Employment Stability, Earnings Dynamics, and Life-Cycle Savings*

Moritz Kuhn† Leanne Nam‡ Gašper Ploj§

March 26, 2026

Abstract

Labor markets feature large heterogeneity in employment stability. Some careers provide lifetime employment with rare separations, while others involve frequent transitions into and out of employment. Such heterogeneity shapes earnings dynamics and labor market risk but implications for households' financial decisions remain poorly understood. We first document two new empirical facts: careers of more stable employment (i) exhibit steeper life-cycle earnings profiles and (ii) are associated with more wealth accumulation per dollar of income. To interpret these facts, we then develop a life-cycle search-and-saving model that embeds heterogeneous employment risk, job-to-job mobility, endogenous human capital and wealth accumulation. The model matches the earnings and wealth life-cycle differences across heterogeneous careers consistent with the empirical evidence. Differential earnings growth across careers accounts for 60% of the differences in wealth accumulation, with consumption smoothing dynamics accounting for the remainder. We demonstrate the importance of heterogeneity in employment stability for the macroeconomy by studying the consequences for earnings losses and wealth inequality.

Keywords: Employment risk, employment stability, consumption-saving behavior
JEL: J64, E21, E24

*We thank Christian Bayer, Richard Blundell, Thomas Hintermaier, Philip Jung, Per Krusell, Iourii Manovskii, Kurt Mitman, Giuseppe Moscarini, Aysegul Sahin, and Abigail Wozniak for very helpful comments and remarks. We also thank seminar and workshop participants at all institutions and workshops where the work has been presented for their comments and remarks to improve the paper. Kuhn thanks the DFG for financial support (DFG No. 433368816). Financial support from the DAAD, the Bonn Graduate School of Economics, and the DFG research training group 2281 "The Macroeconomics of Inequality" is gratefully acknowledged. This paper should not be reported as representing the views of the Banka Slovenije. The views expressed are those of the author.

†U Mannheim & Macro Inequality Lab, CEPR, IZA, mokuhn@uni-mannheim.de, D-68161 Mannheim
‡U Mannheim & Macro Inequality Lab, leanne.nam@uni-mannheim.de, D-68161 Mannheim
§Banka Slovenije, gasper.ploj@bsi.si, 1000 Ljubljana, Slovenia

1 Introduction

Labor markets feature large heterogeneity in employment stability. Some careers provide lifetime employment with rare separations, while others involve frequent transitions into and out of employment. This salient fact has been documented at least since Hall (1982), yet its implications for households' financial decisions remain poorly understood, leaving a gap in the existing literature. We address this gap by revisiting the classic question of how labor-market risk shapes saving behavior within a novel framework of endogenous and heterogeneous earnings and wealth dynamics. We show that heterogeneity in employment stability is a first-order determinant of wealth accumulation, as it shifts the dominant savings motive between life-cycle saving and precautionary saving. We also demonstrate that heterogeneity in employment stability is important from a macroeconomic perspective, as it shapes wealth inequality and the consumption response to changes in unemployment.

Empirically, we document substantial heterogeneity in employment stability using the *Panel Study of Income Dynamics* (PSID) and show that it is systematically related to earnings growth and wealth accumulation. More stable careers are associated with higher earnings growth and greater wealth accumulation, even after controlling for income differences. Life-cycle profiles reveal large differences: by the end of prime working life, stable careers are associated with about 20% higher earnings and wealth equivalent to roughly one and a half years of income. These empirical patterns point to a tight link between employment stability, earnings dynamics, and saving behavior. In particular, they suggest that differences in employment stability are associated not only with differences in earnings levels, but also with differences in earnings growth and the way households accumulate wealth over the life cycle. We confirm this relationship using independent data from the *Survey of Consumer Finances* (SCF).

To interpret these empirical patterns, we turn to theory. Heterogeneity in employment stability affects savings behavior through two opposing forces. Higher earnings growth along stable careers strengthens the life-cycle savings motive, while higher employment risk along unstable careers increases the precautionary savings motive. To study the interaction of these forces and derive their quantitative implications for wealth accumulation, we develop a new search-and-saving model combining endogenous earnings and savings dynamics.

The new search-and-saving model combines a life-cycle labor search model with heterogeneity in employment stability and endogenous human capital investment with a consumption-saving model under incomplete markets. We interpret the empirical heterogeneity in employment stability as arising from heterogeneity in job stability, in line with a large empirical and structural literature (e.g., Hall, 1982; Pinheiro and Visschers, 2015a; Jarosch, 2023; Jung and Kuhn, 2018; Larkin, 2019; Bilal, 2019; Kuhn et al., 2026 forthcoming).[1] In the model, workers search both on and off the

[1]Worker characteristics (e.g., education) also contribute to heterogeneity in employment stability. Our focus is

job, and jobs differ in wages and separation rates (Pinheiro and Visschers, 2015a; Jarosch, 2023; Larkin, 2019; Nam, 2022). Human capital can be accumulated only while employed, generating a wedge between human capital accumulation dynamics on and off the job. We calibrate the model to jointly match average life-cycle labor market dynamics, earnings growth, and wealth profiles. The model also matches untargeted empirical facts on consumption inequality (Aguiar and Hurst, 2013), earnings dynamics (Topel and Ward, 1992; Blundell et al., 2008), earnings losses after job displacement (Jacobson et al., 1993), and the distribution of annual earnings growth (Guvenen et al., 2021a). These validations demonstrate that the model provides an informative quantitative framework to interpret the empirical correlation between employment stability, earnings growth, and wealth accumulation.

The search-and-saving model explains the tight empirical relationship between employment stability, earnings growth, and wealth accumulation. On a stable career path, workers do not fall off the job ladder, improve their jobs when opportunities arise, and invest continuously in human capital, jointly leading to steeper life-cycle earnings profiles and a stronger life-cycle savings motive. Intuitively, workers with stable careers save for when their career is over and smooth life-cycle variation in income.

By contrast, along an unstable career path, where a career does not take off, there is no sustained earnings growth and hence little to smooth over the long run, leaving little incentive to accumulate life-cycle savings. Instead, unstable careers involve repeated falls off the job ladder and employment interruptions that prevent human capital accumulation. Limited earnings growth mutes the life-cycle motive and shifts savings toward precautionary motives aimed at smoothing income fluctuations arising from employment instability. Workers accumulate assets while employed and decumulate them during non-employment, generating a recurring cycle of precautionary savings in which a buffer stock repeatedly fills up and empties out.

Taken together, heterogeneity in employment stability jointly determines earnings risk and earnings growth, thereby shaping the nature of the household saving problem. It shifts the dominant savings motive from life-cycle saving in stable careers to precautionary saving in unstable careers, leading to systematic differences in wealth accumulation.

We demonstrate that heterogeneity in employment stability also has important macroeconomic consequences. It shapes savings behavior and, in turn, the distribution of wealth, and amplifies consumption responses to unemployment shocks. Abstracting from this heterogeneity, as in a representative-worker benchmark, leads to lower wealth inequality and substantially smaller responses to labor market shocks. Nam (2022) further shows that these differences also affect the

on within-worker-group heterogeneity, which we document to be quantitatively as large as the overall dispersion of employment stability. As a robustness exercise, we extend the model to include empirically disciplined worker-type heterogeneity and show that the key model predictions are robust (Section 4.3). Appendix B further discusses heterogeneity in employment stability across workers and jobs in the data.

scope for insurance policies.

The calibrated model matches the untargeted heterogeneity in earnings growth and wealth accumulation between stable and unstable careers. We use the model to decompose the sources of this heterogeneity. For earnings, unstable employment leads to lower human capital accumulation and, as a result, flatter life-cycle earnings profiles. Differences in human capital account for about half of the observed earnings growth gap, while the remainder reflects lower wages along a slippery job ladder in unstable careers. This decomposition provides a micro-foundation for the "stepping-stone versus dead-end" view of careers: stable careers act as stepping stones that facilitate human capital investment and sustained earnings growth, whereas unstable careers resemble dead-end jobs with low earnings, limited career progression, and elevated separation risk.

