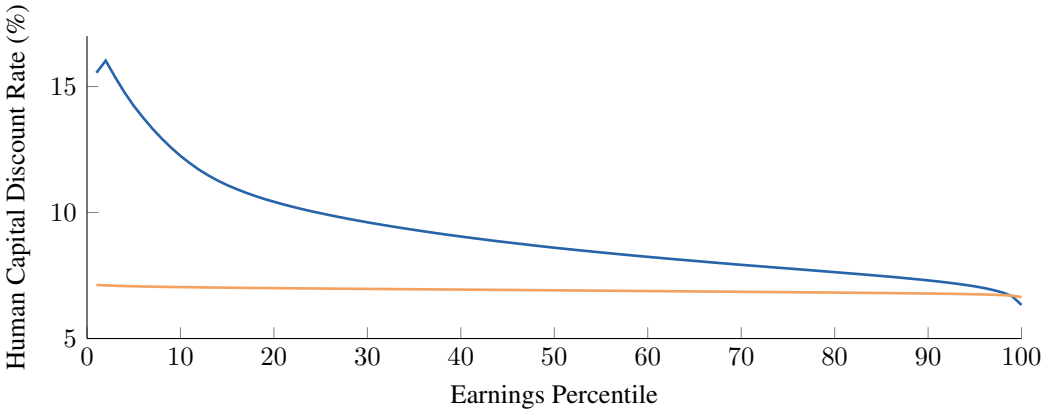


Figure 14: Human Capital Discount Rate

(a) Average Human Capital Discount Rate by Labor Productivity (All Workers)



(b) Average Human Capital Discount Rate by Earnings (Incumbent Workers)



(c) Human Capital Discount Rate: Realized Fluctuations

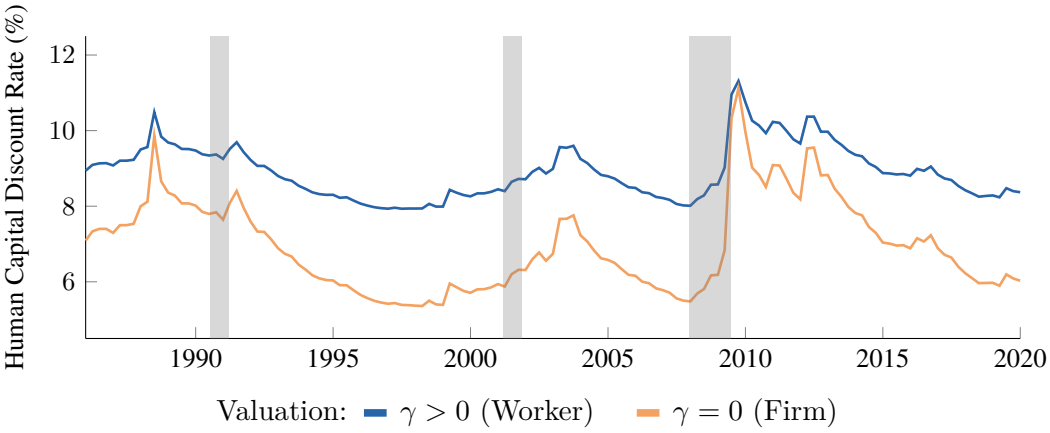


Table 1: Heterogeneity in Pass-Through of Firm TFP Growth

	(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

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

Table 3: 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
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	λ_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	λ_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 4: 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 5: 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, %)	1.4	2.0	1.2	1.8	2.4	3.5
Δ welfare (%)	1.1	0.7	0.4	0.1	1.4	0.8
Δ average GDP (%)	-1.9	-4.0	0.9	1.3	-0.7	-2.0
Δ average consumption (%)	-1.7	-2.6	-0.6	-0.4	-1.9	-2.4
Δ employment-population ratio mean (ppt)	-1.7	-3.5	0.9	1.3	-0.6	-1.9
Δ employment-population ratio vol (ppt)	1.4	2.1	-0.3	-0.4	0.9	1.4
Δ unemployment rate mean (ppt)	3.2	7.2	-0.6	-0.8	2.4	5.9
Δ unemployment rate vol (ppt)	4.6	6.1	-0.4	-0.5	4.0	5.3
Δ income growth vol (ppt)	1.5	2.9	-0.7	-1.0	0.8	1.6

Table 6: Financial Policies

	Fed Put					
	Anticipated			Unanticipated		
	(1)	(2)	(3)	(4)	(5)	(6)
Threshold parameter (δ, %)	0.1	0.25	0.5	5	10	20
Δ welfare (%)	8.3	17.1	22.4	0.3	0.9	3.0
Δ average GDP (%)	5.7	9.0	11.6	0.1	0.4	1.0
Δ average consumption (%)	1.0	1.4	1.6	0.0	0.1	0.2
Δ employment-population ratio mean (ppt)	5.0	8.0	10.3	0.1	0.4	1.0
Δ employment-population ratio vol (ppt)	-0.7	-1.1	-1.5	-0.1	-0.2	-0.6
Δ unemployment rate mean (ppt)	-2.0	-2.8	-3.2	-0.1	-0.3	-0.8
Δ unemployment rate vol (ppt)	-0.8	-1.1	-1.2	-0.1	-0.4	-0.7
Δ income growth vol (ppt)	-3.3	-5.6	-7.5	-0.1	-0.1	-0.7

Online Appendix

A Additional Details on the Empirical Analysis

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

A.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. Then, after constructing all relevant variables, we randomly sample 20% of all workers in each year for inclusion in our final dataset to keep the analysis computationally feasible. 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.

