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

November 2025

Abstract

Using administrative employer–employee linked data on millions of U.S. workers, we show that the pass-through of firm productivity shocks to worker earnings varies substantially over time and increases when risk premia rise. We develop an equilibrium model of directed labor market search with dynamic wage contracts, limited commitment, and time-varying risk premia that endogenously links productivity shocks and financial conditions to worker earnings dynamics. The model matches key labor market facts—including realistic unemployment dynamics with procyclical job finding and countercyclical separations—as well as the heterogeneous pass-through of productivity and financial shocks across workers and over time, large and non-normal earnings risk, and countercyclical tail risk in worker earnings growth. The model implies sizable welfare losses from idiosyncratic labor income risk, leading to depressed valuations of human capital.

*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).

Worker earnings risk is both significant and unevenly distributed across individuals and over time. The distribution of worker earnings exhibits significant deviations from log normality, and its moments exhibits significant heterogeneity across workers and across time (Guvenen, Karahan, Ozkan, and Song, 2021). Earnings risk is closely tied to firm risk: Friedrich, Laun, Meghir, and Pistaferri (2019) show that there is significant pass-through of firm idiosyncratic shocks to workers, particularly for the highest paid workers. At the aggregate level, recessions sharply raise the probability of large income losses (Guvenen, Ozkan, and Song, 2014). Meeuwis, Papanikolaou, Rothbaum, and Schmidt (2025) show that fluctuations in risk premia can generate significant ex-post heterogeneity in worker earnings losses driven by endogenous separations: an increase in discount rates leads to an increased possibility of separation for low-productivity workers, leading to unemployment and significant earnings losses.

In this paper, we show that time-varying risk premia lead to endogenous fluctuations in labor income risk. We start by documenting a new fact: the extent to which firm risk pass through idiosyncratic productivity shocks to worker earnings is not constant over time, but rather it exhibits a significant degree of time variation that is related to fluctuations in risk premia. This fact is a direct implication of the model of Meeuwis et al. (2025). In particular, an increase in risk premia leads firms to be more selective in which workers they keep: they increase the productivity threshold for retaining a worker. This increase implies that workers experiencing adverse productivity shocks are now at greater risk of job loss. Consistent with this prediction, we show that the extent to which firm risk passes through to worker earnings is not constant over time, but rather it exhibits a significant degree of time variation that is positively correlated with risk premium shocks. That is, the pass-through of firm shocks to worker earnings is higher when risk premia increase, particularly for lower-paid workers that are closer to the separation threshold. Even though both the passthrough and risk premia are countercyclical, this positive correlation is robust to controls for the business cycle.

Building on the idea that the degree to which firms insure workers is not constant but varies with risk premia, we develop a labor search model with optimal wage contracting that jointly matches key moments of labor markets as well as worker earnings, and has rich implications for workers' labor income risk, welfare, and the value of human capital. Following Meeuwis et al. (2025), the model features time-varying risk premia, directed labor market search and endogenous separations—given that the possibility of job loss is a key driver of earnings dynamics. The key elements in their model is that firms' and workers' effective level of risk aversion with respect to aggregate shocks varies over time; this time variation in risk premia directly affects the valuation of an employment match compared to the worker's value of unemployment. Increases in discount rates imply a reduction in the value of employment, particularly for low-productivity workers, because the benefits of employment are more backloaded, and hence are more sensitive to discount rates, than the benefits of unemployment. However, we depart from their model in several important ways.

Most importantly, we let firms provide insurance to risk averse workers subject to a limited commitment friction. Specifically, workers are risk averse and face incomplete markets: they have no access to financial markets that allow them to smooth shocks to their labor income; firms fill in that role by offering workers a full set of state-contingent wage payments, subject to the constraint that firms' continuation values are non-negative and workers values exceed their outside option. Allowing firms to smooth worker earnings has a significant effect, both on the dynamics of wages, but also, crucially, on the probability of separation: since job loss is associated with a significant decline in worker earnings, firms are more likely to retain low-productivity workers compared to the case where workers are risk neutral as in Meeuwis et al. (2025). In addition, this optimal dynamic contract has an endogenous aspect of wage rigidity conditional on employment: firms endeavor to keep workers' wages fixed, unless doing so would violate workers' or firms' participation constraints. Hence, the distribution of worker earnings growth is endogenously lepto-kurtotic even though the underlying worker productivity shocks are log normal.

The key insight in our model is that the scope for firms to insure worker risk depends on the firms' valuation of an employment match. Increases in discount rates reduce firms' value of a match surplus, and therefore reduce the scope for insuring workers, both against endogenous termination but also against fluctuations in worker productivity. In particular, if the firm's value of an employment match becomes non-positive following an increase in discount rates, workers either receive a wage reduction, or if there is no feasible new wage that the worker prefers relative to her outside option, the match is dissolved. Hence, increases in risk premia lead to lower job creation and higher job destruction—and both of these channels interact in generating earnings losses for workers when risk premia rise. As a result, increases in risk premia lead to endogenous increases in labor income risk, particularly so for the lower-productivity workers that are closer to the termination threshold.

In addition, our model features a set of additional elements that help the model quantitatively the stylized features of the data. First, we allow workers to receive outside offers while employed similar to Balke and Lamadon (2022). The possibility of receiving an outside offer is increasing in the worker's current productivity. Given that worker wages endogenously respond to fluctuations in workers' outside options, this assumption allows the model to generate a degree of passthrough of firm risk to worker earnings that is increasing in the workers' level of income, consistent with the findings of Friedrich et al. (2019).

We calibrate the model to match the dynamics of asset prices, labor market flows, and the passthrough of productivity and discount rate shocks to worker earnings. We choose parameters governing risk premia to fit key asset pricing moments, following Lettau and Wachter (2007). We set the remaining parameters to match labor market facts, with a focus on cross-sectional and time-series moments of separation and job-finding rates. Overall, the model does a good job fitting the data.

In addition, our model implies that the pass-through of productivity shocks to worker earnings is highly nonlinear and state-dependent. A negative worker productivity shock increases the likelihood of job loss, and therefore has a larger effect on worker earnings than a positive productivity shock of the same magnitude. The level of risk premia determine the likelihood of job destruction, hence this asymmetry becomes starker as risk premia rise. Further, firms aim to smooth wages, but are limited in their ability to do so due to the two-sided lack of commitment. Therefore, smaller shocks to productivity have a (proportionally) smaller pass-through than larger shocks, since these shocks are more likely to lead to binding commitment constraints. The end result is that the distribution of workers' labor income is significantly more skewed and leptokurtotic than the underlying productivity shocks, consistent with the findings of Guvenen et al. (2021). Further, the distribution of earnings growth becomes more negatively skewed as risk premia rise—particularly for low income workers, which in turn leads to an increase in income inequality at the bottom of the distribution.

The model also replicates the realized paths of key labor market variables. Feeding in our empirical risk premium shocks, we generate model-implied series that track the data closely: correlations range from 50–80% and volatilities are comparable. The model captures the slow recovery after the Great Recession, driven by persistently high premia that depress job creation and erode human capital through prolonged nonemployment. An increase in risk premia implies an increase in the probability of job loss, while lowering output and employment, therefore generating counter-cyclical skewness in labor earnings, quantitatively replicating the fluctuations in income risk in the data (Guvenen et al., 2014).

Using the calibrated model, we perform a number of quantitative exercises. First, we compute the welfare losses that accrue to workers as a result of their idiosyncratic labor income risk. That is, we compare their certainty equivalent in our model to a counterfactual in which their risk aversion with respect to idiosyncratic shocks is zero, but otherwise keep the same path for wages and employment status as in our baseline model. The welfare cost of these idiosyncratic shocks is substantial (25 percent). This cost is mostly uniform along the earnings distribution, since all workers in the model are hand-to-mouth and worker payoffs in nonemployment are relatively safe.

Second, we compute the value of workers' human capital, both from the perspective of the worker and but also from the perspective of a diversified shareholder (firms). Naturally, workers have a lower valuation for their own human capital than firms, because labor income has significant idiosyncratic risk which firms can diversify away. However, the magnitude of this discount varies with worker productivity and earnings. Since lower-paid workers have higher idiosyncratic income risk than higher-paid workers, the discount is decreasing with worker earnings. If we were to extend the model to allow workers to invest in their human capital at birth, the model would imply that idiosyncratic income risk leads to under-investment in human capital for lower-paid workers.

In a series of seminal papers using panel data from the U.S. Social Security Administration (SSA), Guvenen et al. (2014) and Guvenen et al. (2021) document a set of stylized facts regarding how the distribution of worker earnings risk varies across workers and across time. Guvenen et al.

(2021) focus on characterizing the moments of individual earnings growth. They first show that earnings growth is extremely volatile—much more than what is typically assumed in macroeconomic models. Volatility itself varies across the cross section: younger and lower-income workers face higher variances in earnings growth than older and higher-income workers. Most importantly, job stayers have much lower dispersion in earnings growth than job movers. In addition, Guvenen et al. (2021) show that the shape of the earnings growth distribution departs sharply from normality. It is highly leptokurtic, with a sharp peak at zero and very fat tails. Last, Guvenen et al. (2021) document significant dispersion in lifetime earnings and lifetime nonemployment rate: for instance, 40% of men experience at most one year of nonemployment, while 18% spend more than half of their working years as nonemployed. Guvenen et al. (2014) focus on the time-series properties of worker earnings risk: they show that the main cyclical change is not in the variance of idiosyncratic shocks, which is only weakly countercyclical, but rather in the skewness of the earnings growth distribution.

A number of papers explore the economic drivers' of worker earnings risk. For example, a number of studies show that firms partially insure workers: using matched employer–employee study for Italy, Guiso, Pistaferri, and Schivardi (2005) show that firms fully insure workers against transitory shocks to firm performance but only partially insure against permanent shocks. Quantitatively, their estimates imply essentially zero pass-through for transitory shocks and a small but significant pass-through for permanent shocks (approximately equal to 0.07). In follow-up work that uses matched employer-employee data from Sweden, Friedrich et al. (2019) show that there is significant pass-through of firm-specific productivity shocks to workers. The pass-through is increasing with income, and is particularly high for the highest paid workers. Chan, Xu, and Salgado (2019) use matched employer-employee from Denmark and find that the pass-through of firm shocks to worker earnings is asymmetric: it is higher for negative than positive firm shocks. Using employer-employee matched data from the Unites States, Meeuwis et al. (2025) show that shocks to risk premia pass through to worker earnings; low-paid workers are significantly more exposed to risk premium shocks than higher-paid workers. Meeuwis et al. (2025) show that the impact of risk premium shocks on worker earnings is mainly driven by separations.

In addition to Meeuwis et al. (2025), the papers that are closest to our work are Kehoe, Midrigan, and Pastorino (2019); Kehoe, Lopez, Midrigan, and Pastorino (2023) and Balke and Lamadon (2022). Similar to our work, Kehoe et al. (2019, 2023) feature time-variation in risk premia in a model of directed labor market search. However, their focus is on aggregate unemployment fluctuations and hence feature complete markers whereas our focus is on worker's labor income risk and hence markets are incomplete in our model and workers are risk averse. Close to our paper, Balke and Lamadon (2022) also consider dynamic contracts in a model of incomplete markets labor market search. However, our work has some significant differences, both in terms of modeling choices and focus. In particular, we are primarily interested in how fluctuations in discount rates endogenously generate labor income risk for workers. Hence, in contrast to Balke and Lamadon (2022), firms cannot commit to retaining (ex-post) unproductive workers. In addition, whereas Balke and Lamadon (2022) focus on generating a realistic passthrough of firm shocks to worker earnings, our focus is on how time-varying discount rates lead to heterogeneity in labor income risk, both over time and also across workers.

1 Pass-Through of Firm Shocks to Worker Earnings

1.1 Data and Methodology

Our empirical analysis combines worker earnings data with firm-level information. We use a random subsample of earnings records from the Longitudinal Employer–Household Dynamics (LEHD) database for the period 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 of the sample construction.

The main outcome variable is cumulative, age-adjusted earnings growth, following Autor, Dorn, Hanson, and Song (2014) and Guvenen et al. (2014):

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

This measure emphasizes persistent changes by averaging earnings over multiple years. The denominator $D(\text{age}_{i,\tau})$ adjusts for the average life-cycle earnings profile. To be included in year t, a worker must be employed by a Compustat firm in that year. We then follow individuals over time regardless of subsequent employment status, so that the growth measure includes earnings from both public and private firms as well as nonemployment spells (zero wage income). We winsorize growth rates at the 1st and 99th percentiles each year.

In our empirical analysis, we link worker earnings to firm-level productivity shocks and aggregate risk premium shocks. To measure firm productivity growth, $\epsilon^{tfp}_{f,t+1}$, we build on İmrohoroğlu and Tüzel (2014) to estimate annual revenue-based total factor productivity (TFPR); Appendix A.2 provides further detail. 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 a broad set of existing measures: 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, Ludvigson, and Ng (2015); the risk appetite index of Bauer, Bernanke, and Milstein (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). Each series is monthly, with signs oriented so that higher values indicate higher risk premia. Appendix A.3 provides additional detail. 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. We denote the resulting series of risk premium shocks as ϵ^{rp}. This factor captures the dominant source of variation: it 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 $t+1$.

1.2 Cross-Sectional Heterogeneity in Pass-Through to Workers

We begin by briefly summarizing existing evidence on cross-sectional heterogeneity in the pass-through of different shocks to worker earnings. We focus on heterogeneity by prior earnings rank, measured relative to other workers within the same firm. Following Meeuwis et al. (2025), we estimate the following specification,

$$g_{i,t:t+h} = \sum_s \{\beta_s \, \epsilon^{tfp}_{f(i,t),t+1} + \gamma_s \, \epsilon^{rp}_{t+1}\} \, \mathbb{1}_{\tau(i,t)=s} + c'\mathbf{Z}_{i,t} + \eta_{i,t+h}. \tag{2}$$

Here, $g_{i,t:t+h}$ is worker i's cumulative earnings growth over horizon h defined in equation (1), $f(i,t)$ 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 $\tau(i,t)$. 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 × age × gender fixed effects. We cluster standard errors by worker and year.

Figure 1a shows that firm productivity shocks are only partially passed through to worker earnings. The effect of TFP growth on cumulative earnings is roughly flat—around 3 percent at a one-year horizon and 5 percent at a five-year horizon—for workers in the bottom 90 percent of the within-firm earnings distribution, and rises sharply at the top. Pass-through reaches 15 to 20 percent for workers in the top few percentiles, indicating that higher-paid workers are considerably more exposed to firm productivity shocks than lower-paid workers. These findings closely mirror those in Friedrich et al. (2019), who document similar heterogeneity in the wage responses of job stayers in Sweden. Like Friedrich et al. (2019), we also find extensive-margin effects of productivity shocks: Figure 1b shows that adverse firm shocks raise nonemployment risk, particularly among low-rank workers. Hence, for these workers, firm downturns translate primarily into separations rather than wage cuts.

Figure 2 presents the analogous estimates for risk premium shocks, closely mirroring the results in Meeuwis et al. (2025). At the one-year horizon, earnings exposure is U-shaped across the within-firm earnings distribution: high for low-wage workers, declining toward the 90th percentile, and increasing again at the top. This pattern is consistent with the U-shaped pattern of worker betas with respect to stock market returns documented by Guvenen, Schulhofer-Wohl, Song, and Yogo (2017). At longer horizons, the exposure of low-wage workers rises further, while that of high-wage workers weakens, rendering the relationship nearly monotonic. As emphasized by Meeuwis et al. (2025), the strong exposure at the bottom primarily reflects extensive-margin effects. Consistent with this interpretation, Figure 2b shows that increases in risk premia are associated with a pronounced increase in separation risk for lower-paid workers, whereas there are no such effects for high-wage workers. Hence, risk premium shocks negatively affect high-wage workers only through the intensive margin.