For wealth accumulation, the model explains the between-career patterns through a shift in the dominant savings motive from life-cycle saving to precautionary saving. In stable careers, rising earnings over the life cycle give rise to a strong life-cycle savings motive, as households accumulate assets to smooth their concave earnings profile. With stable employment, there is no need to decumulate assets due to non-employment spells, so savings dynamics reflect the smoothing of a steeper life-cycle earnings profile. We refer to such behavior as "Modigliani savers," as these careers induce a quasi-deterministic life-cycle consumption-saving problem.

By contrast, in unstable careers, limited earnings growth mutes the life-cycle savings motive, as there is little to smooth over the life cycle when careers do not generate upward-sloping earnings. Instead, savings shift toward precautionary motives to insure against recurring income fluctuations arising from employment instability, giving rise to a buffer-stock saving strategy. While employed, workers accumulate assets to prepare for the next non-employment spell; once displaced, they decumulate these assets to smooth consumption. Because most new job opportunities are unstable and good jobs are hard to find, this process can repeat itself over the entire life cycle, giving rise to unstable career paths. The resulting wealth dynamics take the form of a buffer stock that repeatedly fills up and empties out, and this recurring accumulation and decumulation of assets limits long-run wealth accumulation.

Decomposing the differences in wealth-to-income ratios, we find that earnings growth (life-cycle saving) explains 60% of the observed gap, while the remaining 40% is driven by these consumption-smoothing dynamics, which on net even reduce long-run wealth growth in unstable careers. Thus, the model shows how labor market dynamics intertwine earnings growth and earnings volatility in a way that accounts for the observed wealth accumulation patterns across these career paths.

We use the model to quantify how these findings matter for the macroeconomic consequences of unemployment. From a macroeconomic perspective, separations from stable careers generate large and persistent earnings losses, precisely the type of losses that amplify the costs of business cycles (Krebs, 2007). By contrast, in a representative-worker benchmark that abstracts from heterogeneity

in employment stability, earnings losses are small and transitory (Jung and Kuhn, 2018; Jarosch, 2023). Analyzing a deep economic downturn, modeled as a spike in job displacement that raises unemployment by 5 percentage points—consistent with episodes such as the 1980s recession, the Global Financial Crisis, and Covid-19, we find that consumption responses are almost three times larger in the model with employment stability heterogeneity than in the representative-worker benchmark. This amplification reflects that employment stability heterogeneity generates careers with large and persistent consequences of job loss, whereas in the representative-worker benchmark job losses lead only to transitory earnings fluctuations.

These results highlight that not only the magnitude but also the composition of job loss matters for aggregate consumption dynamics. Recessions that disproportionately destroy stable jobs generate larger and more persistent consumption declines than those concentrated among unstable jobs. This distinction is particularly relevant during episodes of structural change, characterized by sectoral reallocation or the decline of traditionally stable industries (Jaimovich and Siu, 2020; Howes, 2022). More generally, representative-worker labor market models understate both the size and persistence of consumption responses, as they abstract from heterogeneity in employment stability that shapes income risk and consumption behavior.

The shift in savings motives associated with employment stability also shapes wealth inequality. Heterogeneity in employment stability lowers average lifetime earnings and (net) saving rates at the bottom of the earnings distribution, consistent with empirical evidence on saving behavior (Dynan et al., 2004). As a result, it amplifies wealth inequality: the 90–10 wealth ratio in the model is 26% higher than in a model without such heterogeneity. To put this magnitude into perspective, it is comparable to the rise observed in the SCF between 1998 and 2022, where the 90–10 ratio increased by 32%.[2]

Related Literature Our work relates to two large strands of literature: models of consumption-saving behavior in the presence of idiosyncratic income risk and market incompleteness (Bewley, undated; Aiyagari, 1994; Huggett, 1993) and models of labor market mobility (Mortensen and Pissarides, 1994; Burdett and Mortensen, 1998). Existing models of consumption-saving behavior or labor market dynamics treat labor market dynamics and consumption-saving choices largely as orthogonal: models of consumption-saving behavior typically consider earnings as the result of an exogenous stochastic process, and models of labor market dynamics typically abstract from human capital investment and consumption-saving decisions. Only recently, a strand of research emerged that combined models of consumption-saving and labor market behavior (Lise, 2012; Krusell et al., 2010, 2017; Hubmer, 2018; Larkin, 2019; Cajner et al., 2020; Kaas et al., 2023). We add to this literature by exploring the consequences of heterogeneity in job stability. Our paper connects the

[2]As we rule out negative wealth in the data, we also drop these observations in the SCF data when computing the 90–10 ratio.

part of the literature that focuses on macroeconomic dynamics, as in Krusell et al. (2010), with microeconomic behavior, as in Lise (2012).

Lise (2012) explores savings behavior and earnings dynamics in an infinite horizon model with on-the-job search and uniform unemployment risk. His model struggles to simultaneously account for observed labor market mobility and earnings dynamics. While Lise (2012) abstracts from human capital dynamics, we corroborate the argument in Jung and Kuhn (2018) and Hubmer (2018) that human capital accumulation is key to account for the life-cycle dynamics of earnings inequality. Our human capital accumulation process is consistent with the recent empirical finding that workers who are employed in more productive firms are able to accumulate more human capital (Acabbi et al., 2024). Hubmer (2018) explicitly incorporates life-cycle dynamics and a consumption-saving decision in his model but does not discuss the model's fit to the empirical counterparts. Michelacci and Ruffo (2015) consider a life-cycle consumption-saving model with human capital investment where the probability of job loss declines with age but abstract from heterogeneity in job stability across workers of the same age. Larkin (2019) demonstrates the macroeconomic consequences of heterogeneity in unemployment risk for the consumption dynamics during the Great Recession. Cajner et al. (2020) extend the model in Krusell et al. (2010, 2017) to a life-cycle setting and explore the consequences of tax changes for labor supply.

Our labor market model builds on Jung and Kuhn (2018), who develop a life-cycle search model to demonstrate that heterogeneity in job stability is key to account for earnings losses following job displacement. Their model is fully microfounded with worker and match heterogeneity in productivity so that wages and job stability are endogenous model outcomes. Jarosch (2023) also highlights the importance of heterogeneity in job stability to account for observed earnings losses. While heterogeneity in job stability can be microfounded as in Jung and Kuhn (2018), we follow Pinheiro and Visschers (2015b) and Jarosch (2023) and introduce this heterogeneity in reduced form to the job-offer distribution.[3] Guvenen et al. (2021a) explore life-cycle earnings dynamics and document large heterogeneity in life-cycle non-employment spells. They emphasize that this heterogeneity is key to account for life-cycle earnings dynamics. Additional evidence for heterogeneity in job stability comes from Morchio (2020), who documents large heterogeneity in unemployment within cohorts of U.S. workers.