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.

A.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

$$y_{jt} = \beta_0 + \beta_k k_{jt} + \beta_l l_{jt} + \omega_{jt} + \eta_{jt}, \tag{A.1}$$

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

$$\Delta\omega_{jt} = \Delta y_{jt} - \hat{\beta}_k \Delta k_{jt} - \hat{\beta}_l \Delta l_{jt}. \tag{A.2}$$

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. To obtain measures of industry-level or aggregate TFP growth, we compute the weighted average of firm TFP growth where we weight firms by their lagged number of employees.

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, Pistaferri, and Schivardi

(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.

A.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 $t+1$.

A.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.

A.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.

A.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.

A.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.

B Model Appendix

B.1 A Microfoundation for Worker Preferences

We show how the worker preferences in the main text 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 time-varying ambiguity aversion over aggregate risk but standard risk aversion γ over idiosyncratic risk. We argue that, under an approximation described below, the agent's value function can be represented by

$$V_{i,t} = \left((1-\beta) c_{i,t}^{1-1/\phi} + \beta \, E_t\Big[\Gamma_{t+1} V_{i,t+1}^{1-\gamma} \Big]^{\frac{1-1/\phi}{1-\gamma}} \right)^{\frac{1}{1-1/\phi}}, \tag{A.3}$$

where the preference shifter Γ_{t+1} is given by

$$\Gamma_{t+1} = \exp\left(-\tfrac{1}{2} x_t^2 - x_t \, \varepsilon_{A,t+1} \right). \tag{A.4}$$

In the main text, we set $\phi = 1/\gamma$, under which $\frac{1-1/\phi}{1-\gamma} = 1$ and (A.3) simplifies to equation (6).

The agent has Epstein-Zin preferences with certainty equivalent $\mathcal{R}_t$:

$$V_{i,t} = \left((1-\beta) c_{i,t}^{1-1/\phi} + \beta \left(\mathcal{R}_t(V_{i,t+1}) \right)^{1-1/\phi} \right)^{\frac{1}{1-1/\phi}}. \tag{A.5}$$

Under standard expected utility, $\mathcal{R}_t(V) = E_t[V^{1-\gamma}]^{1/(1-\gamma)}$. We modify the certainty equivalent to incorporate ambiguity aversion over aggregate risk.

Entropic penalty for ambiguity. Following Hansen and Sargent (2001), the agent evaluates continuation utility under a worst-case probability measure, penalized by its distance from the reference measure P:

$$\tilde{\mathcal{R}}_t(V_{i,t+1}) = \min_{Q \in \mathcal{Q}} \left\{ E_t^Q[V_{i,t+1}] + \frac{1}{2}\,\Psi_{i,t}\, E_t^Q\left[\log \frac{dQ}{dP}\right] \right\}. \tag{A.6}$$

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,

$$\Psi_{i,t} = \frac{V_{i,t}}{x_t}, \tag{A.7}$$

to preserve scale invariance (homotheticity) of the preferences. 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. In general, the solution to the inner minimization in (A.6) yields an optimal distortion dQ^*/dP that depends on the agent's continuation value $V_{i,t+1}$ and hence on her entire consumption process, including its idiosyncratic components. The degree of pessimism would therefore be endogenous to the agent's state. We simplify by treating x_t as an exogenous process that governs the degree of ambiguity aversion for all agents. This is an approximation. It would be exact if every agent's consumption were proportional to aggregate productivity A_t state-by-state, since then the optimal distortion would depend only on the aggregate state. In our model, workers are hand-to-mouth and their consumption equals their wage, which is not proportional to A_t; however, all wages and flow payoffs in the model scale with A_t and are therefore 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 the habit in their preferences over the exogenous aggregate productivity (or endowment) process rather than the endogenous consumption process.

Under this simplification, when $\varepsilon_{A,t+1} \sim N(0,1)$ under the reference measure P, the optimal distortion takes the form of an exponential tilt:

$$\frac{dQ^*}{dP} = \Gamma_{t+1} = \exp\left(-\tfrac{1}{2}x_t^2 - x_t\,\varepsilon_{A,t+1}\right). \tag{A.8}$$

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:* $E_t^P[\Gamma_{t+1}] = 1$, 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

$$KL(Q^*\|P) = E_t^{Q^*}[\log \Gamma_{t+1}] = E_t^P[\Gamma_{t+1} \log \Gamma_{t+1}] = \tfrac{1}{2}x_t^2, \tag{A.9}$$

   which is increasing in the degree of ambiguity aversion x_t.

From entropic penalty to multiplicative distortion. Given the exogenous Γ_{t+1} derived above, the certainty equivalent under the optimal distortion can be written as $\tilde{\mathcal{R}}_t(V_{i,t+1}) = E_t[\Gamma_{t+1} V_{i,t+1}^{1-\gamma}]^{1/(1-\gamma)}$, where Γ_{t+1} enters multiplicatively inside the CES aggregator. Substituting into the Epstein-Zin recursion yields (A.3) with Γ_{t+1} given by (A.4).