1.3 A Stylized Model of Worker Earnings Dynamics

Meeuwis et al. (2025) show that the empirical patterns documented in Section 1.2 are consistent with a directed search model with time-varying risk premia that jointly generates realistic labor market dynamics and worker earnings pass-through. To motivate our subsequent empirical analysis and theoretical framework, we briefly outline a stylized model of worker earnings, in the spirit of Meeuwis et al. (2025), that links aggregate financial conditions to employment and wage dynamics. The model is deliberately parsimonious and relies on reduced-form assumptions governing job separations, job finding, and wage growth.

The aggregate state of the economy is summarized by aggregate productivity A, which follows a random walk with drift, and by an aggregate financial conditions index x that captures the level of risk premia and evolves as a persistent but mean-reverting process. For simplicity, we assume that innovations to A and x are perfectly negatively correlated.

Individual workers are heterogeneous in their employment status (employed or unemployed) and in the level of their persistent human capital z, which evolves as a mean-reverting process. Unemployed workers find a new job next period with probability $p(x)$, which is independent of individual productivity but decreasing in x. A worker–firm match generates a surplus that is increasing in z, since the worker's outside option is less sensitive to productivity, and decreasing in x, because higher risk premia lower the discounted value of future match surplus. Matches are dissolved whenever the surplus becomes negative, that is, if and only if $z \leq z^*(x)$, where the threshold $z^*(x)$ is an increasing function of x. Consequently, a rise in x leads to additional separations, especially among low-productivity workers close to the threshold.

Finally, conditional on continued employment, log wage growth follows

$$\log(w'/w) = \phi(z)\,\{\log(A'/A) + \log(z'/z)\}, \tag{3}$$

where $\phi(z)$ governs the pass-through of productivity shocks to wages. The function $\phi(z)$ is independent of x but increasing in z, implying that high-productivity workers experience greater wage sensitivity to both aggregate and idiosyncratic shocks.

Appendix Figure A.1 illustrates the pass-through coefficients implied by a calibrated version of this model. We see that the pass-through of firm-level shocks to earnings varies systematically with risk premia. When x is high, separations become more likely and job finding slows, amplifying the earnings response to firm shocks for workers near the separation margin. In contrast, high-productivity workers are largely insulated from this channel, so their earnings are less sensitive to financial conditions.

1.4 Time Variation in Pass-Through to Workers

Motivated by the stylized model in Section 1.3, we next provide new evidence on how the pass-through of firm shocks to worker earnings varies across time and workers. We begin by examining how the pass-through of firm-specific shocks to worker earnings varies over time. Specifically, we estimate the following specification,

$$g_{i,t:t+h} = \beta_t \, \epsilon^{tfp}_{f(i,t),t+1} + \xi_{I(i,t),\tau(i,t),t} + c'\mathbf{Z}_{i,t} + \eta_{i,t+h}, \tag{4}$$

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. In particular, we note that periods when the pass-through of *firm-specific* shocks is high systematically coincide with increases in risk premia.

To explore this relationship more formally, we next estimate a variant of equation (2) that interacts firm productivity shocks with risk premium shocks and allows the coefficients to vary with the worker's prior earnings rank within her firm,

$$g_{i,t:t+h} = \sum_s \{\beta_{0,s} \, \epsilon^{tfp}_{f(i,t),t+1} + \beta_{1,s} \, \epsilon^{tfp}_{f(i,t),t+1} \times \epsilon^{rp}_{t+1} + \gamma_s \, \epsilon^{rp}_{t+1}\} \, \mathbb{1}_{\tau(i,t)=s} + c'\mathbf{Z}_{i,t} + \eta_{i,t+h}. \tag{5}$$

In alternative specifications, we also interact the firm-specific shocks by earnings group with aggregate output growth or the fraction of the year spent in NBER recessions.

Table 1 reports the estimated coefficients β_0 and β_1 from estimates of equation (5) for a three-year horizon. A key pattern stands out: although the average pass-through β_0 of firm-specific shocks increases with workers' prior earnings, this pattern reverses following an increase in risk premia. That is, the estimated interaction coefficients β_1 are large and positive for lower-paid workers and decline sharply with prior earnings. This finding is robust across alternative specifications and remains unchanged when controlling for aggregate output growth or NBER recessions. These estimates imply that higher discount rates are associated with a greater overall pass-through of firm-specific shocks to worker earnings, with this amplification concentrated among low-wage workers. Columns (1)–(3) of Table 2 shows that these estimates are robust across horizons of one to five years.

To understand the source of these earnings losses, we next examine the role of job separations in generating the patterns in Table 1. As before, our proxy for job loss is an indicator that equals one if the worker experiences at least one full quarter with zero wage earnings (a nonemployment spell) over the next h years. Columns (4)–(6) of Table 2 report the estimated coefficients β_0 and β_1 from modified versions of equation (5), where the outcome variable is our measure of job destruction. We consider horizons of one to three years. These estimates show that negative firm productivity shocks are associated with a greater likelihood of job loss for lower-paid workers. Moreover, the interaction coefficients β_1 indicate that this probability rises sharply when risk premia increase, with the effect again concentrated among lower-paid workers.

In Table A.1, we re-estimate equation (5) separately for workers who leave their initial employer (movers) and those who remain employed by it (stayers). A worker is classified as a stayer at horizon h if she continues to earn a positive amount of labor income from her time-t employer in year $t + h + 1$; otherwise, she is classified as a mover. The results show that the increase in the pass-through of firm-specific shocks during periods of elevated risk premia is substantially larger for movers than for stayers, conditional on prior income. In particular, lower-paid movers experience the largest increase in pass-through when discount rates rise.

In sum, the results in this section show that the extent to which firms insure their workers—the pass-through of firm-specific shocks to worker earnings—varies systematically across both time and workers. Increases in risk premia are associated with higher pass-through of firm-specific shocks (that is, weaker firm insurance) than when discount rates fall. This amplification is concentrated among lower-paid workers (relative to their peers within the same firm) and operates primarily through higher rates of job destruction.

2 Model

We begin by presenting a model of the labor market with search frictions (Diamond, 1982; Mortensen, 1982; Pissarides, 1985). The model features heterogeneous risk averse workers that are subject to shocks to their productivity. We model a directed search process in which firms search for workers with different levels of productivity (Montgomery, 1991; Moen, 1997). Worker productivity is stochastic and persistent. Firms can 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. Importantly, the model also features on-tho-job search (Postel-Vinay and Robin, 2002).

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.

2.1 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

$$y_t(\Omega_{i,t}) = A_t\, h_{i,t}\, z_{i,t}\, \lambda_{i,t}. \tag{6}$$

Here, A_t is equal to aggregate productivity, while h, z, and λ are different components of worker productivity; the worker's current level productivity is summarized by the vector $\Omega = (h, z, \lambda)$.

Aggregate productivity A_t follows a random walk with a drift:

$$\Delta \log A_{t+1} = \mu_A + \sigma_A\, \varepsilon_{A,t+1}, \quad \varepsilon_{A,t+1} \sim N(0,1). \tag{7}$$

The first component of worker-specific productivity h is permanent and evolves according to the following process:

$$\Delta \log h_{i,t+1} = g_{i,t} + \sigma_h \, \varepsilon_{h,i,t+1}, \quad \varepsilon_{h,i,t+1} \sim N(0,1). \tag{8}$$

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:

$$\log z_{i,t+1} = \psi_z \log z_{i,t} + (1-\psi_z) \log \bar{z} + \sigma_z \, \varepsilon_{z,i,t+1}, \quad \varepsilon_{z,i,t+1} \sim N(0,1). \tag{9}$$

The key difference between z and h is that z refers to job-specific productivity whereas h refers to 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 $\lambda_{i,t} \in \{\bar{\lambda}_L, \bar{\lambda}_H\}$ is a two-point Markov process; while the worker is employed, λ evolves according to the following transition probability matrix,

$$T_\lambda = \begin{bmatrix} 1-f & f \\ 0 & 1 \end{bmatrix}. \tag{10}$$

The component λ is meant to capture an element of job-specific experience. Hence, as long as the worker is employed, there is a flow probability $f \, \Delta t$ that an inexperienced worker becomes experienced at her job. An experienced worker retains her experience level as long as she remains on her job. 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 non participation).

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

Worker Preferences and Payoffs

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

$$V_{i,t} = \left\{ c_{i,t}^{1-\gamma} + \beta \, (1-\zeta) \, \mathbb{E} \left[\Gamma_{t+1} \, V_{i,t+1}^{1-\gamma} \, | \, \mathcal{F}_t \right] \right\}^{\frac{1}{1-\gamma}}, \tag{11}$$

Equation (11) 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: the effective degree of worker risk aversion towards aggregate shocks is time-varying. In particular, worker's attitude towards aggregate risk are modified by the process

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

where

$$\log x_{t+1} = \psi_x\,\log x_t + (1-\psi_x)\log\bar{x} - \sigma_x\,\varepsilon_{A,t+1}. \tag{13}$$

Examining our specification for worker utility (11)–(13), 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 (12), 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 (13), 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 and hence consume their flow wage each period $c_{i,t} = w_{i,t}$. Workers that are out of a job engage in home production: they receive and consume a benefit that depends on their labor market participation status. Workers that decide to enter the unemployment pool and actively search for a new job consume

$$b_t(\Omega) = A_t\,h\,(\bar{b}_0 + \bar{b}_1\,z). \tag{14}$$

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

[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 (12) capturing the extent of this behavioral distortion.

a new job in the upcoming search-and-matching stage, consume a benefit given by

$$n_t(\Omega) = \bar{n}\, A_t\, h. \tag{15}$$

The main difference between the flow benefits of unemployment (14) and non-participation (15) is that the latter do 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).

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 hold a well diversified portfolio of firms and can save in a risk-free security that is in zero net-supply. Shareholders have the same preferences as the workers (11), but in contrast to workers, they have a risk aversion of zero with respect to idiosyncratic shocks $\gamma = 0$, which 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

$$P_t = \mathbb{E}_t \left\{ \sum_{\tau=1}^{\infty} \left(\prod_{s=1}^{\tau} \Lambda_{t+s} \right) X_{t+\tau} \right\}, \tag{16}$$

where Λ_{t+s} is the one-period shareholders' stochastic discount factor (SDF) between periods $t+s-1$ and $t+s$ given by

$$\Lambda_{t+1} = \beta \, \exp \left\{ -\frac{1}{2}\, x_t^2 - x_t\, \varepsilon_{A,t+1} \right\}. \tag{17}$$

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 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

$$\chi(\Omega) = \frac{z^{\bar{\chi}_1}}{\bar{\chi}_0 + z^{\bar{\chi}_1}}. \tag{18}$$

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

$$\kappa_t(\Omega) = \bar{\kappa}_0 \, A_t \, h \, z^{\bar{\kappa}_1}. \tag{19}$$

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

$$m(u, \nu) = \frac{u\,\nu}{(u^\alpha + \nu^\alpha)^{\frac{1}{\alpha}}}. \tag{20}$$

This matching function ensures that job-finding and vacancy-filling rates are bounded between zero and one. Specifically, equation (20) 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 $p(\theta) = \theta(1 + \theta^\alpha)^{-1/\alpha}$.

2.2 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^C of staying in the current match decides to break up the match and leaves for nonemployment if and only if V^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:

$$R_t(\Omega, V) \equiv \sup_{\mathcal{V}} \; p(\theta_t(\Omega, \mathcal{V})) \, \frac{\mathcal{V}^{1-\gamma} - V^{1-\gamma}}{1-\gamma}. \tag{21}$$

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

$$\tilde{p}_t(\Omega, V) \equiv 1 - \chi(\Omega) \, p_t^*(\Omega, V). \tag{22}$$

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

$$V_t^O(\Omega) = \max\{V_t^U(\Omega), V_t^N(\Omega)\}. \tag{23}$$

A worker that decides to be out of the labor force collects the non-participation benefit (15) and, conditional on surviving, will be a nonemployed worker in the next period. Her continuation value is therefore equal to

$$V_t^N(\Omega) = \Big\{ n_t(\Omega)^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\,V_{t+1}^O(\Omega')^{1-\gamma}\Big]\Big\}^{\frac{1}{1-\gamma}}. \tag{24}$$

An unemployed worker collects unemployment benefits (14) and, conditional on surviving, actively searches for a job during the search-and-matching stage of the next period. Hence, her continuation value is

$$V_t^U(\Omega) = \Big\{ b_t(\Omega)^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{V_{t+1}^O(\Omega')^{1-\gamma} + (1-\gamma)\,R_{t+1}\Big(\Omega', V_{t+1}^O(\Omega')\Big)\Big\}\Big]\Big\}^{\frac{1}{1-\gamma}}. \tag{25}$$

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:

$$\mathbb{1}_t^C(\Omega, V^C) = \begin{cases} 1 & \text{if } J_t^{FC}(\Omega, V^C) \geq 0 \text{ and } V^C \geq V_t^O(\Omega) \\ 0 & \text{otherwise.} \end{cases} \tag{26}$$

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,

$$J_t^{FC}(\Omega, V^C) = \max_{\pi_j, V_j^S} \quad \sum_{j=1}^{2} \pi_j\,\tilde{p}_t(\Omega, V_j^S)\,J_t^{FS}(\Omega, V_j^S) \tag{27}$$

$$\text{s.t.} \quad V^C \leq \left(\sum_{j=1}^{2} \pi_j \Big\{ (V_j^S)^{1-\gamma} + \chi(\Omega)\,(1-\gamma)\,R_t(\Omega, V_j^S) \Big\} \right)^{\frac{1}{1-\gamma}}. \tag{28}$$

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

$$\Pi_t(\Omega, V) = -\kappa_t(\Omega) + q(\theta_t(\Omega, V))J_t^{FS}(\Omega, V). \tag{29}$$

Firms create such a vacancy if and only if the benefit (29) 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:

$$\Pi_t(\Omega, V) \leq 0 \quad \forall(\Omega, V), \tag{30}$$

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

$$\begin{aligned} V^S = \Big\{ & w^{1-\gamma} + \beta\,(1-\zeta)\,\mathbb{E}_{t,\Omega}\Big[\Gamma_{t+1}\Big\{V^O_{t+1}(\Omega')^{1-\gamma} + \\ & + (1-s)\mathbb{1}^C_{t+1}(\Omega', V^{C\prime})\Big((V^{C\prime})^{1-\gamma} - V^O_{t+1}(\Omega')^{1-\gamma}\Big)\Big\}\Big]\Big\}^{\frac{1}{1-\gamma}}. \end{aligned} \tag{31}$$

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,

$$J^{FS}_t(\Omega, V^S) = \max_{w,\{V^{C\prime}\}} \; y_t(\Omega) - w + (1-\zeta)(1-s)\,\mathbb{E}_{t,\Omega}\Big[\Lambda_{t+1}\mathbb{1}^C_{t+1}(\Omega', V^{C\prime})\,J^{FC}_{t+1}(\Omega', V^{C\prime})\Big], \tag{32}$$

subject to the promise-keeping constraint (31), which holds as an equality along the equilibrium path.

2.3 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.

3 Model Calibration, Fit, and Mechanisms

3.1 Calibration

We calibrate the model in two steps. Table 3 summarizes our baseline calibration of the model.