Our work also relates to research on heterogeneity in earnings risk, as in Low et al. (2010), Karahan and Ozkan (2013), Caplin et al. (2023), Arellano et al. (2017), De Nardi et al. (2019), and Guvenen et al. (2021b). Low et al. (2010) explore a model with labor market search, employment risk, and consumption-saving decisions. They abstract from heterogeneity in job stability, and earnings dynamics are predominantly governed by an exogenous stochastic productivity process. Karahan and Ozkan (2013) estimate a stochastic earnings process with age-dependent parameters

[3]Heterogeneity in job stability across regional labor markets has recently been highlighted in Bilal (2019), Jung et al. (2023), and Kuhn et al. (2026 forthcoming) as the main driver of spatial unemployment rate differences.

and find that the variance and persistence of the process vary with age. They find that the welfare consequences of market incompleteness are substantially lower in a model with an age-varying income process compared to a model with age-invariant income risk. Using administrative data, Caplin et al. (2023) show that the probability of job separation is a main driver of subjective earnings risk. Their model suggests that heterogeneity in job-match quality across workers generates a large part of earnings risk. Arellano et al. (2017) estimate an earnings process with nonlinear persistence and conditional skewness, showing that the consequences of earnings shocks depend on earnings histories and these nonlinearities generate heterogeneous consumption responses. Complementing this work, our paper provides a structural labor-market mechanism that generates such heterogeneous earnings paths and income risk, and studies how this shapes consumption and life-cycle saving behavior. Using administrative data, Guvenen et al. (2021b) document that earnings changes are distributed non-normally, exhibiting negative skewness and high kurtosis with significant variation across ages and earnings levels. These findings challenge standard earnings processes used in macroeconomic models and underscore the importance of modeling nonlinear income processes. De Nardi et al. (2019) develop a flexible income process with persistent and transitory components that allows for age dependence, non-normality, and nonlinear dynamics. Their framework provides a substantially improved fit to the evolution of cross-sectional consumption inequality over the life cycle and to the degree of individual-level consumption insurance against persistent earnings shocks. Our paper complements these papers by providing a search-and-saving model that endogenously generates nonlinear earnings dynamics and heterogeneity in earnings and consumption responses without imposing nonlinear persistence exogenously.

Section 2 documents the empirical relationship between employment stability, earnings growth, and wealth accumulation. Section 3 presents the model. Section 4 explores the life-cycle consequences of heterogeneity in employment stability for income growth and wealth accumulation. Section 5 studies macroeconomic implications. Section 6 concludes.

2 Heterogeneity in employment stability and wealth accumulation

This section provides new evidence for the salient fact of large heterogeneity in employment stability in the U.S. labor market. In a second step, we provide novel evidence on the tight correlation of employment stability with earnings growth and wealth accumulation. This empirical evidence serves as the motivation and will guide our model building in the next step. The empirical analysis relies on data from the *Panel Study of Income Dynamics* (PSID) and consists of two steps. In the first step, we show empirical evidence for large heterogeneity in employment stability in the U.S. labor market. In the second step, we document the correlation between employment stability, earnings growth, and wealth accumulation. In Appendix C.2, we provide corroborating evidence for both steps based on independent high-quality data from the *Survey of Consumer Finances* (SCF).

The PSID is the longest-running household panel for the United States beginning in 1968. The surveys were conducted annually until 1997 and biennially thereafter. Questions on household wealth have been available since 1985 at five-year intervals and biennially since 1999. Exploiting the panel dimension of the PSID, we track the labor market status and the financial situation of households over time. We restrict the sample to household heads aged 25 to 55 and exclude individuals who are permanently out of the labor force. In addition, we retain only households for which we observe the full prime working-age period. The resulting sample consists of 1,949 households observed in panel form.

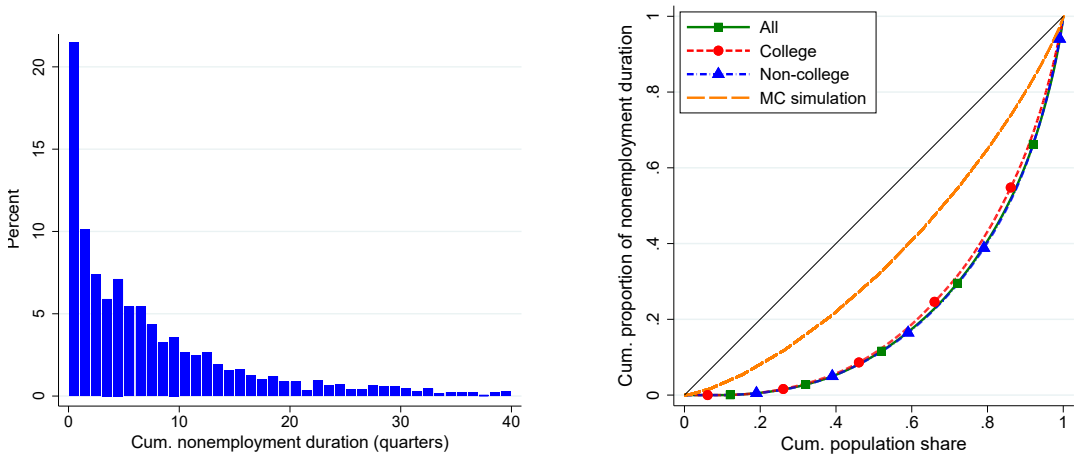
For our analysis of employment stability, one key observation in the PSID data is the labor market status of the individuals. In the PSID, individuals report the number of weeks of employment in the last year. Using this information, we construct non-employment duration by subtracting the weeks employed from the total number of weeks of a year. We assume that an individual was employed for the whole year if the reported weeks of employment are larger than 47 weeks to account for holidays. Starting in 1999, we observe the non-employment duration of workers only biennially. Beginning in 1999, non-employment duration is observed only biennially. Appendix A.1 discusses this feature of the data.

2.1 Heterogeneity in employment stability

To quantify heterogeneity in employment stability, we start in Figure 1 with a histogram of the accumulated life-cycle non-employment duration across worker careers. Figure 1a shows the distribution of non-employment duration at age 50 that workers have accumulated between ages 25 and 50. Strikingly, we find that two out of ten workers experience no spell of non-employment over their prime-age working life. Around half of all workers accumulate a non-employment duration between 1 and 10 quarters, whereas the remaining 30% of workers have accumulated non-employment duration of more than 10 quarters. Such a non-employment duration of 10 quarters corresponds to an individual lifetime non-employment rate of 8 percent.

Figure 1b uses standard inequality measurement statistics to summarize how unequally the non-employment duration is distributed across workers by showing the Lorenz curves for non-employment duration from Figure 1a. Looking first at all workers (green line with squares), we find that around 30% of workers account for more than 70% total non-employment duration. Hence, labor market mobility is concentrated in a small group of workers. To contrast this inequality to a situation with homogeneous employment stability, we show as dashed yellow line the Lorenz curve from a Monte Carlo simulation that abstracts from heterogeneity in employment stability ("representative worker case"). The Lorenz curve without heterogeneity is much closer to the identity line. In this case, 30% of workers account for only 40% of the accumulated non-employment duration so the distribution of non-employment is much more equally distributed across workers.

Figure 1: Histogram and Lorenz curve of non-employment duration



(a) Histogram

(b) Lorenz curve

Notes: Panel (a) shows the histogram of non-employment duration (in quarters). Panel (b) shows the Lorenz curve of non-employment duration for all, college-educated, and non-college households, and a Monte Carlo simulation where all workers have the average age-dependent labor market transition rate (separation and job-to-job transitions). Non-employment duration is measured at age 50 of individuals. Data are from the Panel Study of Income Dynamics.

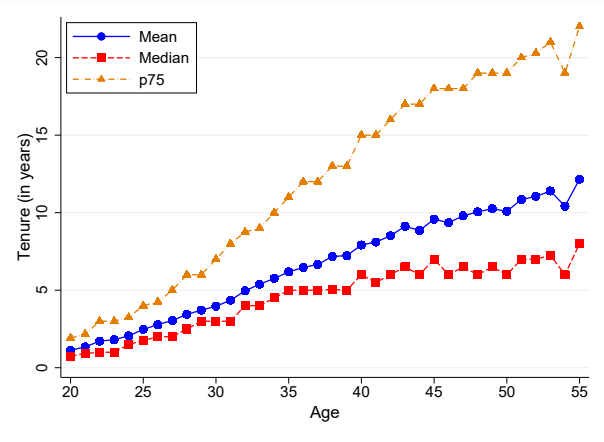
Besides heterogeneity in job stability, worker heterogeneity could be another important driver of such heterogeneity in non-employment episodes. The arguably most important dimension of worker heterogeneity is educational attainment. Figure 1b shows as dotted red line with circles the Lorenz curve for non-employment duration of college graduates, and as dashed blue line with triangles, the corresponding Lorenz curve for non-college workers. On average, college-educated workers accumulate about 20 months of non-employment, whereas non-college workers experience roughly 50% more non-employment over the life cycle. Strikingly, we find that despite very different average transition rates between these groups, the within-group inequality in non-employment duration is virtually identical. Hence, even within two groups of workers with arguably very different average employment stability, we find the same amount of inequality in employment stability. In our structural model, we will therefore focus on employment stability stemming from heterogeneity in job stability that is consistent with large within-worker-group heterogeneity in job stability and abstract from between-worker heterogeneity.[4]

Another approach to document heterogeneity in employment stability is to look at the tenure distribution (Hall, 1982; Jung and Kuhn, 2018). Figure 2 shows life-cycle profiles for tenure in the PSID data.[5] We find that the tenure profile is positively correlated with age. Looking at the mean,

[4]Appendix 4.3 explores the consequences of such heterogeneity through the lens of the model.