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.

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. Sahin (2017, November). Gross worker flows over the business cycle. *American Economic Review 107*(11), 3447–76.

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.

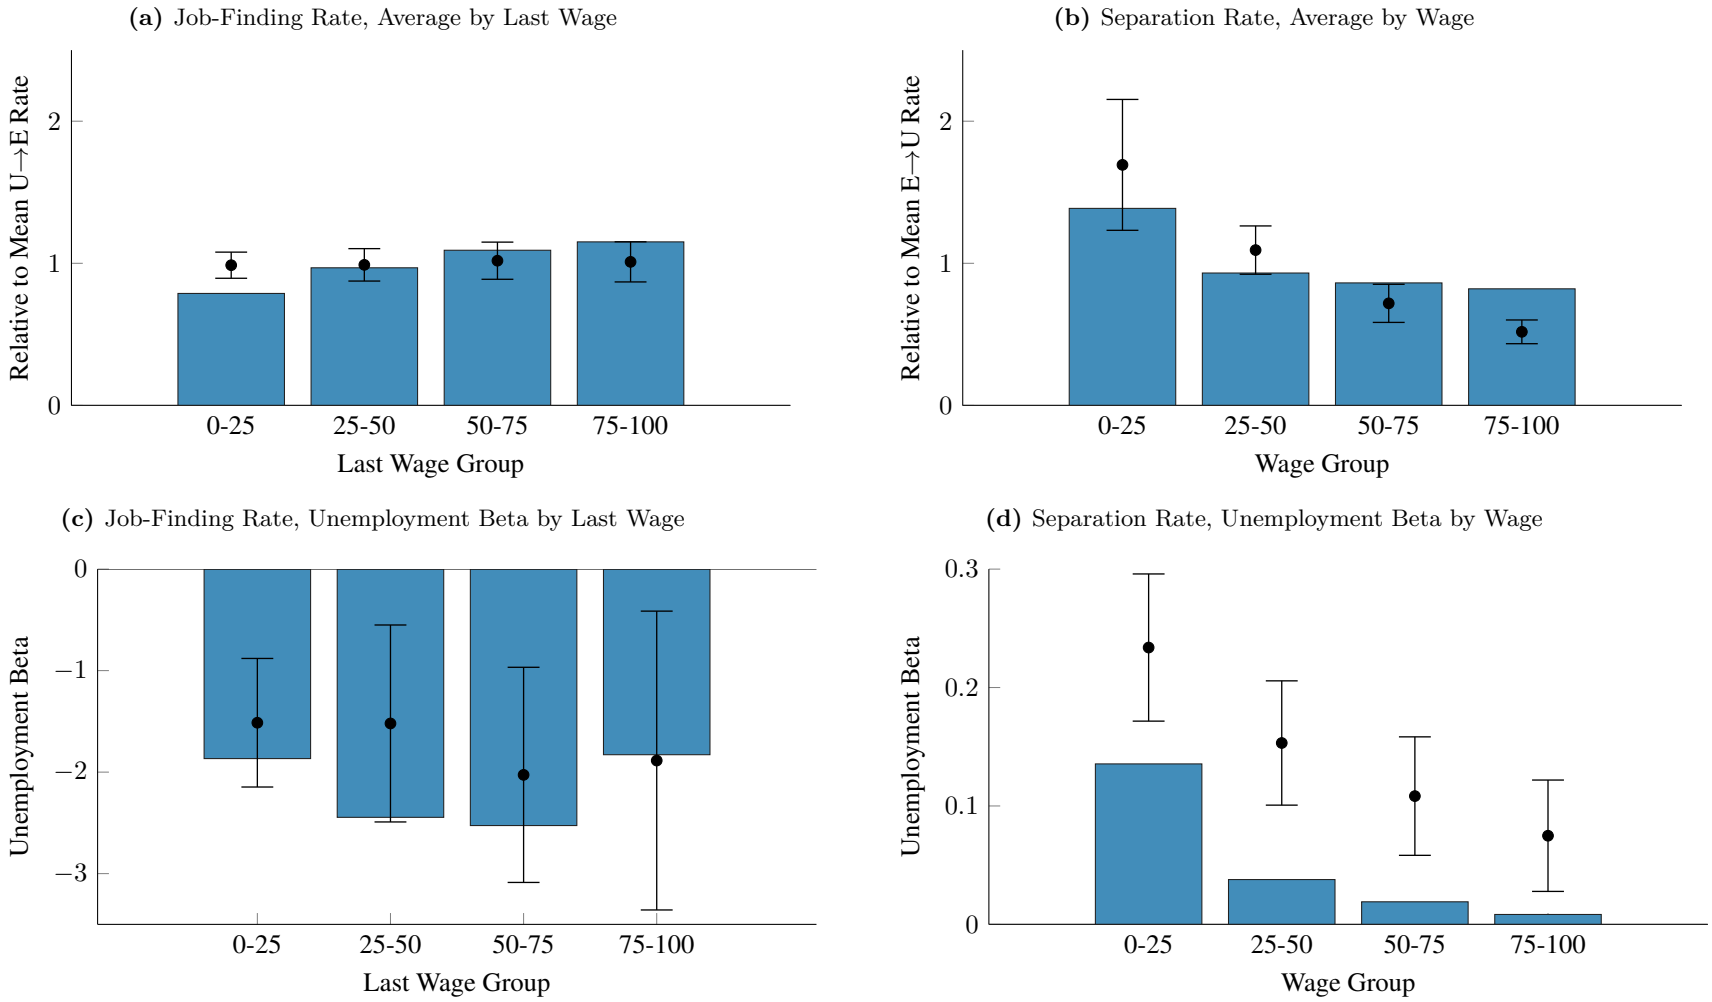
Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*(6), 1263–1297.