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 19 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 1.5% per year (1.4% in the data). The calibrated values $\mu_E = 0.001, \bar{x} = 0.074$, $\psi_x = 0.993$, and $\sigma_x = 0.119$ 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.4); 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.5).

We set the exogenous separation rate s to 0.70 percent, targeting average separation rates of higher-wage workers. The unemployment benefit parameters ($\bar{b}_0 = 2.19, \bar{b}_1 = 0.50$) and the nonparticipation payoff $\bar{n} = 2.50$ 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, $g_O = 0.51$ percent 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 ($\bar{\kappa}_0 = 0.017, \bar{\kappa}_1 = 1.51$) 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 3.2.

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

$$\varepsilon_{h,i,t} = \rho_h \, \tilde{\varepsilon}_{h,j,t} + \sqrt{1-\rho_h^2}\, \varepsilon^{\perp}_{h,i,t} \tag{33}$$

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

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

$$\epsilon^{tfp}_{j,t+1} = \log\left(\int_{I_{j,t}} y_{i,t+1}\,di\right) - \log\left(\int_{I_{j,t}} y_{i,t}\,di\right) + \sigma_k\,\tilde{\varepsilon}_{k,j,t+1}, \tag{35}$$

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 $\sigma_k = 0.065$. On the worker side, we target the average volatility of annual worker earnings growth (53%) reported by Guvenen et al. (2014) and set $\sigma_h = 0.03$.

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 $\lambda_H = 2.5$ governs the upper level of worker experience, which creates scope for firms to insure workers against shocks. The parameters $\bar{\chi}_0 = 49$ and $\bar{\chi}_1 = 2.4$ 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 $\rho_h = 0.02$ and $\rho_z = 0.25$ 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. Section 3.3 discusses the resulting model's fit to the empirical pass-through regression estimates.

3.2 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.2 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.

3.3 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 1. We estimate the same regression specification (5) 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 increases with worker income, consistent with 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). 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.

3.4 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_{t+1}^{rp} from Section 1.1 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 71% 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 3.2. 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 (4) on the realized paths, yielding a measure of the pass-through of firm-level productivity shocks for each period. We run these regressions separately on the three-year earnings growth of 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 (60%), 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.

3.5 Discussion of Model Mechanisms

The model features several sources of worker heterogeneity: earnings are driven by idiosyncratic shocks to z, h, and λ. Figure 7 illustrates how each of these shocks feeds into worker earnings. Shocks to z have the biggest effect on earnings (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 stronger wage pass-through 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.

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 Implications

4.1 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: it generates a negatively skewed and highly leptokurtic distribution of earnings growth—even though individual productivity growth is normally distributed. Appendix Figure A.4 shows that the model also replicates the log density of earnings growth remarkably well, including the heavy tails that deviate sharply from a Gaussian distribution. The excess tail risk in worker earnings arises endogenously from the joint dynamics of labor markets and wage contracts. 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.

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. Panels (b)–(d) of Figure 9 assess the model's ability to replicate these cross-sectional patterns. Panel (b) shows that the model quantitatively matches the declining pattern of earnings volatility from the bottom up to about the 90th percentile of the earnings distribution, but it does not reproduce the sharp rise in volatility at the very top. The model also matches the fact that earnings growth for job switchers is much more volatile than for job stayers, particularly for low-income workers (see Appendix Figure A.5). Panels (c) and (d) of Figure 9 show how skewness and kurtosis vary across workers. The model qualitatively matches the hump-shared pattern of skewness and kurtosis, though not with the same magnitudes as in the data.

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 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

Panels (e) and (f) of Figure 9 replicate the analysis of lifetime earnings growth and employment rates in Guvenen et al. (2021). Panel (e) plots cumulative earnings growth between ages 25 and 55 against workers' lifetime earnings percentiles. The upward slope is not surprising and partly mechanical—faster growth implies higher overall earnings—but the amount of heterogeneity is striking: average earnings grow by 163% at the median, 720% at the top 5%, and 1971% at the top 1%. Panel (f) plots the distribution of the total number 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.

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 3.4. 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 58% for the left tail and 33% for the right tail.

4.2 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 (11) 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

$$CE_{i,t} = (1 - \beta (1 - \zeta))^{1/(1-\gamma)} V_{i,t}. \tag{36}$$

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. To obtain a stationary measure of welfare, we normalize by aggregate productivity and evaluate $CE_{i,t}/A_t$. We then average this ratio across workers and over time to characterize welfare in each version of the model.

On average, certainty-equivalent consumption in the baseline valuation is about 25 percent lower than in the counterfactual. This implies that workers would require a permanent increase in consumption of roughly one third to be as well off as under risk neutrality with respect to idiosyncratic shocks. These large welfare losses reflect the substantial lifetime variation in consumption generated by the model's realistic earnings dynamics and the assumption that workers are hand-to-mouth.

Figure 11a compares welfare in the baseline model relative to the counterfactual across workers sorted by current labor productivity. Welfare losses in terms of certainty-equivalent consumption are sizable but fairly uniform across the distribution—around 25 percent. Because flow payoffs for nonemployed workers are less volatile, losses are slightly smaller for the lowest-productivity workers. 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.

4.3 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.

Figure 11b plots 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 modell. 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 panel (b) of Figure 11, 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. Panel (c) 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.

Conclusion

This paper documents a new empirical fact: the pass-through of firm productivity shocks to worker earnings varies substantially over time. We show that periods of elevated risk premia feature weaker firm insurance and greater exposure of low-wage workers to firm-specific shocks, primarily through higher job destruction.

We develop a directed search model with dynamic wage contracts, limited commitment, and time-varying risk premia that endogenously links productivity shocks and aggregate financial conditions to worker earnings dynamics. The model matches key labor market facts—including unemployment fluctuations driven by procyclical job finding and countercyclical separations—as well as the heterogeneous pass-through of shocks to worker earnings, large and non-normal earnings risk, and countercyclical tail risk in labor income.

Our framework emphasizes financial conditions as a central driver of worker earnings dynamics and labor income risk. Large, time-varying, and undiversifiable idiosyncratic income risk generates sizable welfare losses for workers and depresses the valuation of human capital.

References

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.

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

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

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

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.

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.

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.

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.

Guvenen, F., S. Schulhofer-Wohl, J. Song, and M. Yogo (2017). Worker betas: Five facts about systematic earnings risk. *American Economic Review 107*(5), 398–403.

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.

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.

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

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.

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.

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.

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

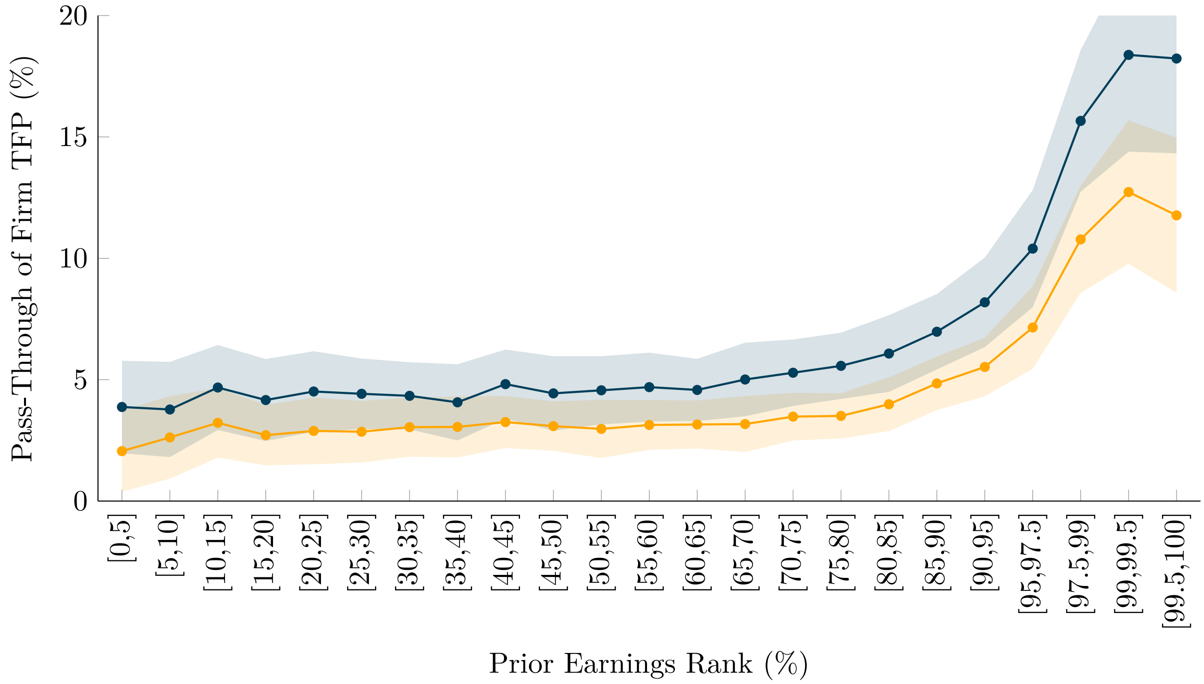

(b) Extensive Margin (Probability of Zero-Earnings Quarter)