[5]Employer tenure is widely observed in cross sectional datasets and provides therefore an attractive variable to

Figure 2: Tenure over the life cycle



Notes: The figure shows the life-cycle evolution of the cross-sectional distribution of tenure (in years). Data are from the Panel Study of Income Dynamics.

the median, and the 75th percentile of the tenure distribution, we observe a spreading out of the distribution as workers age. Consistent with the finding from Figure 1 and as already pointed out in Hall (1982), the typical U.S. worker has a stable employment history with one in five workers experiencing no non-employment over their entire prime-age working life. Looking at the tenure distribution at age 60, we find that more than 50 percent of workers have been with their employer for 10 years, and almost a quarter of workers at age 60 have been at the same employer for at least 25 years. It is this large group of workers to which Hall (1982) refers to as having lifetime jobs.

2.2 Employment stability and earnings dynamics

What are the implications of the observed heterogeneity in employment stability on earnings and wealth accumulation? To address this question, we leverage that the PSID data tracks labor market histories, earnings, and wealth accumulation jointly over time. In the first step, we examine the relationship between labor income growth and stable and unstable employment histories. We group employment histories in the unstable group if the accumulated non-employment duration between age 25 and 50 of the worker is part of the top 25% of employment histories sorted by cumulative non-employment duration at age 50. The remaining employment histories are assigned to the stable group. This sorting conditions on the ex-post realization of the employment history.[6]

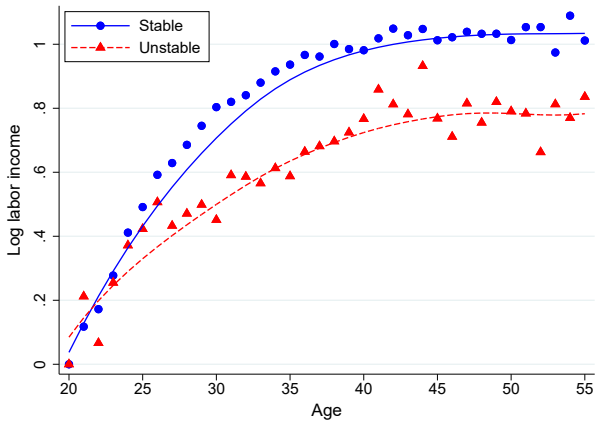
explore the relationship between wealth accumulation and employment stability. We use tenure as an alternative measure for employment stability when studying the relationship with wealth accumulation in the SCF data.

[6]One potential concern is that our measure of labor market history in the PSID may be biased, leading to misclassification due to the biennial sampling frequency since 1999. In Appendix A.1, we conduct a counterfactual analysis using the quantitative model presented in Section 3 and show that the share of misclassified individuals in a biannual sample is relatively small, at approximately 2%.

As an alternative, we document that results are robust to sorting employment histories ex ante by looking at the accumulated non-employment duration of labor market histories at worker age 30.

We also check the mobility between the groupings using non-employment duration at age 30 and 50. We find that only 13% of employment histories from the stable group at age 30 are classified as unstable group when looking at age 50. It is tempting to interpret the high persistence in the grouping as evidence for a strong worker fixed effect, however, we show below that our model with only job heterogeneity is consistent with this evidence. The reason is that stable jobs have low outflow rates and hence make the identification between a worker and a job component challenging. In Appendix B, we discuss evidence that looks at cases when stable jobs end, for example following a displacement event, and consistent with the interpretation of an important job component we find employment histories with higher non-employment rates after the job loss.

Figure 3: Log labor income



Notes: The figure shows the life-cycle profiles of log labor income for stable and unstable careers in terms of employment stability. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. The profiles of log labor income are normalized to zero at age 20. Data are from the Panel Study of Income Dynamics.

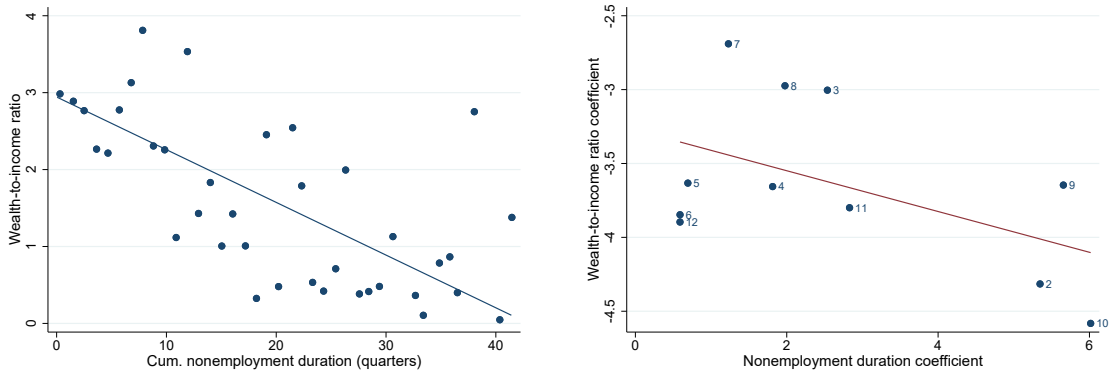
Figure 3 shows the life-cycle profiles of log labor income for workers with stable and unstable employment where we have removed initial differences to focus on the differences in income growth. Strikingly, we observe a large difference in earnings growth over the life cycle. At age 50, the average labor income of the stable group has grown 20 log points more compared to the group of workers with unstable employment history. Accounting for household-level income further amplifies this difference, as shown in Appendix A.2. This result speaks against the common practice to consider the average life-cycle earnings growth and earnings risk as independent. Our results suggest that being more often non-employed does not only lead to an income loss during non-employment, but also to lower income growth over the life cycle. What we cannot rule out is that a repeated permanent negative shocks from job loss leads to lower earnings over the life cycle. What speaks

against this interpretation of repeated large negative and persistent shocks is, however, that the empirical literature finds the large and persistent shocks for long-term employed workers from stable employment relationships (Jacobson et al., 1993; Davis and von Wachter, 2011). Put differently, only workers in a currently stable employment relationship experience large and persistent earnings losses. We revisit this fact below in the context of our model when discussing the macroeconomic implications of heterogeneity in employment stability.

2.3 Employment stability and wealth accumulation

How do these differences in labor market experiences affect workers saving behavior? The PSID data allow us to trace wealth accumulation of workers with different employment trajectories over time. In Figure 4, we consider the relationship between the accumulated non-employment duration from Figure 1b and wealth-to-income ratios and non-employment.

Figure 4: Wealth-to-income ratio and non-employment duration



(a) All

(b) By occupation

Notes: Panel 4a presents a binned scatter plot of wealth-to-income ratios against cumulative non-employment duration for workers at age 50. Panel 4b displays a scatter plot of occupation effects on wealth-to-income ratios and non-employment duration, obtained from separate regressions controlling for demographic characteristics. The baseline occupation (occupation 1) is omitted. Occupations are classified as follows: (1) Agriculture, forestry, and fisheries; (2) Mining; (3) Construction; (4) Manufacturing; (5) Transportation, communications, and other public utilities; (6) Wholesale and retail trade; (7) Finance, insurance, and real estate; (8) Business and repair services; (9) Personal services; (10) Entertainment and recreation services; (11) Professional and related services; and (12) Public services.