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.

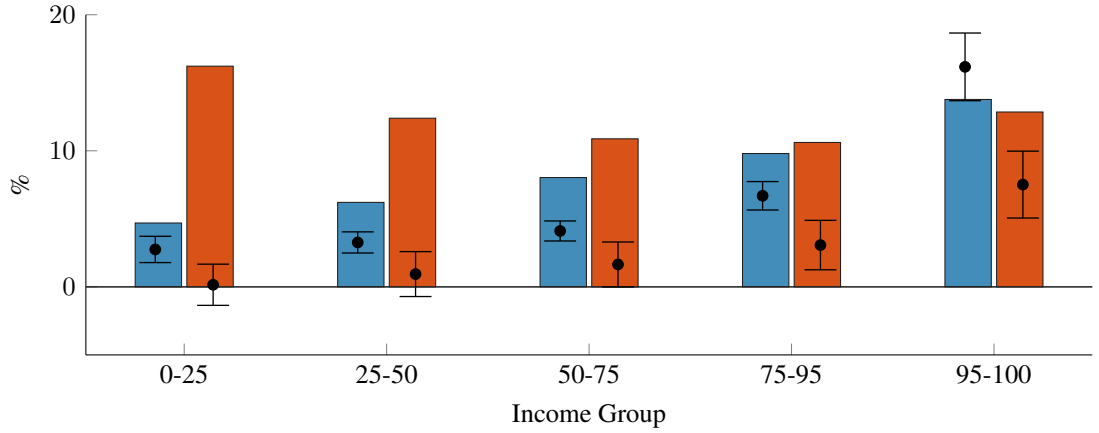
Figure A.1: Separation and Job-Finding Rates by Worker Income: Model vs. Data



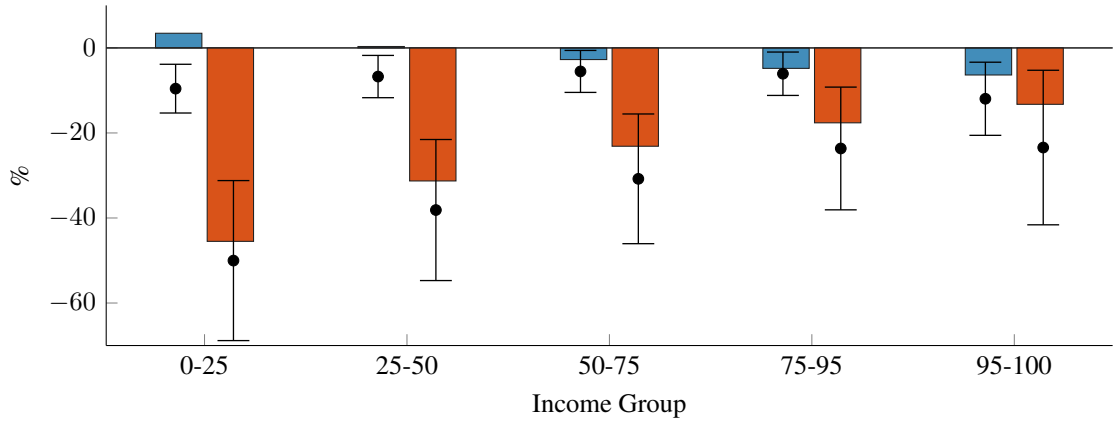
This figure compares the average and cyclicality (unemployment beta) of the job-finding rate (U→E) and the separation rate into unemployment (E→U) by income group in the model and in the data. The empirical counterparts are computed from the SIPP, adjusted for flow level differences from the CPS. Unemployed 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.2: Response of Wage Earnings, Stayers vs. Movers: Model vs. Data