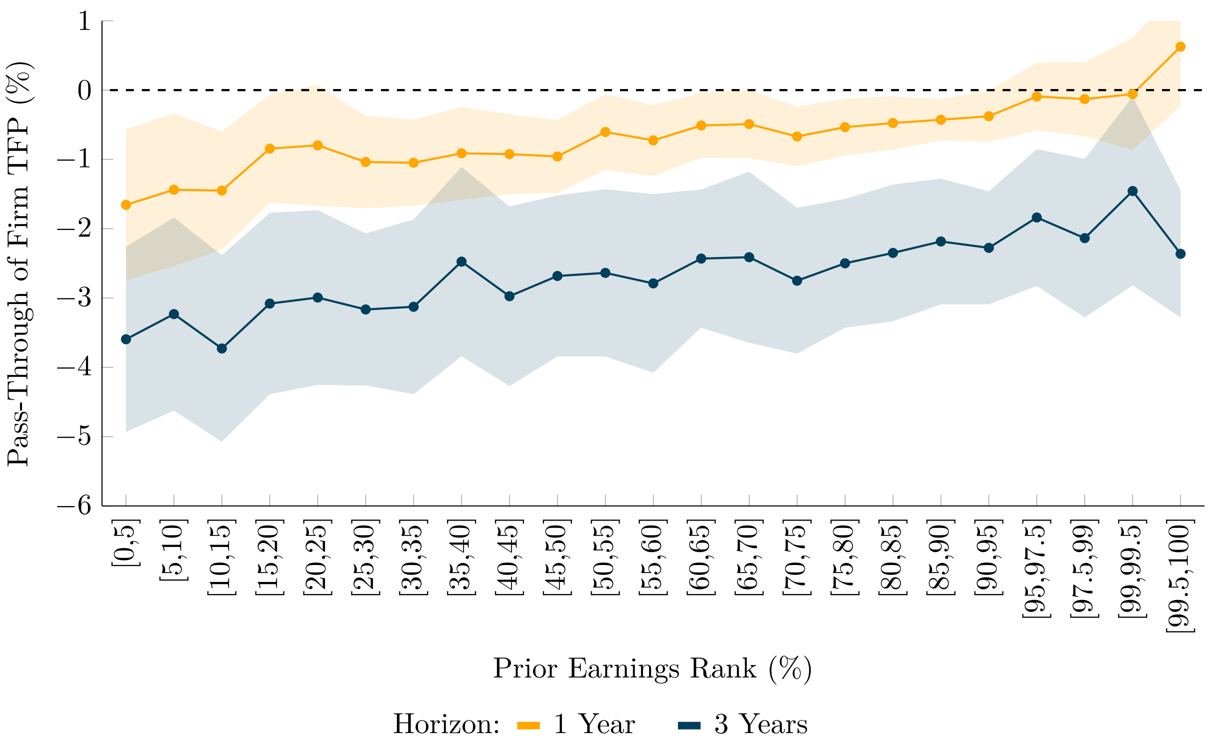

Figure 2: Pass-Through of Risk Premium Shocks to Workers

(a) Wage Earnings

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

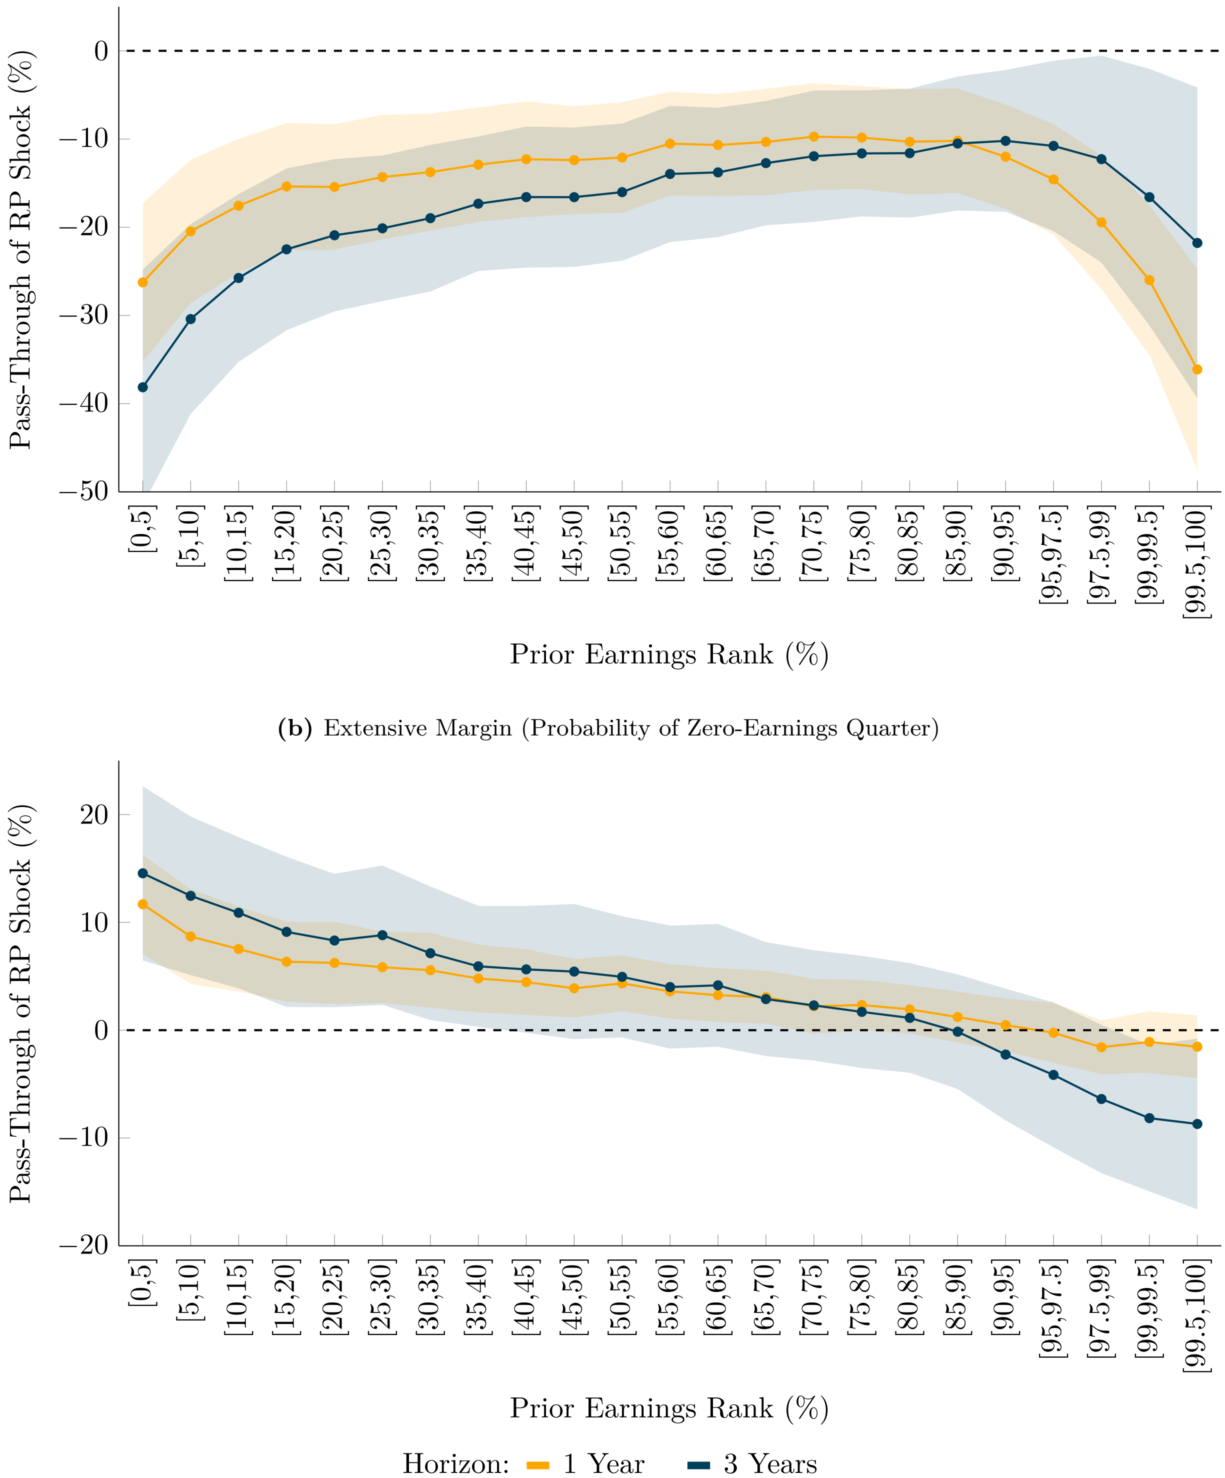

Figure 3: Time-Varying Pass-Through of Firm TFP Growth

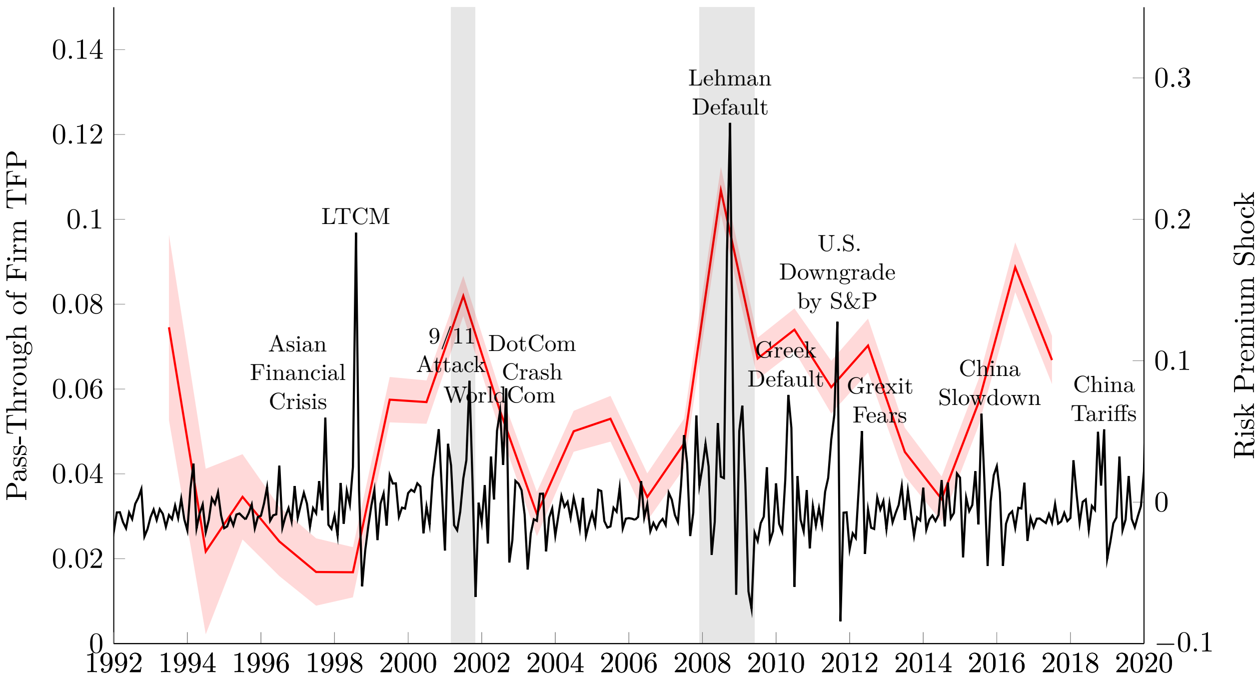

Figure 4: Response of Wage Earnings: Model vs. Data

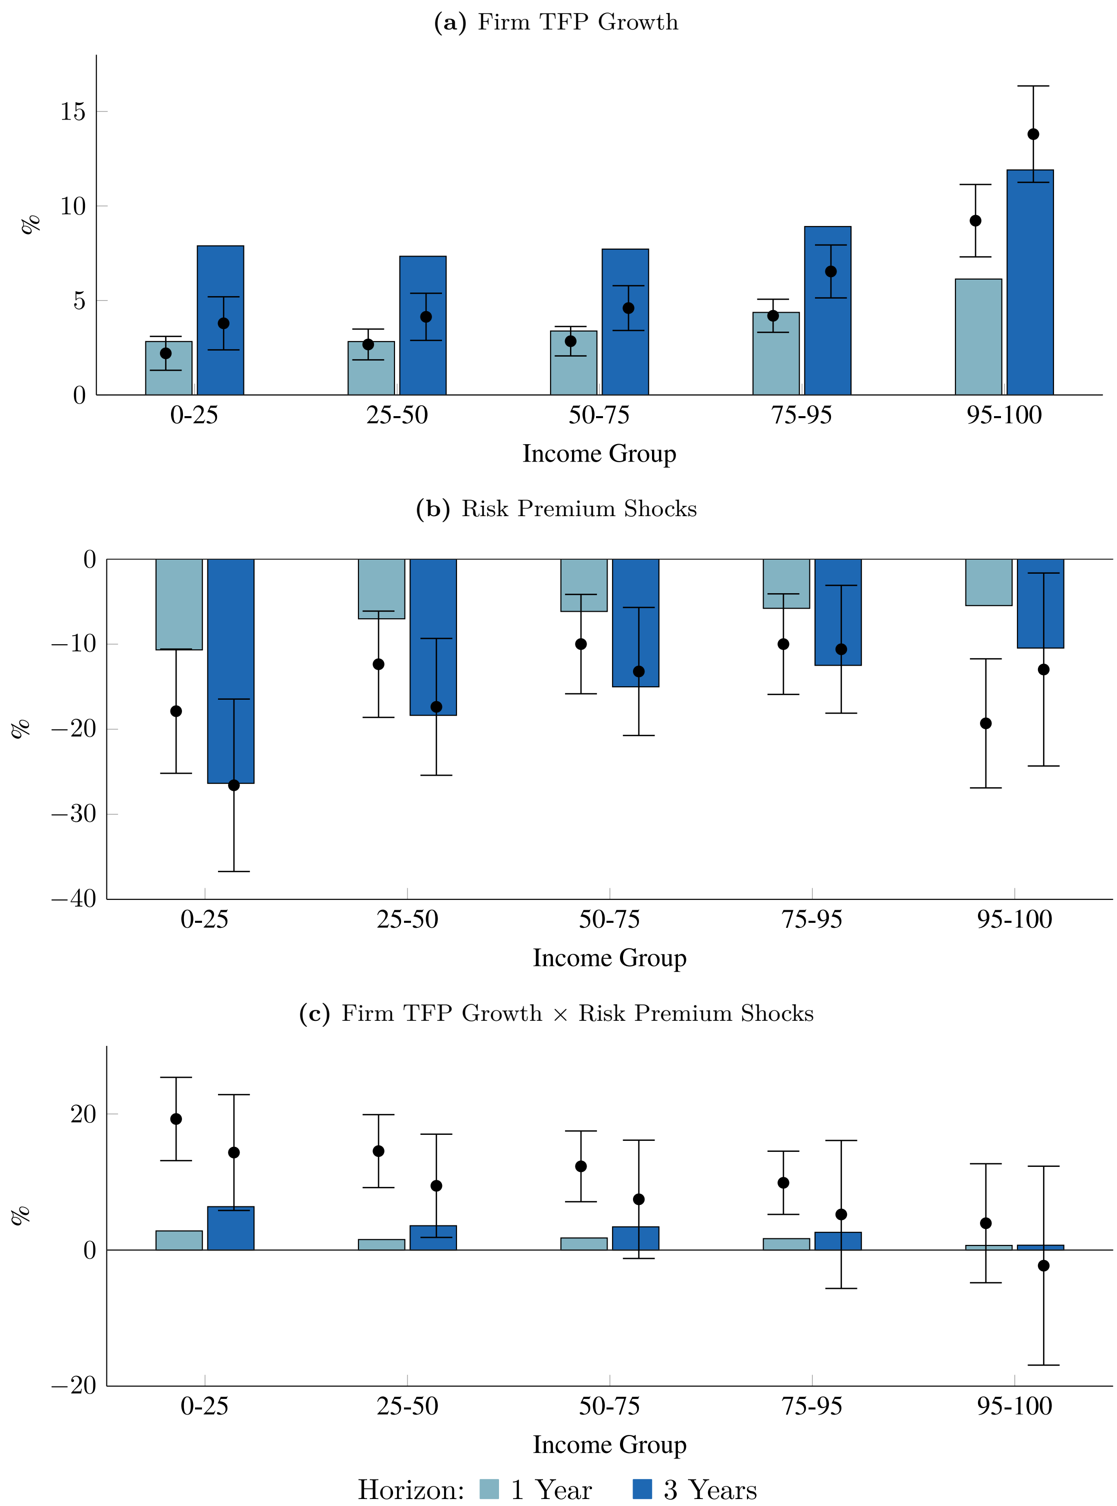

Figure 5: Response of Job Destruction: Model vs. Data