We look at wealth-to-income ratios already removes the effect of lower income on wealth accumulation that we have documented in Figure 3. For the non-employment duration, we look at the accumulated duration at age 50. Figure 4a shows a strong negative relationship with wealth-to-income ratios. Qualitatively, a negative slope implies that per dollar of income, workers in more stable employment have more wealth, or, in short, workers with more stable employment history

are wealthier. Quantitatively, the observed slope is economically meaningful. The slope implies that being non-employed for four years will, on average, lead to additional wealth corresponding to roughly one year of income.[7] One concern might be other confounding factors from fixed worker differences, most prominently, education. In Appendix C.1, we repeat the analysis with additional controls for education and industry. We find that the documented relationship is both qualitatively and quantitatively robust.

Table 1: Wealth-to-income ratio and non-employment

Wealth-to-income ratio	(1)	(2)	(3)	(4)	(5)
Non-employment (in quarters)	-0.07^{***}	-0.06^{***}	-0.04^{***}	-0.03^{**}	-0.14
	(0.01)	(0.01)	(0.01)	(0.01)	(0.08)
Education		1.21^{***}	0.92^{***}	1.05^{***}	
		(0.31)	(0.01)	(0.01)	
Non-employment x education		-0.03		-0.02	
		(0.03)		(0.03)	
Other worker characteristics	No	No	Yes	Yes	Yes
Observations	1537	1537	1537	1537	11
R^2	0.02	0.04	0.09	0.09	0.25

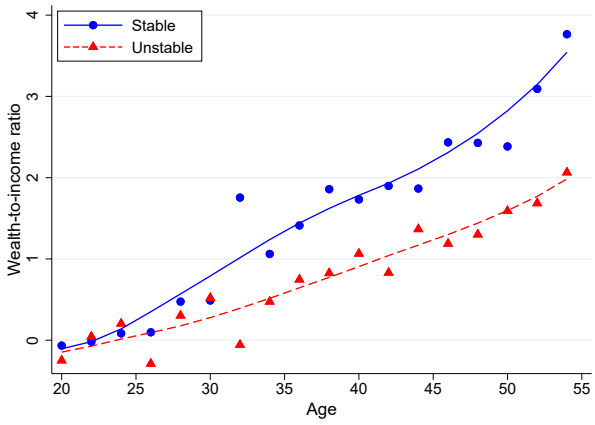
Notes: The dependent variable is wealth-to-income ratio. The sample consists of workers at age 50. Column (1) shows the regression results of wealth-to-income ratio on non-employment duration of workers in quarters. Column (2) includes education and the interaction of education and non-employment duration as additional controls. In columns (3) and (4), we include other worker characteristics. Column (5) presents the regression of the wealth-to-income ratio on non-employment duration by occupation, corresponding to Figure 4b. Other worker characteristics refer to the characteristics of the household head other than education, including sex, race, region, marital status, family size, and the industry with the longest employment duration. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In Table 1, we present the results of a regression of wealth-to-income ratio on non-employment duration of workers. Column 1 shows that a non-employment duration of one year (4 quarters) is associated with a 0.3 lower wealth-to-income ratio. While there is a significant negative relationship between employment stability and wealth-to-income ratio, worker heterogeneity such as education is another important determinant of differences in wealth (Bartscher et al., 2020). We consider education as an additional control in column 2. Higher educational attainment is associated with higher wealth-to-income ratios, but we still find a negative and significant relationship between non-employment and wealth-to-income ratio. Even after including other worker characteristics, the coefficient of non-employment remains negative and statistically significant.

[7] In line with this finding, Iacono and Ranaldi (2020) report for Norwegian data a negative correlation between wealth and unemployment.

Alternatively, we isolate the variation in employment stability and wealth building by looking at the variation across occupation groups. This allows us to most directly analyze the role of job characteristics in shaping household wealth and non-employment outcomes. We control again for a rich set of individual demographic characteristics. Specifically, we estimate regressions of the household wealth-to-income ratio and cumulative non-employment duration on a full set of occupation fixed effects, while controlling for education, gender, race, region, marital status, and family composition. By conditioning on these observable individual characteristics, the estimated occupation effects capture systematic differences across jobs rather than compositional differences across workers. Figure 4b plots, for each occupation, the estimated occupation coefficient from the wealth-to-income regression against the corresponding coefficient from the non-employment regression. The figure reveals a pronounced negative relationship between the two sets of occupation effects. Occupations characterized by longer cumulative non-employment spells are also associated with substantially lower wealth-to-income ratios, whereas industries with more stable employment are associated with higher relative wealth accumulation. Notably, the slope is with an absolute value of 0.14 even larger than in the raw data (Table 1). This pattern indicates that differences in wealth accumulation are strongly driven by job-specific labor market risk rather than by individual characteristics. More broadly, the results highlight the importance of occupation-level employment stability as a key determinant of household wealth outcomes.

Figure 5: Wealth-to-income ratio



Notes: The figure shows the life-cycle evolution of the wealth-to-income ratios for stable and unstable careers in terms of employment stability. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. The wealth-to-income ratios are normalized to zero at age 20. Data are from the Panel Study of Income Dynamics.

In the final step, we look at the life-cycle wealth accumulation associated with different employment histories. Figure 5 shows the life-cycle wealth-to-income ratio profiles for workers with stable and unstable employment histories as in Figure 3. We find a persistent gap in wealth-to-income

ratios for workers with stable and unstable jobs over the life cycle. At age 50, workers with more stable careers are wealthier by one and a half years of income compared to workers with unstable employment. The differences in wealth accumulation are apparent over the entire life cycle starting at age 25. The constant differences in slopes are consistent with the idea that once a difference in job stability has materialized it remains very persistent over the working life.

To summarize, we document that in the data, employment stability is very heterogeneous across careers and is as large in the entire population of workers as within groups of workers with on average different labor market mobility pattern. The novel facts with respect to the consequences of this heterogeneity is that differences in employment stability are systematically related to earnings growth and wealth accumulation. The differences are economically significant. The regression analysis shows that, while education is an important driver of wealth differences across workers, there is also a substantial within-group heterogeneity in wealth-to-income ratio among college-educated workers and workers without college education.

2.4 Sensitivity and Robustness

In Appendix C.2, we further corroborate a positive relationship between employment stability and wealth using data from the SCF that provides information on job tenure and the number of employers that workers had over their career. All results strongly support the idea that more stable employment leads to more wealth accumulation over the life cycle.

3 Search-and-saving model

In this section, we develop a search-and-saving model that combines a standard model of household saving behavior with a labor market search model that features heterogeneity in employment stability. In what follows, we focus on the within-group heterogeneity in analyzing the relationship between employment stability and wealth. We interpret the source of heterogeneity as coming from job heterogeneity in line with a large strand of the literature (Hall, 1982; Pinheiro and Visschers, 2015b; Jung and Kuhn, 2018; Jarosch, 2023; Larkin, 2019). In Appendix E.1, we describe a model version without heterogeneity in employment stability that we refer to throughout the discussion.

The model is populated by risk-averse agents who maximize expected lifetime utility. Agents derive utility from consumption and disutility from effort required to accumulate human capital. Labor supply at the intensive margin is inelastic so that each employed worker supplies one unit of time.