(a) Firm TFP Growth



(b) Risk Premium Shocks



(c) Firm TFP Growth × Risk Premium Shocks

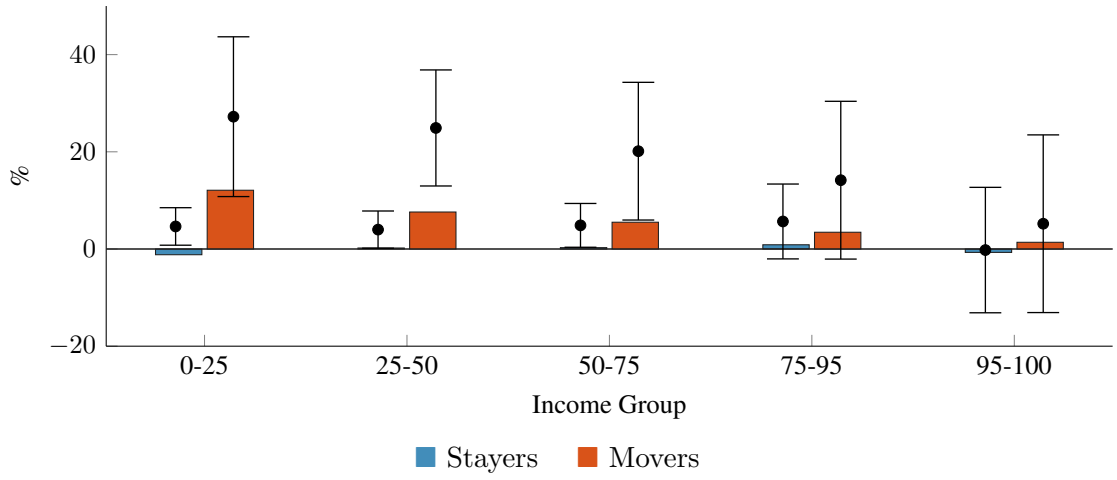
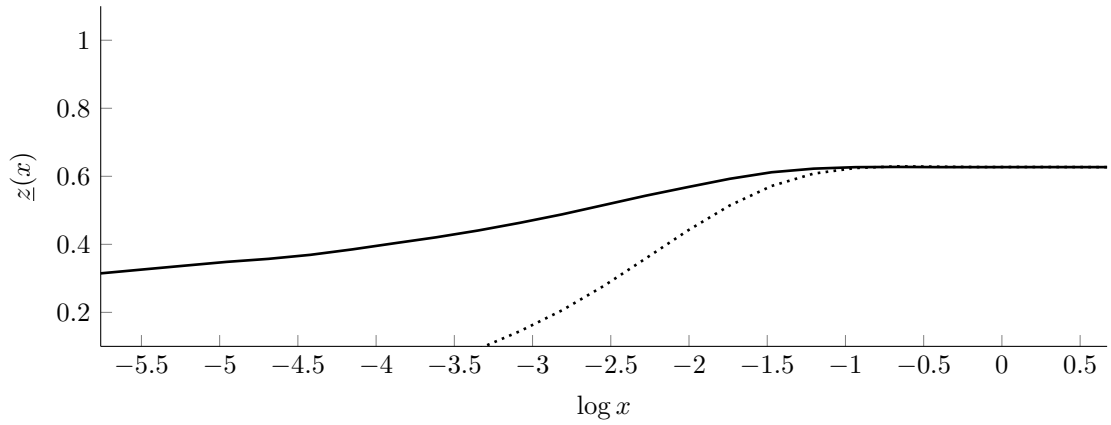
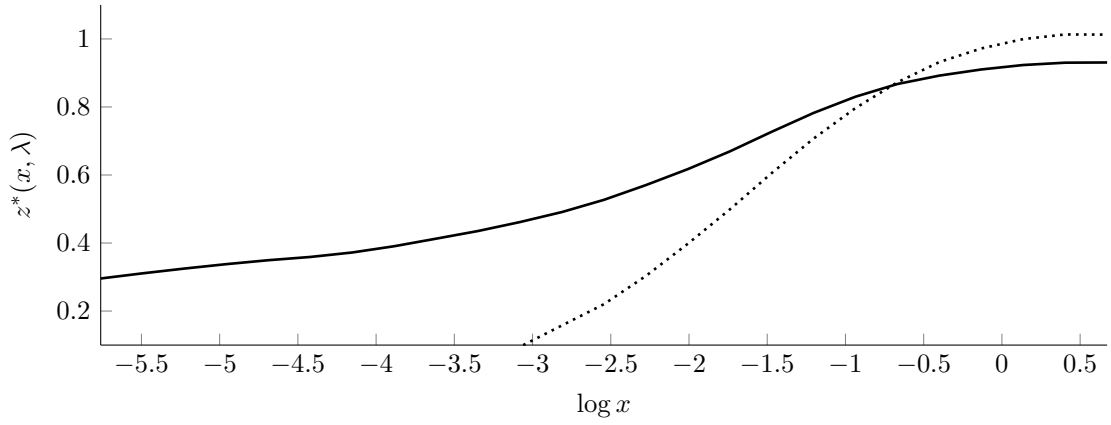


Figure A.3: Risk Premia and Job Destruction and Labor Force Participation Decisions