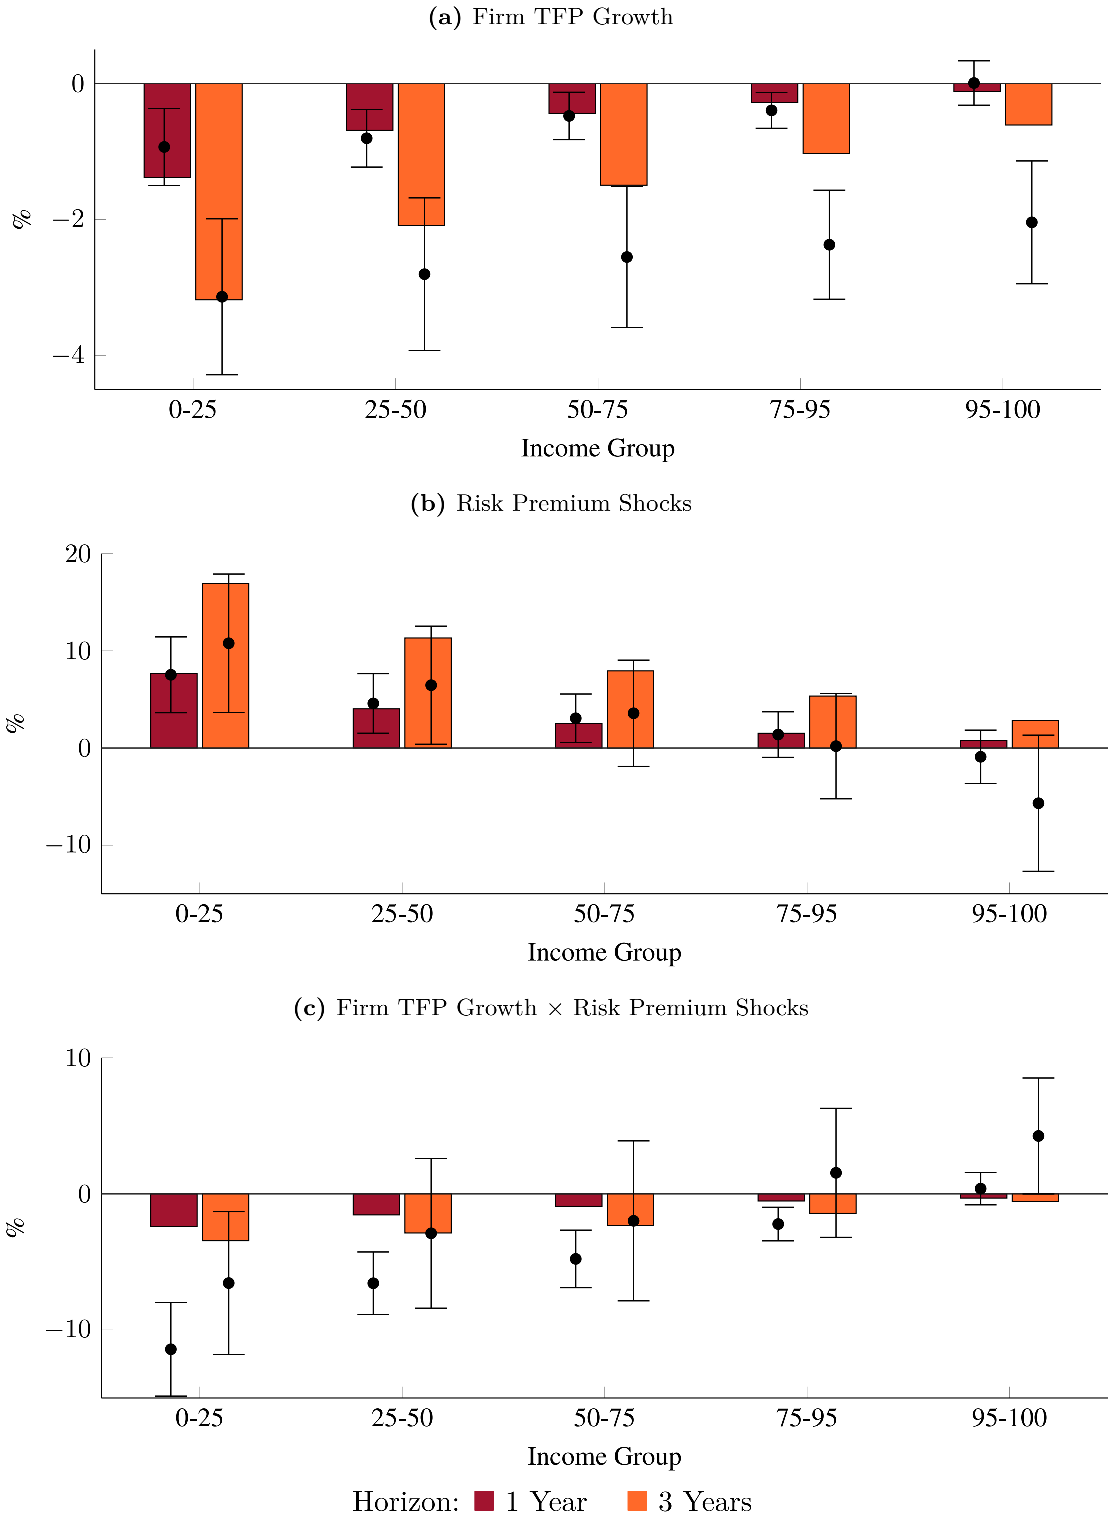

Figure 6: Realized Fluctuations in Labor Markets: Model vs. Data

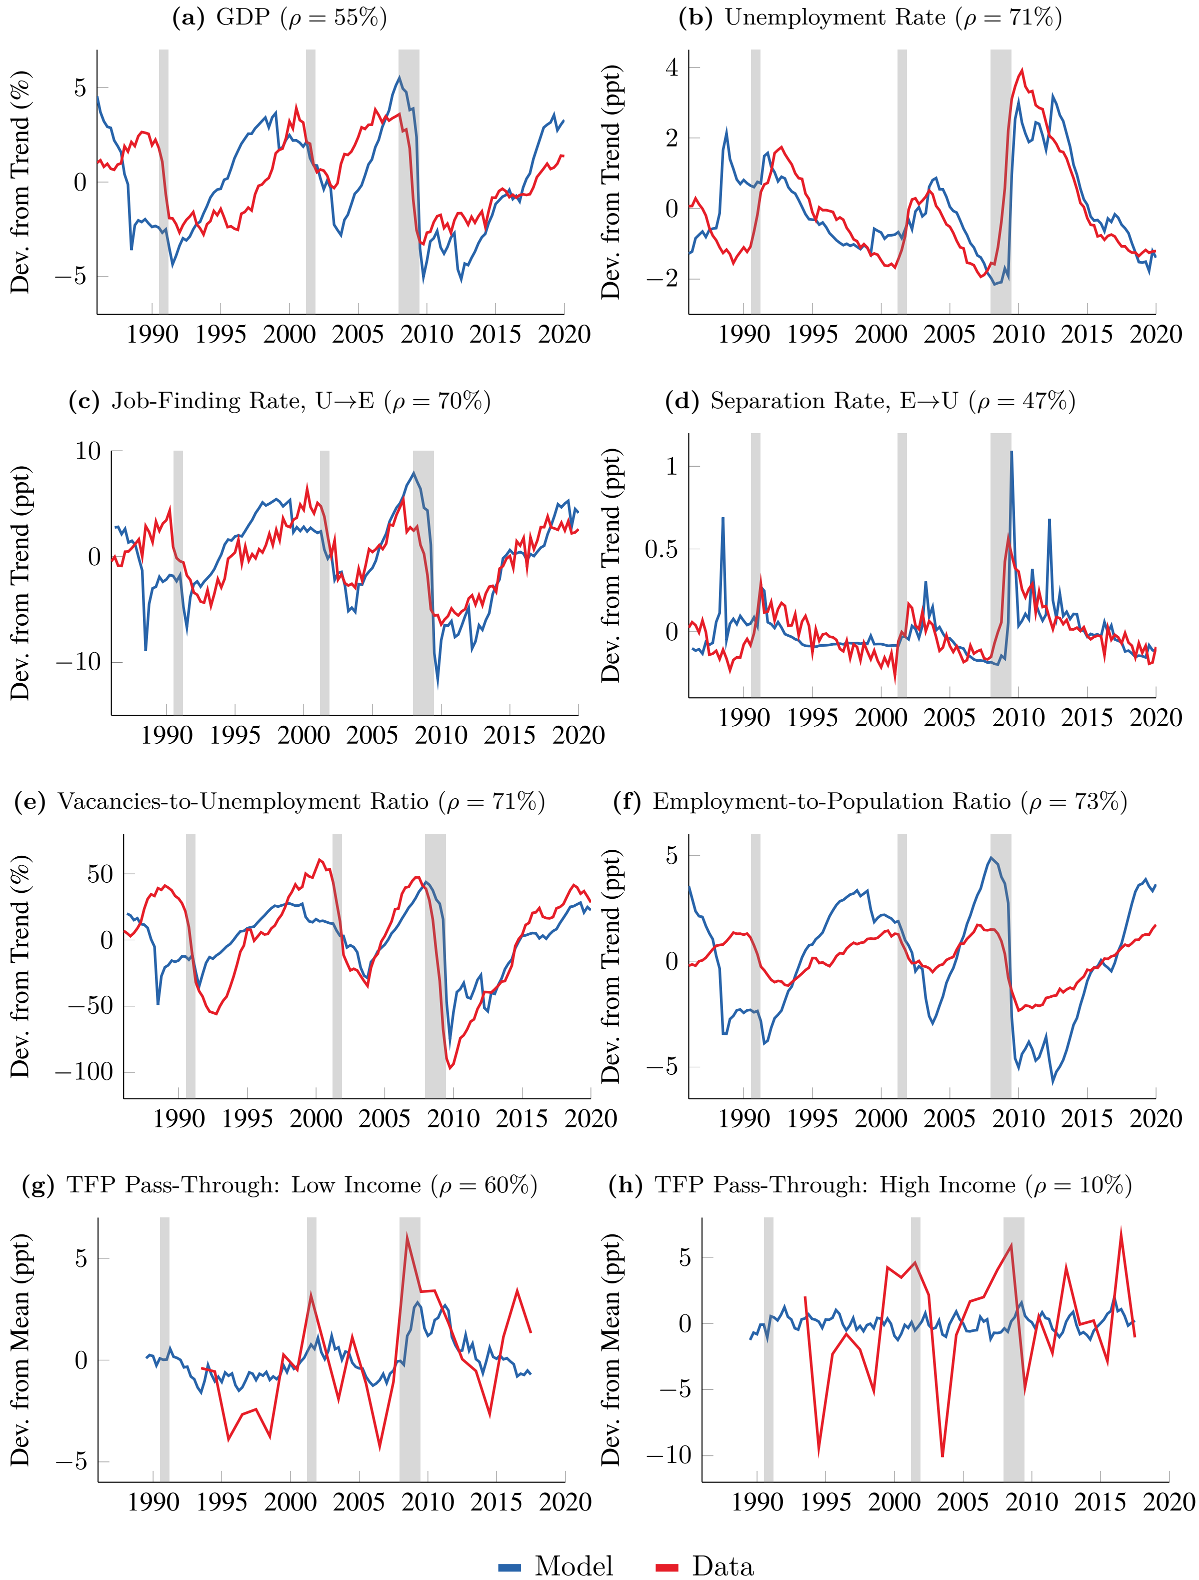

This figure compares the realized paths of key variables between the model and the data. We directly feed into the model our (scaled) empirical measures of risk premium shocks ϵ^{rp}. We detrend all series using an HP filter with quarterly smoothing parameter 10^5.