We denote a worker's age by j and split a worker's life cycle into three phases: a *working phase*, a *transition phase*, and a *retirement phase* (Krebs et al., 2015). Workers start their life in the

working phase that lasts for T^W periods. At the end of the working phase, workers move to the transition phase that is of stochastic length with expected duration T^T. In the end, workers leave the transition phase to the retirement phase that lasts for T^R periods. In each period before the retirement phase, a worker is either employed or non-employed. We denote the agent's employment status by ε with $\varepsilon \in \{e, n\}$ where e stands for employed and n for non-employed. If the worker is employed, her job is characterized by a bundle (w, λ) where w denotes the wage and λ the separation rate where the wage w captures the rental rate of human capital on the current job. We discretize wages and separation rates to grids $\{w_k\}_{k=1}^K$ and $\{\lambda_l\}_{l=1}^L$ and assume that $w_k < w_{k+1}$ for all k and $\lambda_l < \lambda_{l+1}$ for all l. To economize on notation, we denote the wage-separation rate bundle at age j only by $\{w_j, \lambda_j\}$. The difference in λ capture the heterogeneity in employment stability. Each worker holds assets denoted by a and a stock of human capital denoted by h. The period budget constraint is

$$a_{j+1} + c_j = (1 + r)a_j + y(w_j, h_j, \varepsilon), \tag{1}$$

where r denotes the risk-free rate on the economy's single risk-free asset and y denotes current period labor income including transfers. If the agent is employed in the current period, then the worker's income is $y(w_j, h_j, e) = w_j h_j$, the wage rate times the stock of human capital. If the agent is non-employed, she initially receives transfer income proportional to her last employment income $y(w_j, h_j, u) = b w_j h_j$ where b denotes the replacement rate and w_j is the wage on the last job. These benefits decline each period if the agent remains non-employed. We capture declining benefits by lowering the last wage on the grid from w_k to $\max\{w_{k-1}, w_1\}$.

We assume that human capital stays constant during non-employment so the current stock of human capital h_j corresponds to the human capital stock when last employed. During retirement, agents receive social security benefits proportional to their stock of human capital prior to retirement times the economy-wide average wage $y(w_j, h_j, n) = s\bar{w}_j h_j$ where $s \in (0, 1)$ denotes the replacement rate of the old-age social security system (Borella et al., 2023; Nam, 2022).

When the worker is in the working or transition phase, we split each period into four stages: *separation*, *investment*, *production*, and *search*. At the separation stage, employed agents separate from their job with probability λ. If the agent separates, she becomes non-employed and moves to the production stage. Employed agents who do not separate move to the investment stage where human capital investment decisions are made. At the production stage, employed agents receive earnings, the job's wage rate times the worker's stock of human capital, and non-employed agents receive benefits proportional to earnings on their last job. At the search stage, employed and non-employed agents receive job offers. We allow for different job-offer arrival rates on the job and in non-employment. We take job-offer arrival rates as exogenous and denote the arrival rate on the job by π_e and the arrival rate in non-employment by π_n. Job offers, combinations of a wage rate w and a separation probability λ, for employed and non-employed workers are drawn from the same

joint distribution $f(w, \lambda)$ (Pinheiro and Visschers, 2015b; Jarosch, 2023; Larkin, 2019). An agent who receives a job offer decides to reject or accept the job offer. If the agent accepts the job offer, she will be employed at the beginning of the next period in the new job. If the agent rejects the job offer, she remains non-employed (employed in her current job) and there is no recall of previous job offers.

Only employed workers have the opportunity to invest in their human capital. At the investment stage, the agent decides if she wants to exert effort for human capital investment. Effort provision for human capital accumulation is a choice $t \in [0, 1]$ (training). Disutility from effort enters the utility additively separable as quadratic cost κt^2. Nonemployed agents do not have the opportunity to accumulate human capital. If agents do not exert effort, their human capital stays constant at level h until the next period.[8] One interpretation of this effort provision is as career investment with the current employer (e.g., unpaid overtime, higher work intensity, on-the-job training, or committee work). We assume that human capital levels are discrete and are members of an ordered set with largest (smallest) element h^{max} (h^{min}). We denote by h^+ the immediate successor of human capital level h and by h^- the immediate predecessor of h. Human capital investment is risky. An agent at human capital level h exerting effort t to accumulate human capital has a probability $p_H(t, j)$ of reaching human capital level h^+. We allow for age dependence of $p_H(t, j)$. The law of motion for human capital if the agent exerts effort ($t > 0$) is

$$h_{j+1} = \begin{cases} h_j^+ & \text{with probability } p_H(t, j) \\ h_j & \text{with probability } 1 - p_H(t, j). \end{cases}$$

This structure of the human capital process is an extension to Jung and Kuhn (2018) endogenizing the human capital accumulation decision.

The consumption-saving decision is standard. The agent chooses next period's asset level given her current state and facing a no-borrowing constraint. Agents make savings decisions at the production stage before knowing the outcome of the search stage. We denote the period utility function over consumption c by $u_j(c)$. The working and the transition phase differ only in the possible continuation states. A worker in the transition phase either remains in the transition phase or transits to the retirement phase. A worker in the working phase ages deterministically and transits at the end of prime-age working life to the transition phase. We do not allow workers from the transition phase (retirement phase) to transit back to the working (transition) phase.

Transiting from the transition phase to the retirement phase is stochastic and happens with probability ψ. Upon reaching the retirement phase, workers leave the labor market and receive

[8] Although we do not assume human capital depreciation during non-employment, there is on average relative depreciation of human capital because employed workers invest and accumulate human capital while non-employed workers do not.

social security benefits. Agents do not face any labor market risk during retirement and solve a deterministic, finite-horizon consumption-saving problem.

We formulate the agent's decision problem recursively. The state of an agent is described by her age j, her employment state ε, her current asset holdings a, her current or last wage wage w, the separation probability λ if employed, and her level of human capital h. We formulate separate value functions for employed and non-employed workers so that we drop the employment state from the state vector. We use primes to denote next period's states. In a slight abuse of notation, we drop the primes in case variables do not change between periods.

The value function of an employed worker at the beginning of the period V_e is given by the expectations over the employment status as an outcome of the separation stage,

$$V_e(a, w, \lambda, h, j) = \lambda V_n^P(a, w, h, j) + (1 - \lambda)V_e^I(a, w, \lambda, h, j), \tag{2}$$

where V_n^P denotes the value function of a non-employed worker at the production stage and V_e^I denotes the value function of an employed worker at the investment stage. Note that the value function of a non-employed worker at the production stage V_n^P is identical to the value function at the separation stage V_n because for already non-employed workers, nothing happens at the separation stage.

At the investment stage, an employed agent makes her human capital investment decision. The realization of the stochastic human capital accumulation happens at the beginning of the production stage:

$$V_e^I(a, w, \lambda, h, j) = \max_{t \in [0,1]} -\kappa t^2 + p_H(t, j)V_e^P(a, w, \lambda, h^+, j) + (1 - p_H(t, j))V_e^P(a, w, \lambda, h, j). \tag{3}$$

The Bellman equation of an employed agent at the production stage is

$$\begin{aligned} V_e^P(a, w, \lambda, h, j) = \max_{\{c, a' \geq 0\}} & \; u_j(c) + \beta\Big(\pi_e V_e^S(a', w, \lambda, h, j) + (1 - \pi_e)V_e(a', w, \lambda, h, j + 1)\Big) \\ s.t. \qquad & c = (1 + r)a + y(w, h, e) - a', \end{aligned} \tag{4}$$

where V_e^P denotes the employed agent's value function at the production stage, V_e^S denotes the employed agent's value function at the search stage, and V_e denotes the value function of an employed worker at the beginning of the period. The time discount factor is denoted by β. The first line of equation (4) is composed of the flow utility for the current period and the discounted expected utility from the search stage. The probability of receiving a job offer is π_e. The distribution over job offers is $f(w, \lambda)$, so that for the value function of an employed worker at the search stage,

we get

$$V_e^S(a', w, \lambda, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max\Bigg\{ \underbrace{V_e(a', w, \lambda, h, j+1)}_{\text{staying in current job}}, \underbrace{V_e(a', w_s, \lambda_k, h, j+1)}_{\text{accepting outside offer}} \Bigg\} f(w_s, \lambda_k), \quad (5)$$

where N_w is the number of wage realizations in the support of the offer distribution and N_λ is the number of realizations for separation rates in the support of the offer distribution. The value function at the search stage captures the acceptance-rejection decision for outside job offers and the expectations over job offers.