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



(b) Worker Termination Threshold, Low Experience $z^*(x, \lambda_L)$



(c) Worker Termination Threshold, High Experience $z^*(x, \lambda_H)$

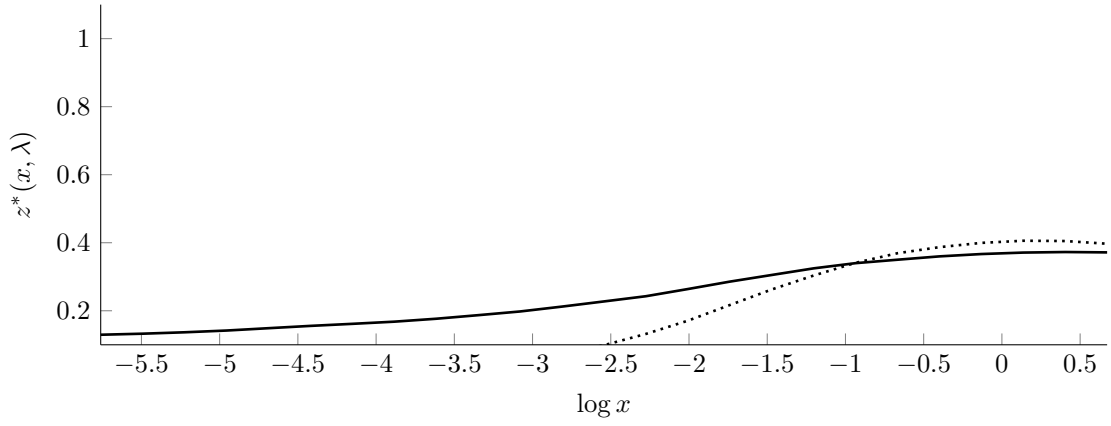