Figure 7: Pass-Through of Shocks to Workers in Model

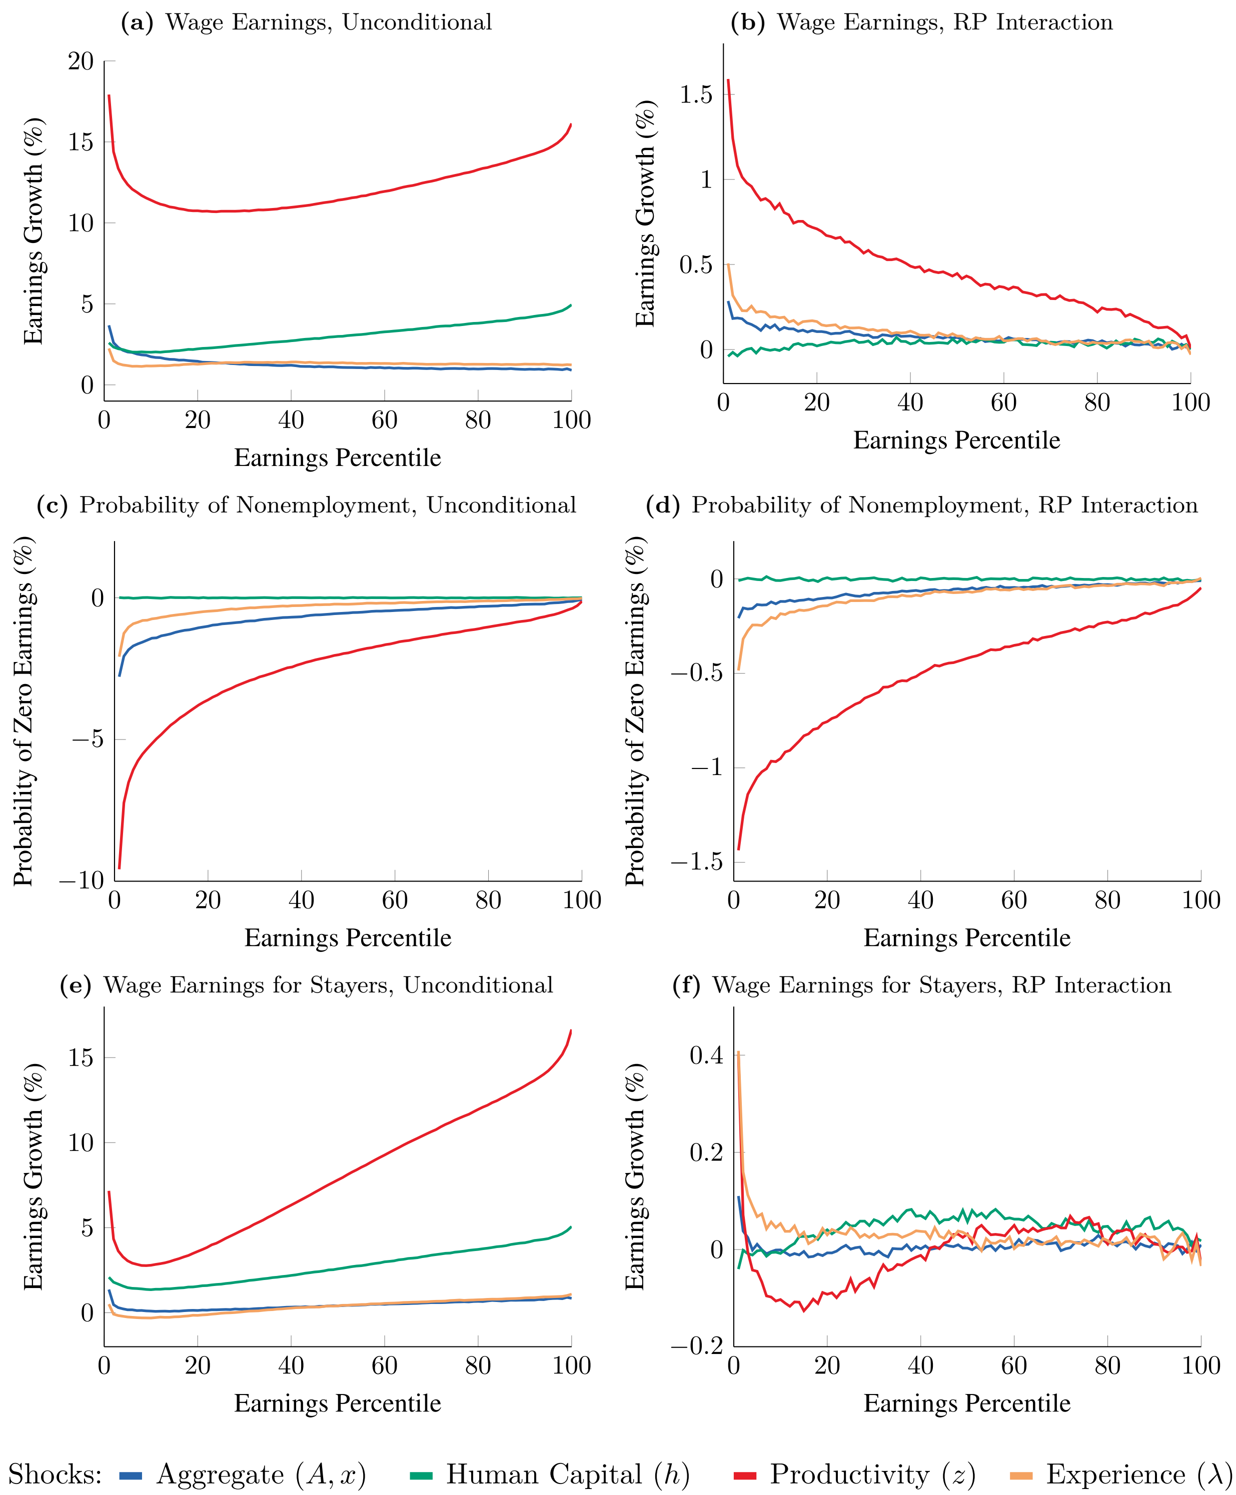

Figure 8: Impulse Responses to Aggregate Shock in Model

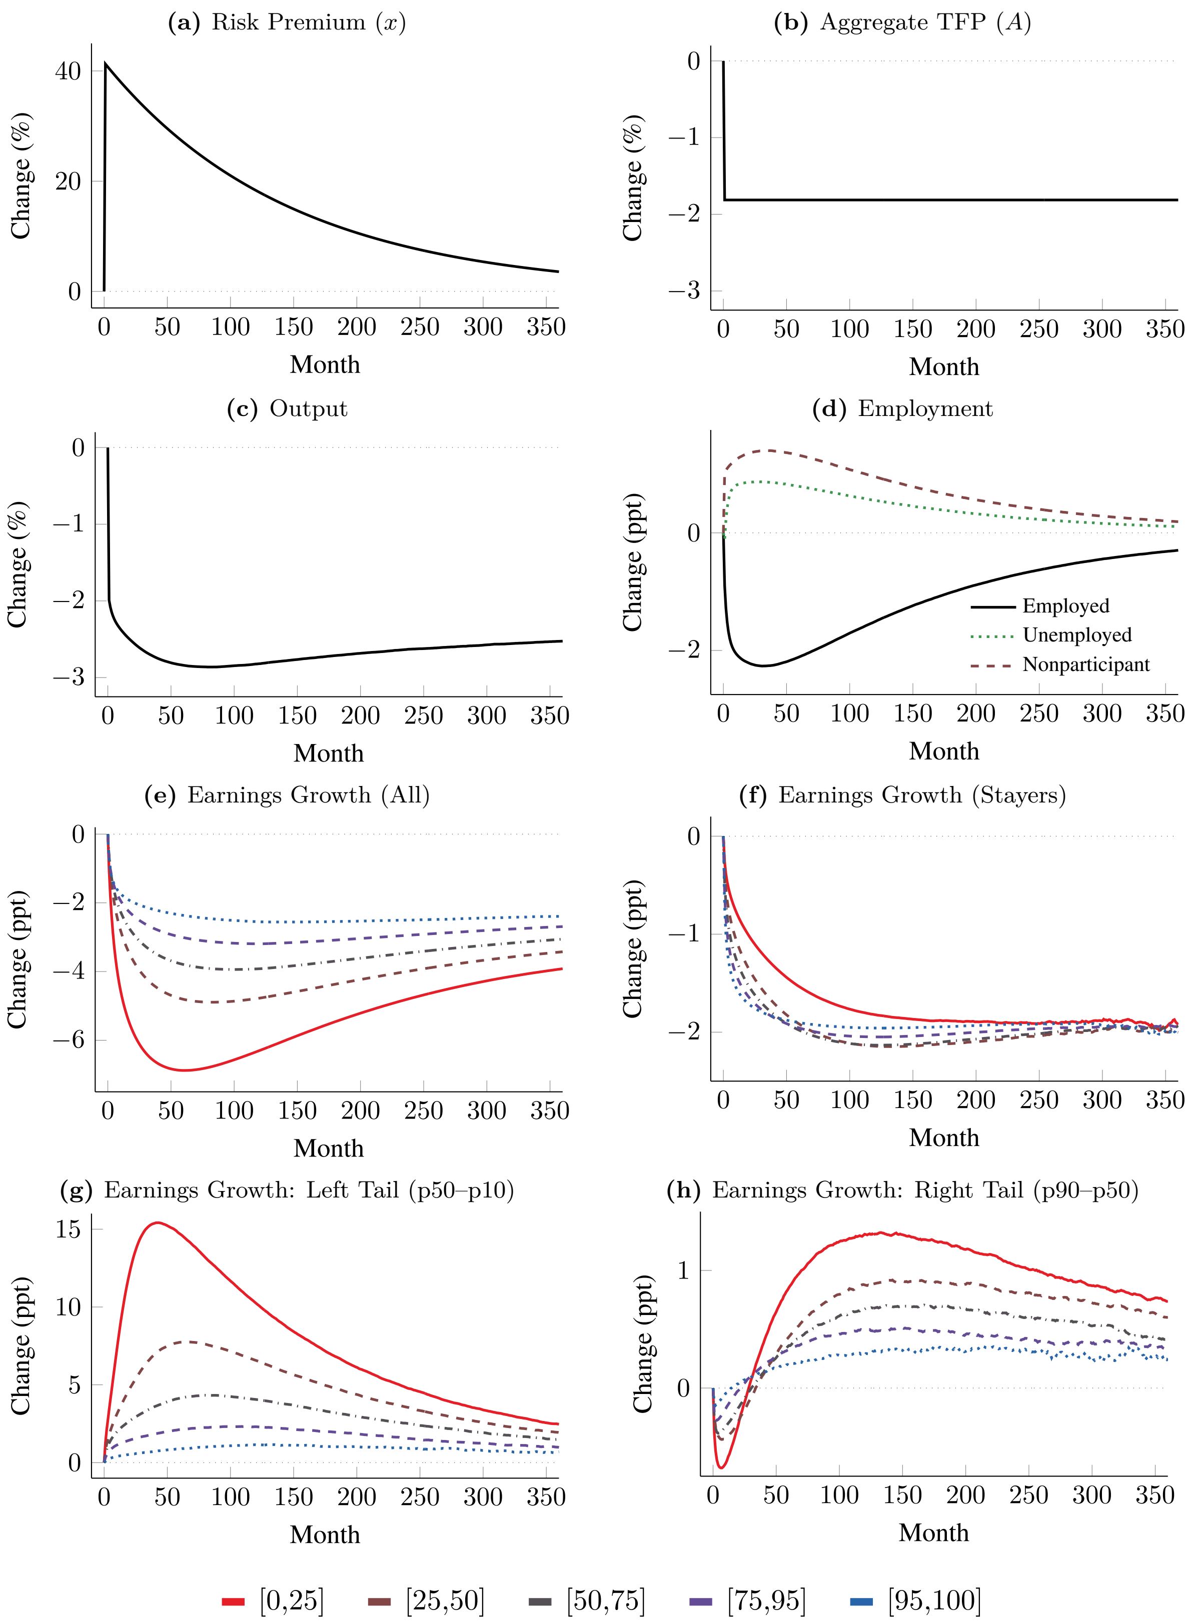

This figure shows the impulse responses of key model quantities following an aggregate shock of one annual standard deviation.

Figure 9: Labor Income Risk: Model vs. Data

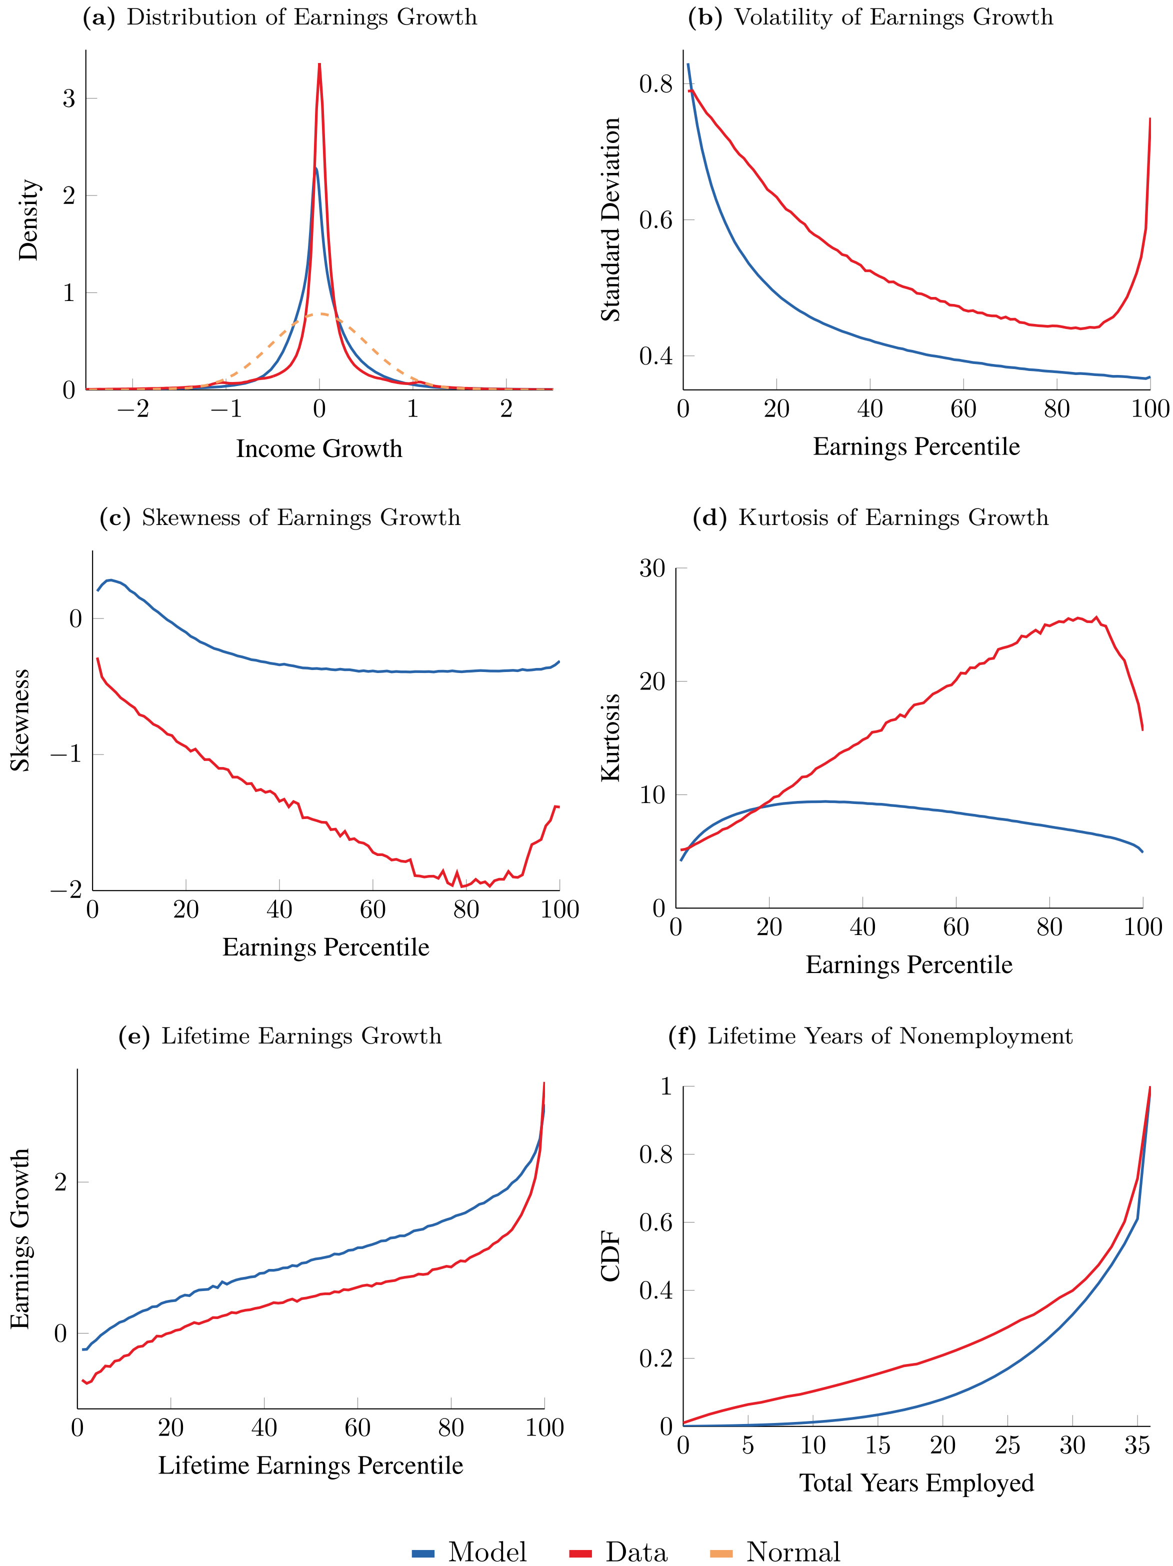

Figure 10: Realized Fluctuations in Earnings Risk: Model vs. Data

(a) Left Tail, p50–p10 ($\rho = 58\%$)

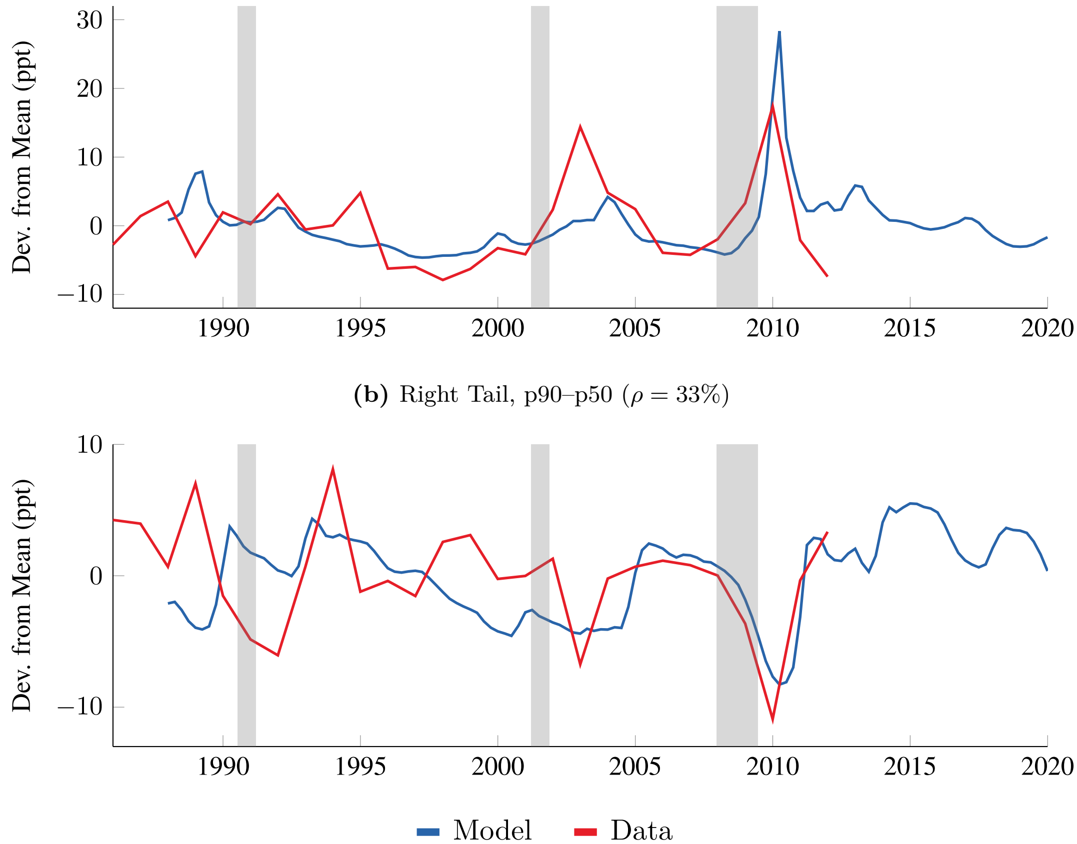

This figure compares the realized paths of key variables between the model and the data. We directly feed into the model our (scaled) empirical measures of risk premium shocks ϵ^{rp}.

Figure 11: Welfare and Value of Human Capital in Model

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

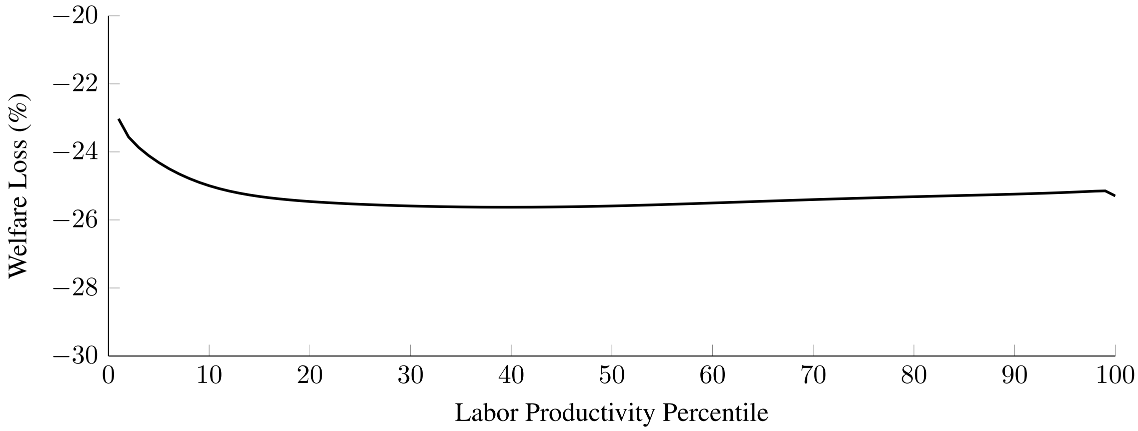

(b) Human Capital to Consumption Ratio by Labor Productivity (All Workers)

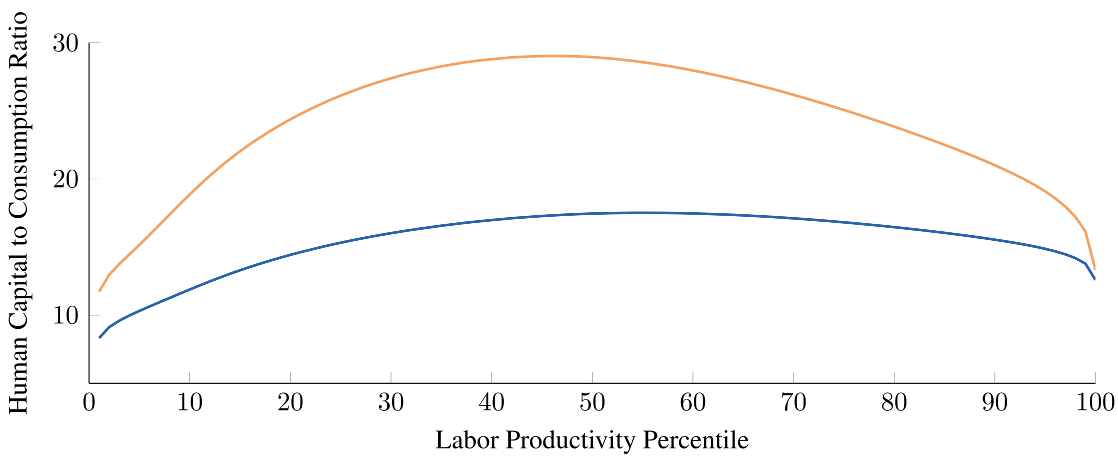

(c) Human Capital to Consumption Ratio by Wage (Incumbent Workers)