The value function of a non-employed worker at the production stage is

$$\begin{aligned} V_n^P(a, w, h, j) &= \max_{\{c, a' \geq 0\}} u_j(c) + \beta \Big(\pi_n V_n^S(a', w, h, j) + (1 - \pi_n) V_n(a', w^-, h, j+1) \Big) \quad (6) \\ s.t. \quad & c = (1+r)a + y(w, h, u) - a', \end{aligned}$$

where declining benefits are captured by a transition from w to w^- where w^- denotes the next lower wage level.

For the value function of a non-employed worker at the search stage captures again the acceptance-rejection decision over job offers, it reads

$$V_n^S(a', w, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max \Bigg\{ \underbrace{V_n(a', w^-, h, j+1)}_{\text{staying non-employed}}, \underbrace{V_e(a', w_s, \lambda_k, h, j+1)}_{\text{accepting job offer}} \Bigg\} f(w_s, \lambda_k). \quad (7)$$

The value functions for the transition phase directly follow the value functions of the working phase. The only difference is that they comprise a probability ψ that at the end of the period, the worker retires and goes to the retirement phase. All decisions are otherwise identical to the working phase. We show value functions for the transition phase in Appendix D.1.

During the retirement phase, agents receive retirement benefits and do not face any income risk. At the end of the retirement phase, everyone dies. We normalize utility in this case to zero. As we abstract from a bequest motive, we get that at the end of the life cycle, all agents will have zero assets. The Bellman equation for retirement reads

$$V_r(a, w, h, j_r) = \max_{a' \geq 0} u((1+r)a + y(w, h, n) - a') + \beta V_r(a', w, h, j_r + 1). \quad (8)$$

We solve the model using backward induction and grid search for the consumption-saving and effort choice decisions. We provide further details on the numerical implementation in Appendix D.2.

3.1 Model calibration

We make the following assumptions on parameters, functional forms, and the human capital process to bring the model to the data. We set a model period to correspond to one quarter and assume the utility function for consumption is

$$u_j(c) = \log(c/\phi_j) \cdot \phi_j$$

where ϕ_j is the household equivalence scale obtained from the PSID data. Human capital takes on discrete values $h_{i,t} \in \{h_1, .., h_{N_h}\}$. The human capital grid comprises N_h states that we set equidistant in log space between $h_1 = 1$ and $h_{N_h} = 6.5$. The probability $p_H(t)$ is age dependent and declines geometrically according to rate ρ,

$$p_H(t, j) = \rho^{j-1} \times t \times \bar{p}_H,$$

with effort provision t and baseline level $\bar{p}_H$. We discuss below that, together with the specification for wages, the human capital process matches the stylized empirical facts on earnings growth and its composition.

At labor market entry, each agent is endowed with the lowest level of human capital $h_1 = 1$ and initial assets $a_0 = 0$. We set some parameters to conventional values or to match external targets. We set the replacement rate in non-employment to 0.4, as in Shimer (2005), and in retirement to 0.45, in line with the OECD estimate for the mean net pension replacement rate in the United States (OECD, 2015). We set working life T^W to 35 years, the duration of the transition phase between employment and retirement T^T to an expected duration of 10 years, and the retirement phase T^R to a duration of 20 years. Labor market entry happens at age 20. We focus our discussion of the model prediction to working life, i.e., between ages 20 and 55.

We denote the support of wages by $[\underline{w}, \overline{w}]$ and of job stability by $\left[1-\overline{\lambda}, 1-\underline{\lambda}\right]$. We set $N_w = 5$, $\underline{w} = 1$, and $\overline{w} = 1.85$, in line with the empirical support of mean log earnings, and use equidistant grid points in logs. For job stability $1-\lambda$, we set $N_\lambda = 10$ and set $\overline{\lambda} = 0.35$ and $\underline{\lambda} = 0.006$ to represent lifetime jobs with an expected duration of 42 years. We set the remaining grid points nonlinearly between the most and least stable jobs, with more grid points toward the least stable job.[9] For the functional form of the job-offer distribution $f(w, \lambda)$, we assume that the marginal distributions of wages and job stability $(1-\lambda)$ follow a truncated exponential distribution. We map both supports to the unit interval $[0, 1]$ denoting by $w^* \in [0, 1]$ the standardized wage and by $1-\lambda^*$ standardized job stability. The density of w^* is $f(w^*) = (1-\exp(-\psi_w))^{-1}(\psi_w \exp(-\psi_w w^*))$ where ψ_w determines the shape of the density. The density of standardized job stability $1-\lambda^*$

[9] Specifically, we set the second grid point at $\lambda_2 = 0.05$ and the remaining grid points according to the nonlinear rule $\lambda_j = \underline{\lambda} + \left(\frac{j-1}{N_\lambda - 1}\right)^{0.6} \times (\overline{\lambda} - \underline{\lambda})$.

follows accordingly with shape parameter ψ_λ. To determine the joint distribution, we parameterize the correlation between the marginal distributions by constructing the joint distribution using a copula C_θ, where the value of θ determines the correlation between w^* and $1-\lambda^*$.

We determine parameters within the model using a simulated method of moments that minimizes the difference between model moments and empirical moments. For the empirical moments, we use the life-cycle profiles of (log) earnings (mean and variance), labor market transition rates, tenure (mean, median, 75th percentile), and of the wealth-to-income ratio. For labor market transition rates, we rely on estimated life-cycle profiles from Jung and Kuhn (2018) based on CPS data.[10] In Appendix D.2, we provide additional details on the implementation of the estimation procedure and offer an intuitive discussion of how the empirical life-cycle profiles help identify the model's free parameters. Rather than presenting a formal proof of identification, we conduct a comparative statics analysis in Appendix D.2.3. This analysis illustrates which parameters are disciplined by which data moments and how changes in parameter values affect the model-generated moments. Table 2 presents the model parameters together with their estimated values.

Table 2: Estimated parameters

Parameter	Value	Description
β	0.994	Quarterly discount factor
κ	0.398	Utility cost of effort
π_e	0.427	Probability of a job offer when employed
π_u	0.851	Probability of a job offer when nonemployed
ψ_w	0.536	Marginal distribution of w^*
ψ_λ	0.504	Marginal distribution of $1-\lambda^*$
θ	0.489	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.071	Skill upgrading probability
ρ	0.984	Persistence of skill upgrading probability

The value of the quarterly discount factor β corresponds to an annualized value of 0.97, which is well within the range of conventional values in the macroeconomic literature. The utility cost parameter κ implies average utility costs of effort provision measured as lifetime consumption-equivalent variation between 0.35 percent during the first ten years of working life and less than one-tenth of a percent during the last ten years of working life.[11] For labor market parameters,

[10] We refer to Jung and Kuhn (2018) for details on the construction of labor market mobility rates. Following their approach, we do not distinguish between separations into unemployment and separations to out of the labor force. See Jung and Kuhn (2018) and Kudlyak and Lange (2014) for more discussion.