Figure A.4: Pass-Through of Shocks to Workers in Model

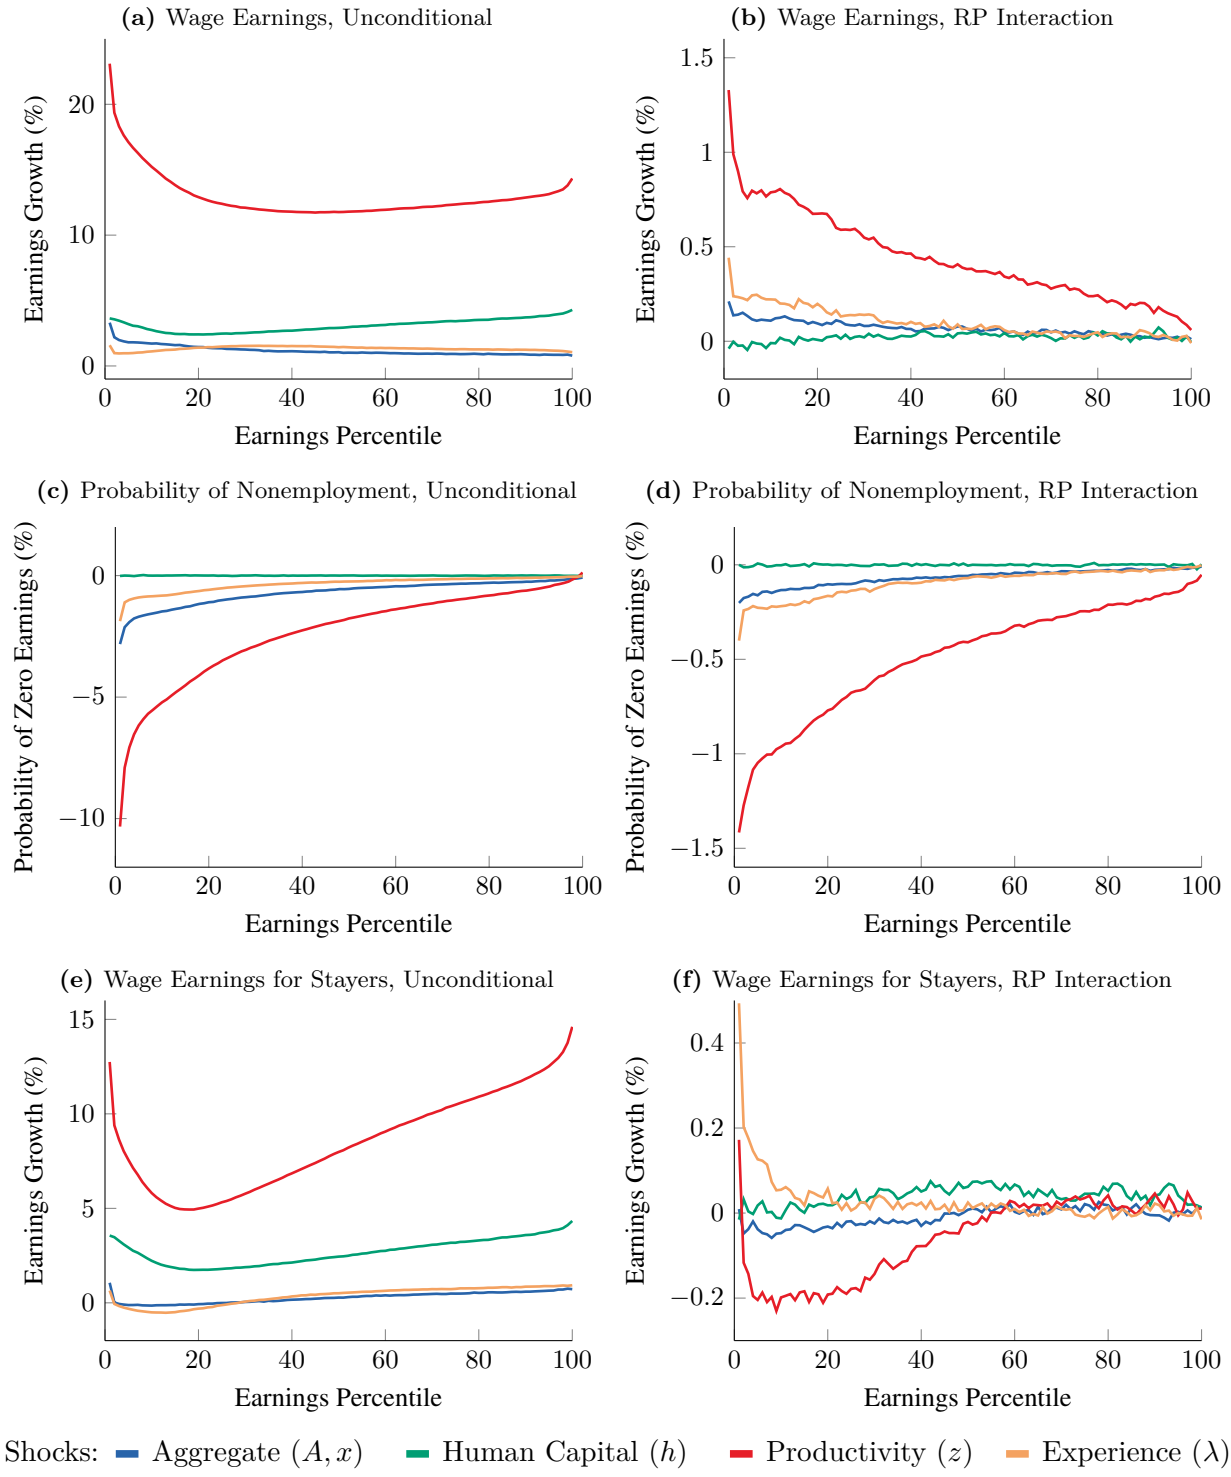


Figure A.5: Labor Income Risk, Higher Order Moments: Model vs. Data

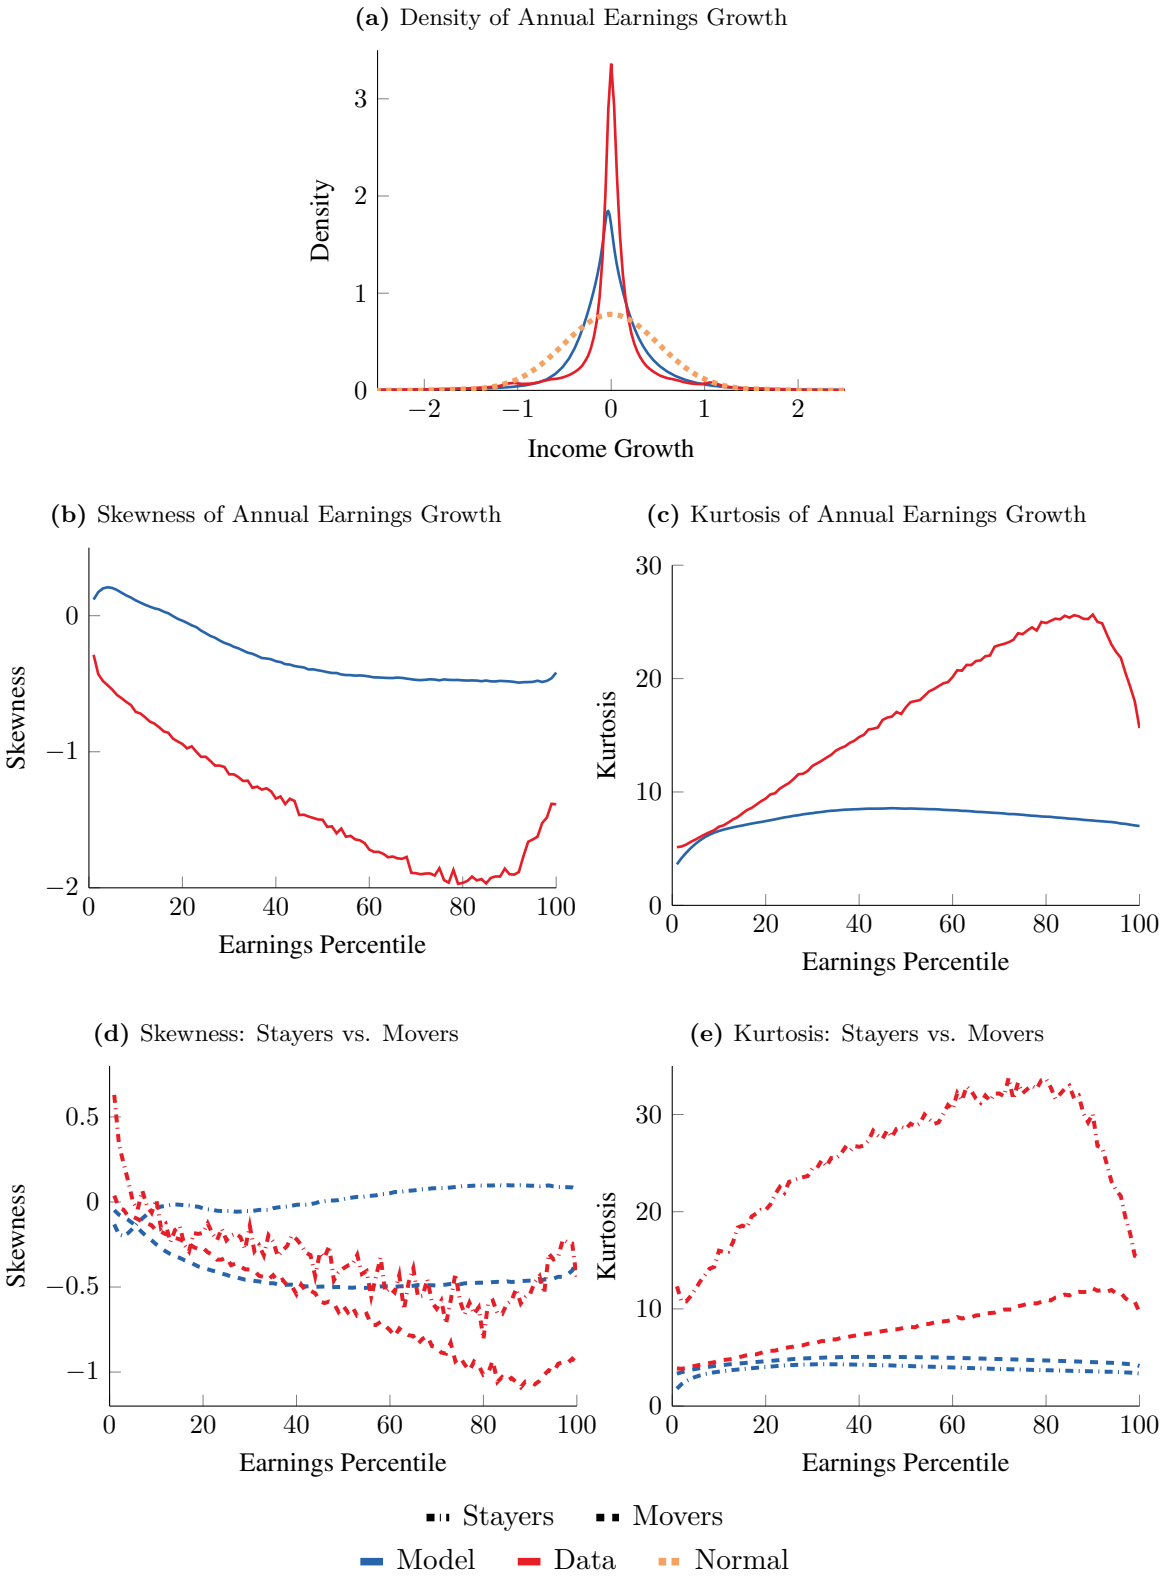
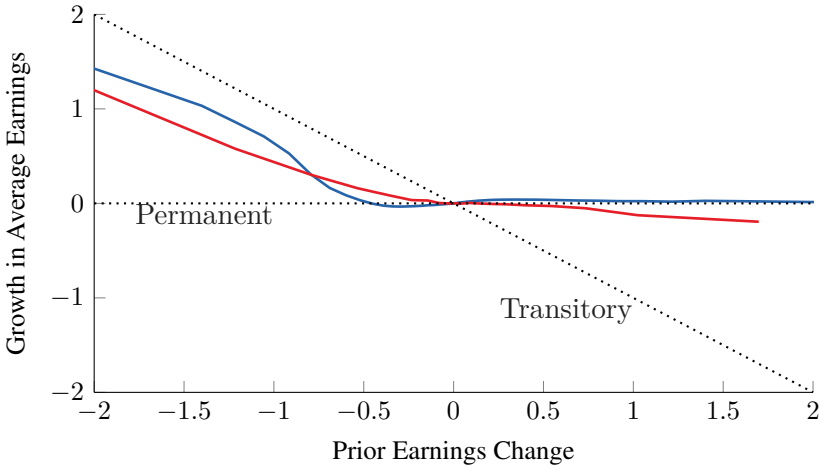