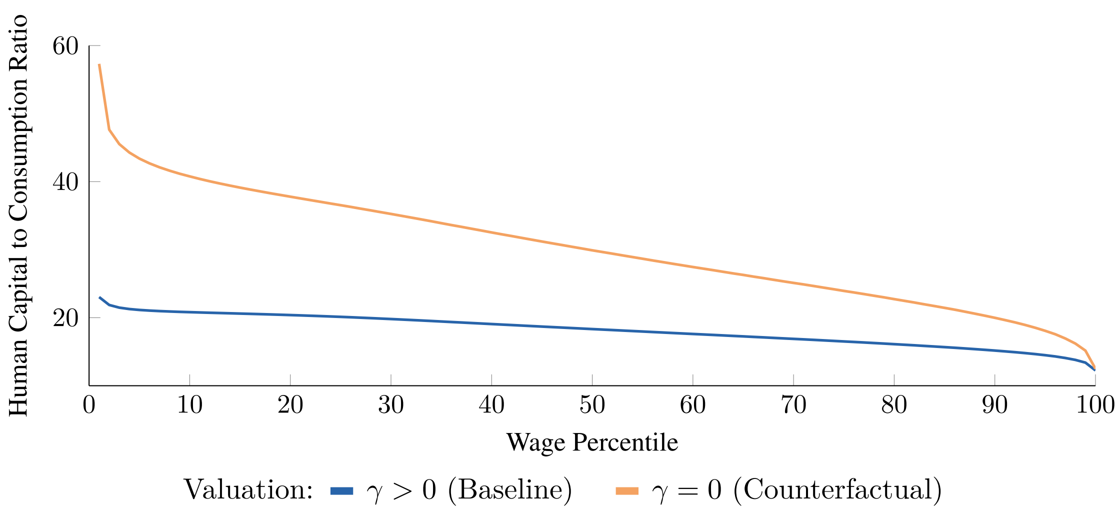

Table 1: Heterogeneity in Pass-Through of Firm TFP Growth

	(1)	(2)	(3)	(4)	(5)	(6)
ϵ^{tfp}						
$\times$ 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)
$\times$ 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)
$\times$ 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)
$\times$ 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)
$\times$ 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)
$\epsilon^{tfp} \times \epsilon^{rp}$						
$\times$ Worker Earnings, 0–25th Percentile		14.33		11.52	5.76	5.61
		(3.30)		(11.46)	(4.93)	(4.42)
$\times$ Worker Earnings, 25–50th Percentile		9.44		7.52	1.67	1.98
		(2.44)		(9.28)	(1.76)	(1.90)
$\times$ Worker Earnings, 50–75th Percentile		7.46		5.92	0.82	0.81
		(1.68)		(7.80)	(0.92)	(0.83)
$\times$ Worker Earnings, 75–95th Percentile		5.23		4.25	-2.05	-2.59
		(0.94)		(4.86)	(-1.99)	(-2.30)
$\times$ Worker Earnings, 95–100th Percentile		-2.29		-1.44	-10.08	-10.27
		(-0.31)		(-0.67)	(-3.94)	(-3.74)
Controls:						
Earn Grp $\times\, \epsilon^{rp}$	✓	✓				
Earn Grp $\times\, \epsilon^{tfp}$ $\times$ GDP Growth					✓	
Earn Grp $\times\, \epsilon^{tfp}$ $\times$ NBER Recession						✓
Fixed Effects:						
NAICS2 $\times$ Age $\times$ Gender	✓	✓	✓	✓	✓	✓
NAICS2 $\times$ Earn Grp	✓	✓				
NAICS2 $\times$ Earn Grp $\times$ 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	2.78
B. Parameters Calibrated to Asset and Labor Markets	**Symbol**	**Value**
Time preference parameter	β	0.999
Average dividend growth (%)	μ_E	0.10
Persistence of price of risk	ψ_x	0.993
Average price of risk	$\bar{x}$	0.074
Volatility of price of risk (%)	σ_x	11.91
Vacancy posting cost, scale ($\times$ 100)	$\bar{\kappa}_0$	1.66
Vacancy posting cost, elasticity to z	$\bar{\kappa}_1$	1.51
Exogenous separation rate (%)	s	0.70
Unemployment benefit, intercept	$\bar{b}_0$	2.19
Unemployment benefit, dependence on z	$\bar{b}_1$	0.50
Nonparticipation benefit	$\bar{n}$	2.50
Growth of h in nonemployment (%)	g_O	0.04
Volatility of h (%)	σ_h	3.01
High value of λ	λ_H	2.49
Worker risk aversion	γ	0.48
On-the-job search, intercept	$\bar{\chi}_0$	48.6
On-the-job search, dependence on z	$\bar{\chi}_1$	2.43
Correlation of h in firm (%)	ρ_h	1.98
Correlation of z in firm (%)	ρ_z	24.7
Volatility of orthogonal firm TFP shocks (%)	σ_k	6.50

This table reports the parameter values in our baseline calibration of the model. Model is calibrated at a monthly frequency. See Section 3.1 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 (%)	1.40	1.44	0.94	0.97	1.00	1.00
Long-term unemployment share (%)	6.05	5.78	0.76	0.97	2.47	3.45
Employment-to-population ratio (%)	2.48	1.08	0.96	0.97	-1.60	-0.72
Labor force participation rate (%)	1.60	0.35	0.95	0.91	-0.85	-0.07
Labor market tightness (log V/U ratio, %)	21.68	37.71	0.91	0.97	-13.42	-25.32
B. *Job Flows*						
Job-finding rate (%)	3.69	2.93	0.80	0.92	-1.95	-1.91
Separation rate into unemployment (%)	0.14	0.17	0.48	0.83	0.05	0.10
C. *Decomposition of Unemployment Rate*						
Unemployment rate w/ constant separations (%)	0.75	0.79	0.96	0.97	0.51	0.51
Unemployment rate w/ constant job finding (%)	0.44	0.61	0.84	0.94	0.23	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.

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

We construct the risk premium shock as the first principal component of the AR(1) residuals of each individual series. We follow Bauer et al. (2023) in dealing with missing observations to obtain a complete time series. The resulting series is highly positively correlated with each component, with a minimum correlation of 51% and an average correlation of 75%.

A.4 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.5 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.6 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

Here, we show an example of how our worker preferences can arise if workers have time-varying ambiguity over aggregate risks.

Proposition 1. *Agents who face ambiguous uncertainty over aggregate risk but perfect clarity over uncertainty in idiosyncratic risk face the following Epstein-Zin value function:*

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

where

$$\log(\Gamma_{t+1}) = -\frac{1}{2}x_t^2 - x_t\,\varepsilon_{A,t+1} \tag{A.4}$$

Proof. The agent faces Knightian uncertainty about the aggregate shock but faces regular uncertainty about their idiosyncratic risk. To capture ambiguity aversion, we want to find the functional form of some distortion to the continuation value. With Epstein-Zin preferences, we have the expression:

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

Typically, we set $\mathcal{R}_t(V_{i,t+1}) = E_t[V_{i,t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}}$. However we want to shift preferences to reflect time-varying attitudes towards ambiguity. We can impose an entropic penalty on utility. In Skiadas (2013b) it can be shown that this entropic penalty is equivalent to a multiplicative distortion. This multiplicative distortion is in line with the axioms from Skiadas (2013a) that states that the functional form of ambiguity-aversion must be time-varying and source dependent while preferences remain scale invariant. Specifically, we set

$$\tilde{\mathcal{R}}_t(V_{i,t+1}) = E_t[\Gamma_{t+1}V_{i,t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}} \tag{A.6}$$

We will impose a series of restrictions on Γ_t. The first restriction comes from the fact that we want the preference distortion to be time-varying. In order for the distortion to be time-varying, we require that the functional form depends on aggregate shocks. The second restriction is that as a distortion, $E_t[\Gamma_{t+1}] = 1$ must also hold. The third follows from Skiadas (2013a) and states that the ambiguity aversion must be source dependent. Since only aggregate risks are knightian, we want Γ_t to only move with the aggregate state. The last comes from Skiadas (2013b). Suppose we define $KL_t(Q||P)$ as the conditional Kullback–Leibler divergence or relative entropy of our alternative one-period distortion Q relative to the actual probability distribution. We can show that

$$KL_t(Q||P) = E_t^Q\left[\log\frac{\partial Q_{t+1}}{\partial P_{t+1}}\right] = \frac{1}{2}x_t^2 \qquad \text{where} \qquad \frac{\partial Q_{t+1}}{\partial P_{t+1}} = \Gamma_{t+1} \tag{A.7}$$

and thus

$$E_t^Q[\log\Gamma_{t+1}] = E_t^P[\Gamma_{t+1}\log\Gamma_{t+1}] = -\frac{1}{2}x_t^2 \tag{A.8}$$

The above implies that the expectation of the log of the multiplicative distortion over the distorted distribution is equal to $-\frac{1}{2}x_t^2$. The four restrictions can be summarized as follows:

1. Γ_{t+1} is a function of $\mathcal{F}_t$

2. $E_t[\Gamma_{t+1}] = 1$

3. $\Gamma_{t+1} \not\perp$ agg. risk, but $\Gamma_{t+1} \perp$ idio. risk

4. $E_t^Q[\Gamma_{t+1}] = \frac{1}{2}x_t^2$

Our specification (A.4) satisfies all four of these restrictions. Here, we note that the conditional ambiguity of Epstein and Schneider (2003) makes the next period set of conditional probabilities rectangular and state-dependent. From the mean-reverting random walk process of the state-dependent x_t, we see that the mean and conditional variance of the ambiguity are also state dependent. Under Gaussian innovations, the worst-case one-period distortion is an exponential tilt, which produces the preference shifter in the form of (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.

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.

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

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.

Haltiwanger, J., R. S. Jarmin, R. Kulick, and J. Miranda (2017). High growth young firms: Contribution to job, output, and productivity growth. In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*. University of Chicago Press.

Heathcote, J., F. Perri, and G. L. Violante (2020). The rise of US earnings inequality: Does the cycle drive the trend? *Review of Economic Dynamics 37*, S181–S204.

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

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.

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: Pass-Through of Worker Productivity Shock in Stylized Model

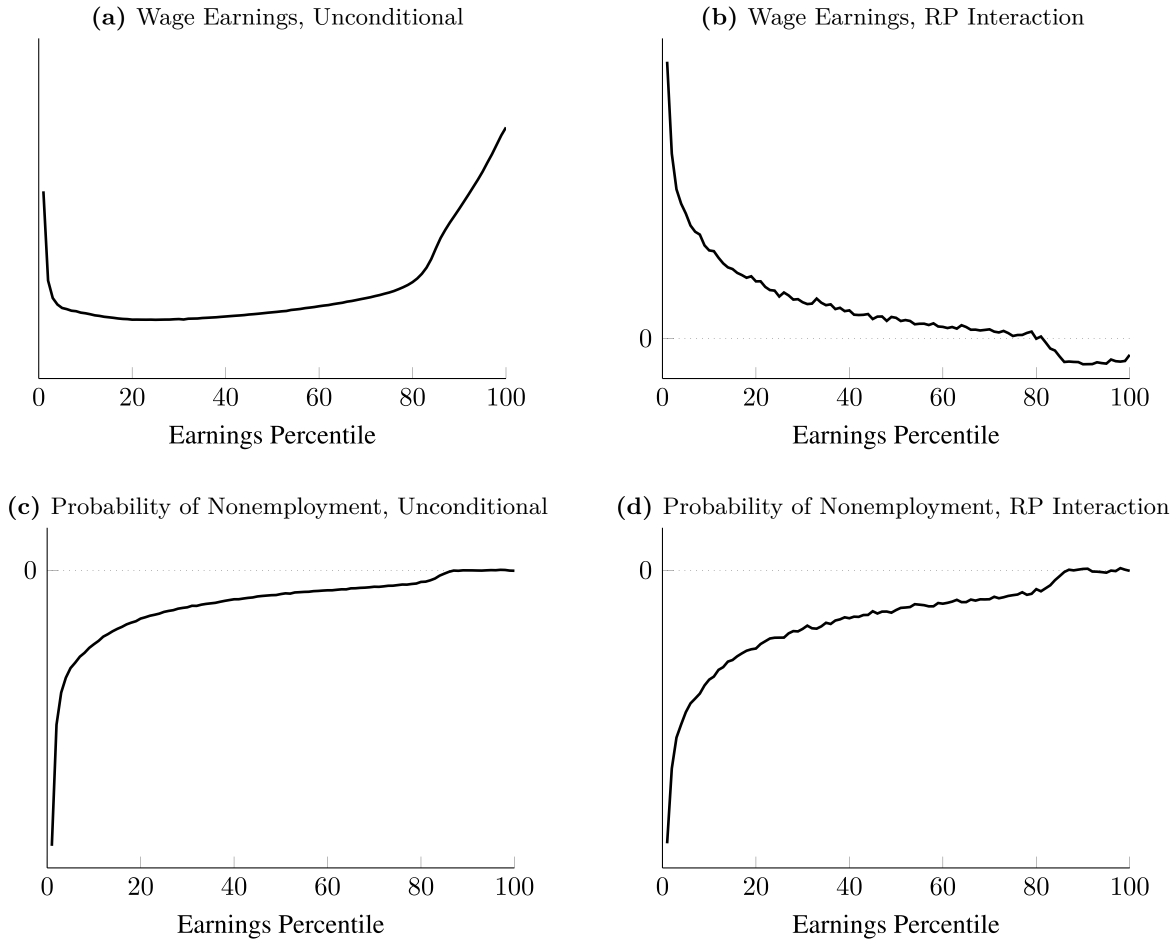

Figure A.2: Separation and Job-Finding Rates by Worker Income: Model vs. Data

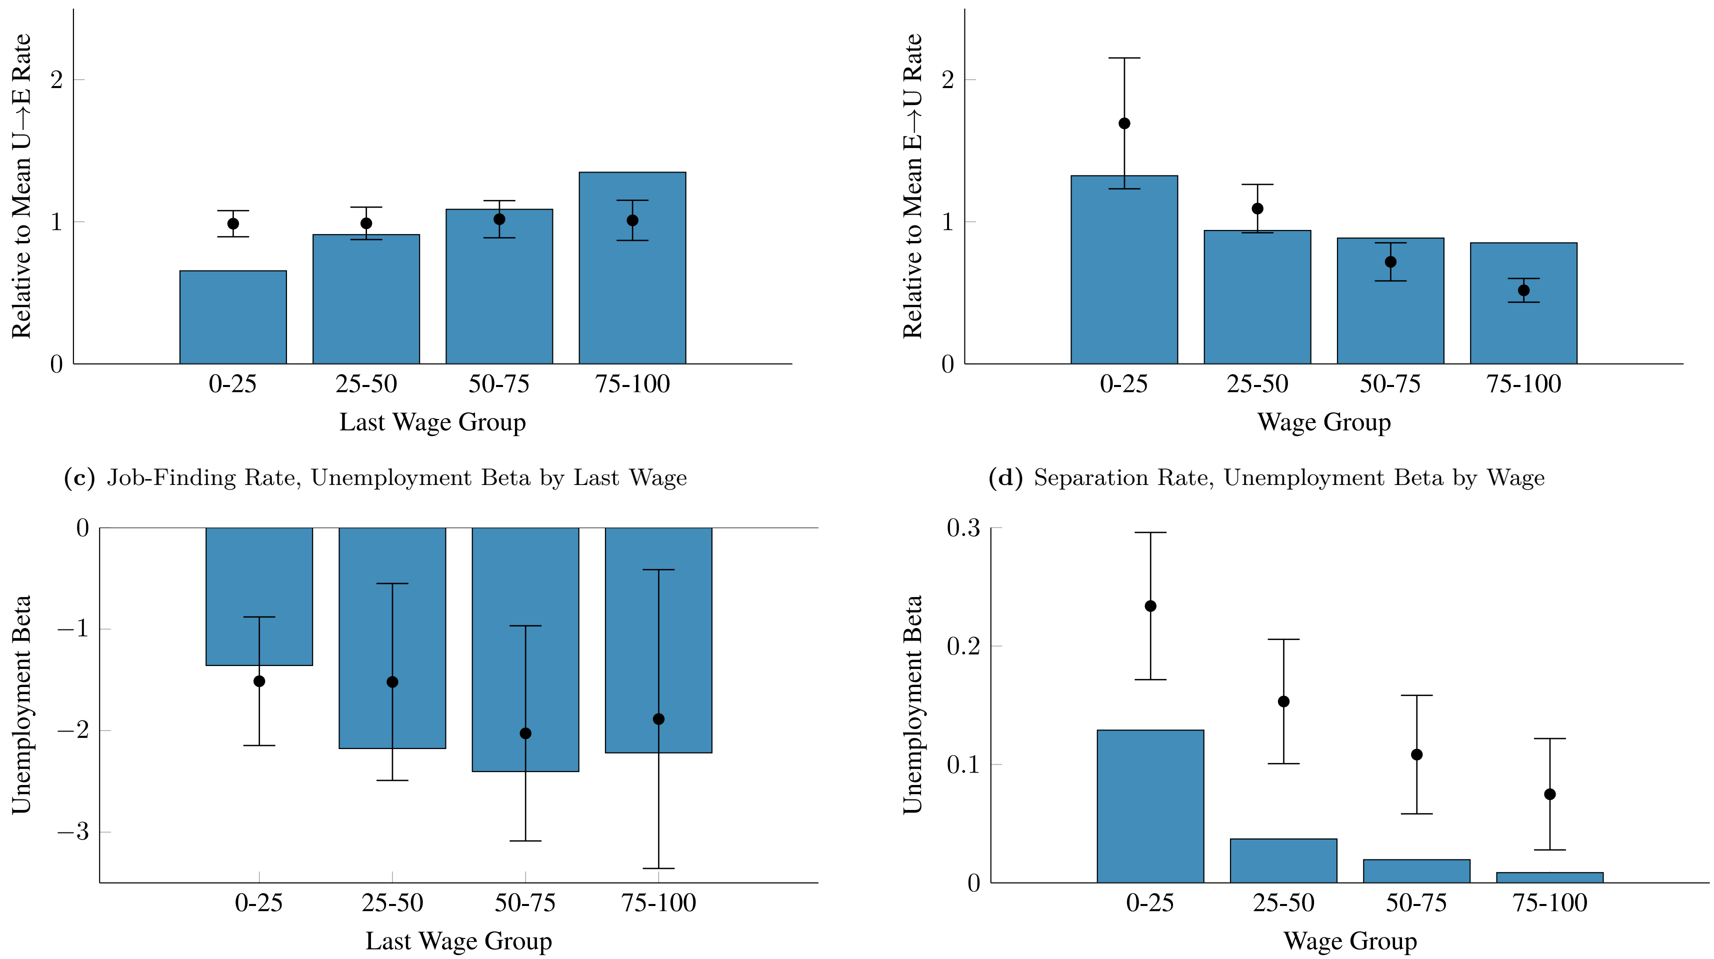

(a) Job-Finding Rate, Average by Last Wage

(b) Separation Rate, Average by Wage

(c) Job-Finding Rate, Unemployment Beta by Last Wage

(d) Separation Rate, Unemployment Beta by Wage

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.3: Response of Wage Earnings, Stayers vs. Movers: Model vs. Data

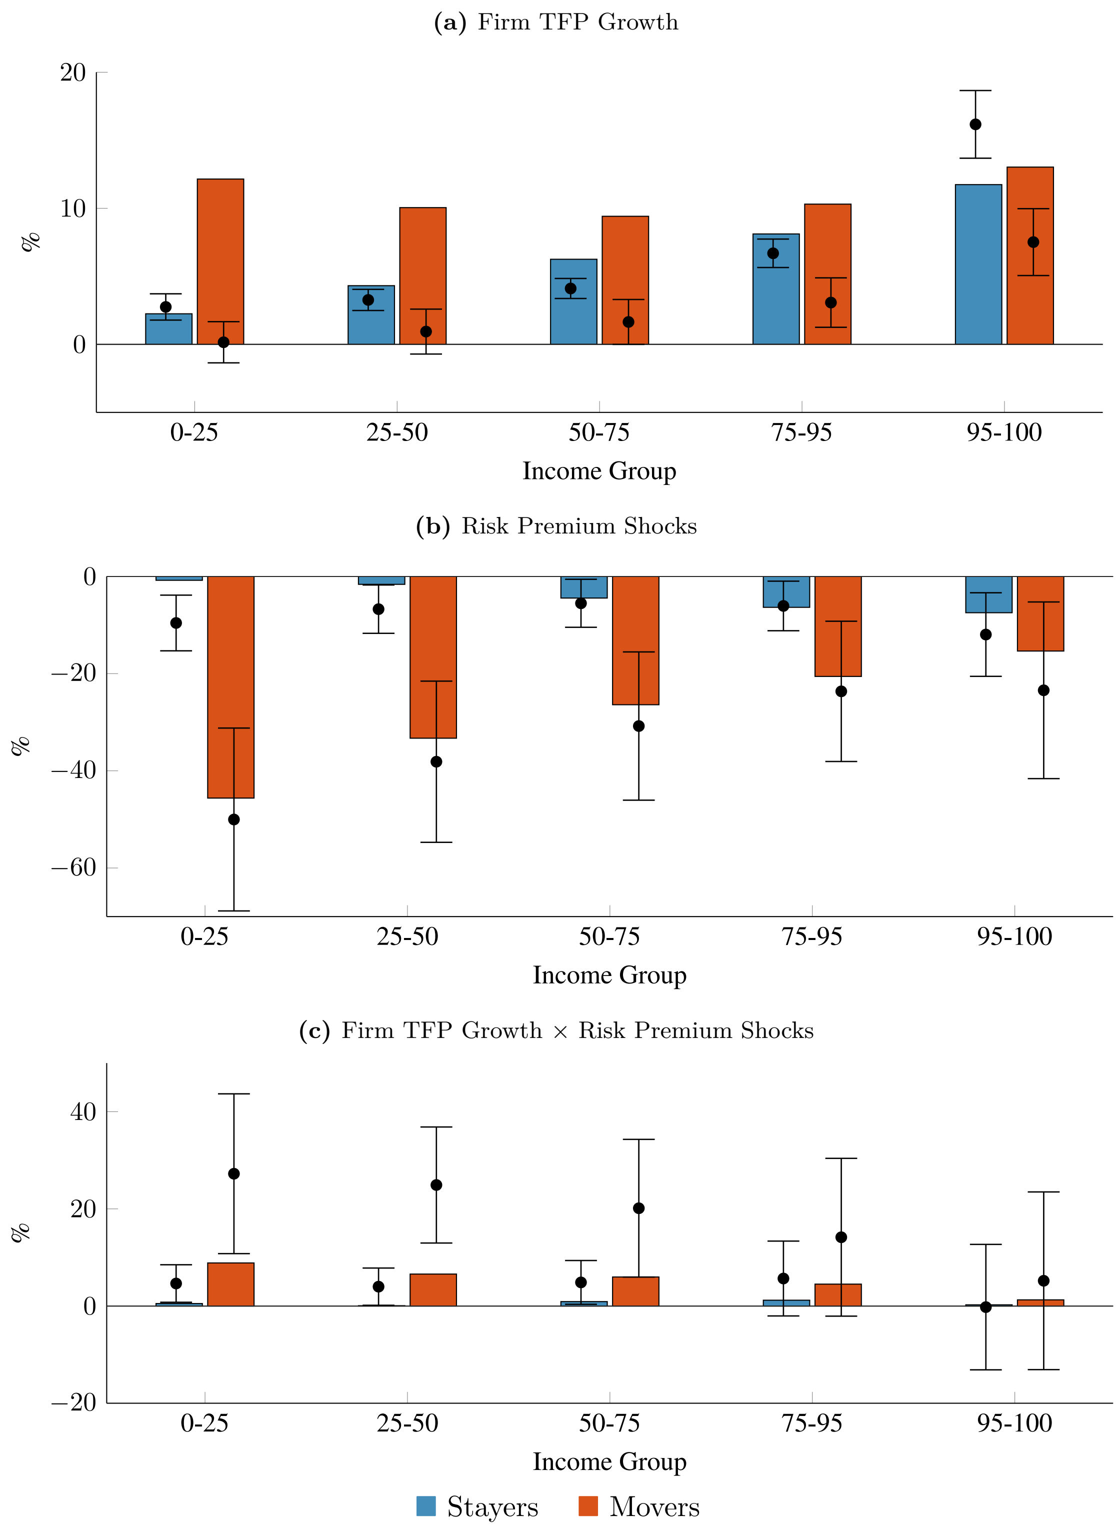

Figure A.4: Log Density of Earnings Growth: Model vs. Data

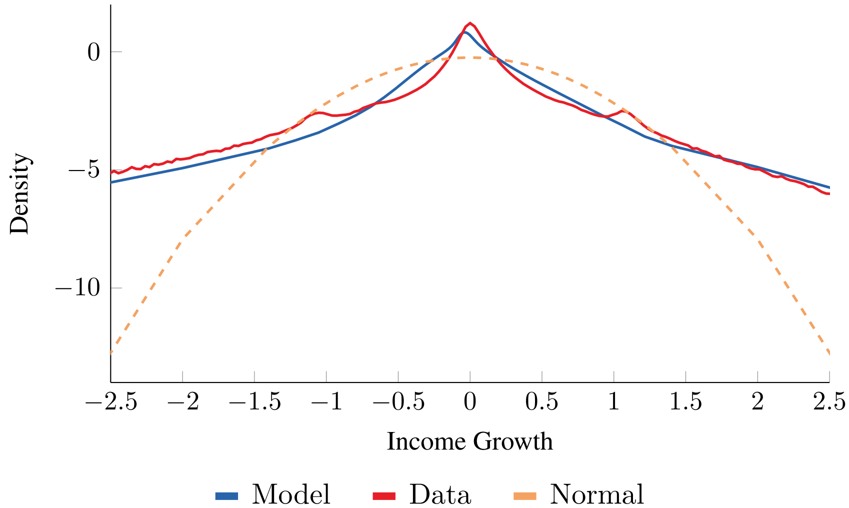

Figure A.5: Labor Income Risk, Stayers vs. Movers: Model vs. Data

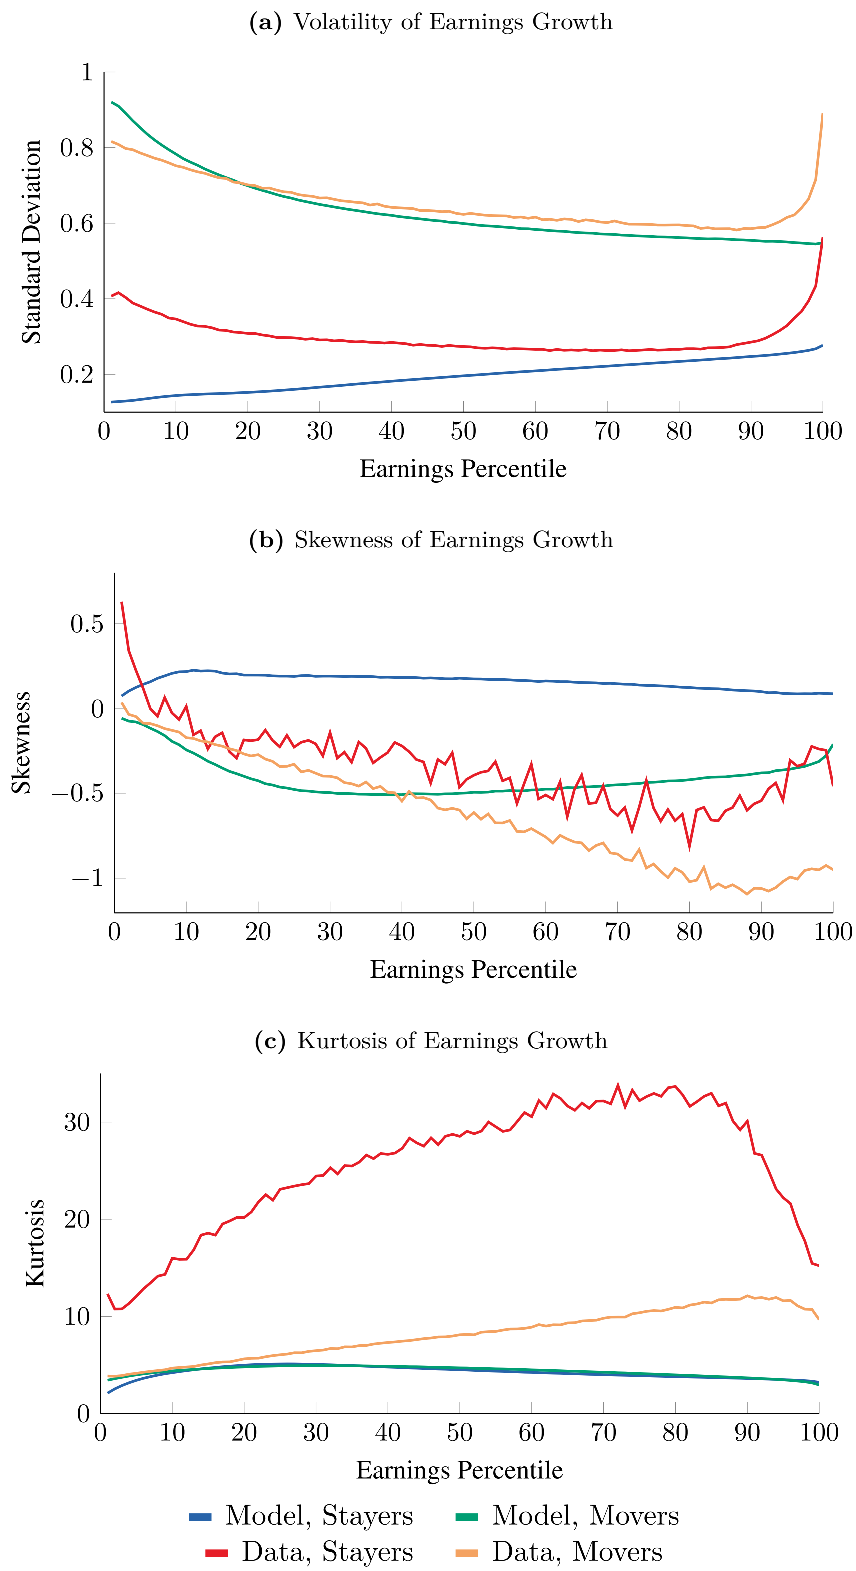

Table A.1: Pass-Through of Firm TFP Growth across Horizons: Stayers versus Movers

	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