Figure A.6: Impulse Responses of Earnings Growth

(a) Workers with Low Recent Earnings



(b) Workers with High Recent Earnings

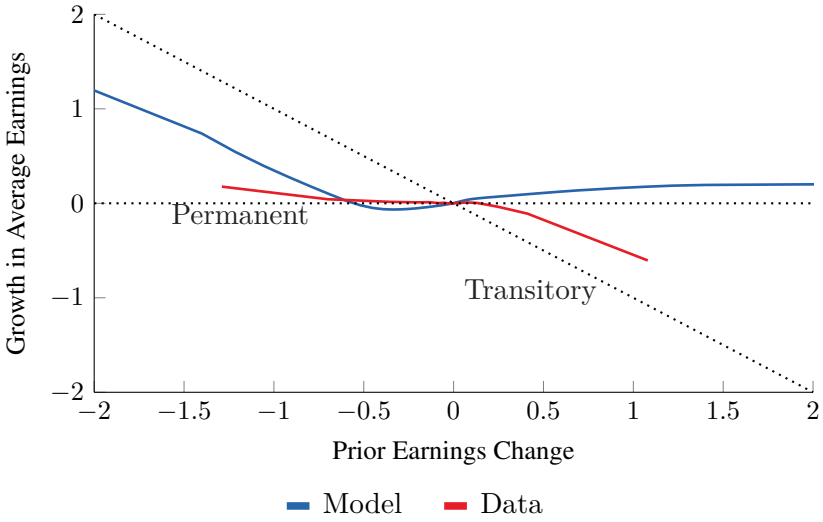


Table A.1: Pass-Through of Firm TFP Growth across Horizons: Stayers versus Movers

	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