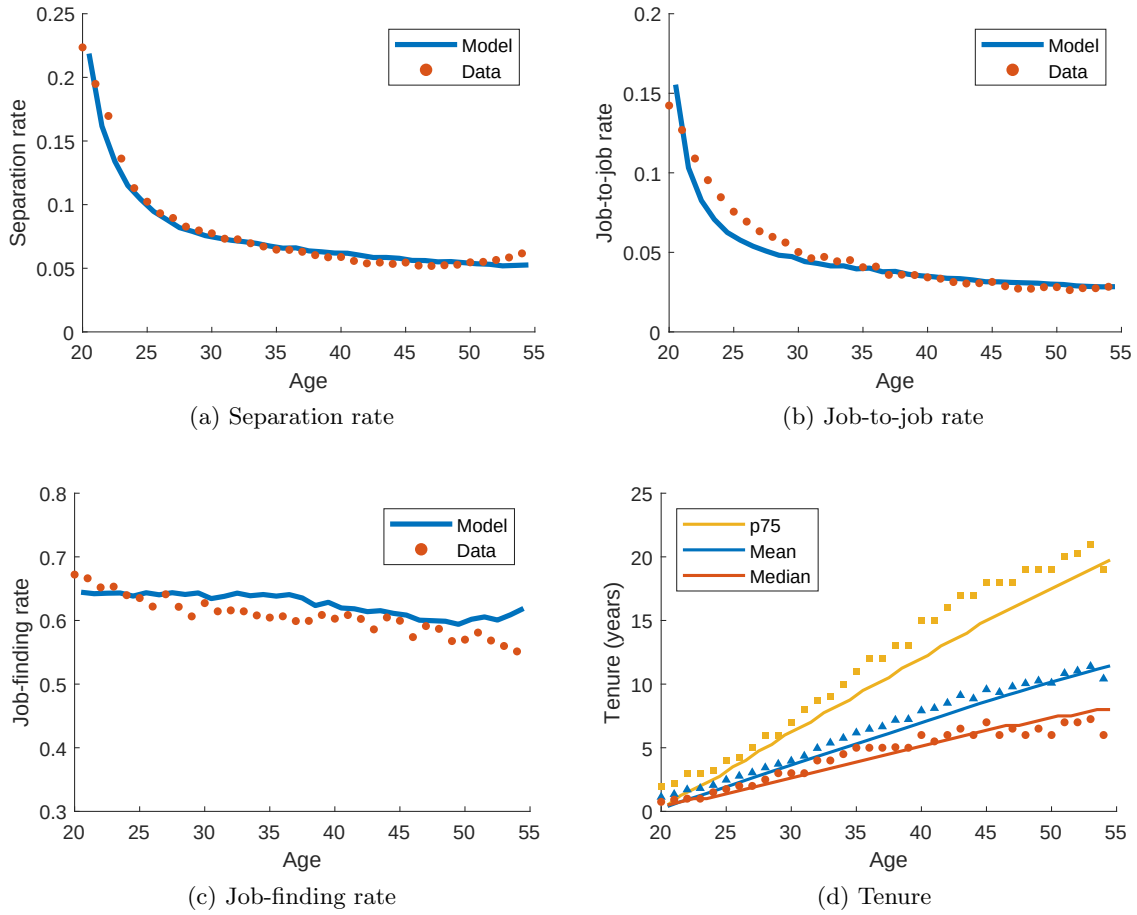
[11] Utility costs as a share of current period consumption are substantially larger and amount to 7.7 percent at age 40 with a steep age gradient.

we get that job offer probabilities in non-employment π_u have to be roughly twice as high than in employment π_e, to match the high quarterly job-finding rates, but even during employment, workers frequently get job offers. Such a difference between contact rates is qualitatively and quantitatively consistent with the calibration in Hornstein et al. (2011) for the United States. The shape parameters of the marginal distributions, ψ_w and ψ_λ, determine the relative frequency of the different wage and job stability levels in the offer distribution. They imply that one-third of job offers come with the lowest wage, and less than one out of twelve job offers come with the highest wage level. For job stability, we get that less than one out of 20 jobs are the most stable lifetime jobs, whereas almost one out of six job offers are of the least stable type, lasting in expectation for one quarter only. The copula parameter θ implies a positive correlation between wages and job stability. If wages and job stability were independent, the probability of the least-stable lowest-paying job would be 5.1 percent, but given θ, it is almost 50 percent higher with 7.4 percent. For the most-stable best-paying job, the offer probability is 0.9 percent, highlighting that stable and high-paying jobs are hard to find. Such a correlation between wages and job stability is also in line with the empirical evidence in Jung and Kuhn (2018) that high-wage jobs are more stable (lower separation rates). Figure D.7a shows the estimated joint job-offer distribution over wages and separation rates and the marginal distributions of separation rates at different wage levels. The joint distribution is clearly asymmetric, with most of the probability mass concentrated at low-wage, unstable jobs. Figure D.7b shows that the marginal distribution of separation rates in low-wage jobs always first-order stochastically dominates the distribution of separation rates in high-wage jobs. The parameter $\bar{p}_H$ for the human capital process implies that for a labor market entrant maximum effort provision during the first year ($t = 1$) yields a 20 percent probability of career progression (human capital increase). The decay in the human capital investment technology ρ implies that after 10 years in the labor market, the same effort provision will imply a 11 percent probability of career progression. While all parameters appear economically reasonable in isolation, we will now demonstrate that they yield a close fit between the model and data along targeted and untargeted dimensions.

3.2 Model fit

Figure 6 shows the empirical life-cycle profiles for separation, job-to-job, and job-finding rates and their model counterparts. Looking at the separation rate in Figure 6a, we see that its evolution is matched very closely. Model and data show a strong decline up to age 30 and constantly falling separation rates between ages 30 and 50. The model mechanism to match these declining life-cycle profiles consists of workers climbing the job ladder and finding more stable and better-paying jobs. Note that this sorting of workers into more stable jobs over the life cycle requires that workers can leave unstable jobs. Fixed worker heterogeneity in separation rates cannot account for the life-cycle decline of separation rates as workers will take their worker type with them while climbing

Figure 6: Transition rates and tenure



(a) Separation rate

(b) Job-to-job rate

(c) Job-finding rate

(d) Tenure

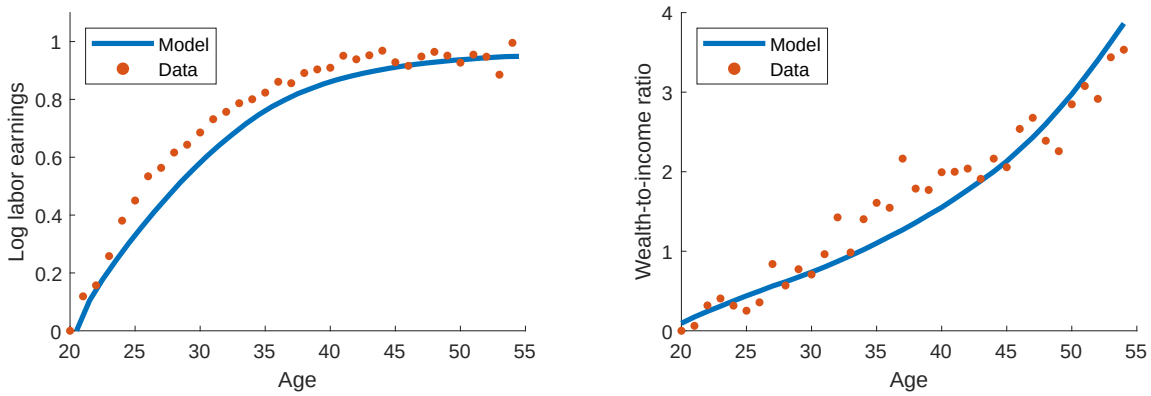
Notes: This figure shows quarterly life-cycle transition rates and tenure in years by age. The dots show the empirical profiles, while the solid lines show the corresponding model profiles. Empirical transition rates are computed using data from the CPS. The empirical tenure profiles are computed using data from the PSID.

the job ladder. Figure 6b shows that the model also matches the life cycle of job-to-job rates very well, with only a slightly steeper decline of job-to-job rates between ages 20 and 30 compared to the data. Job-finding rates in Figure 6c are matched well in level and trend and generally show little life-cycle variation. Finally, Figure 6d shows the life-cycle profiles of mean, median, and the 75th percentile of the tenure distribution. The model closely matches the empirical increase and heterogeneity in job stability. Importantly, this demonstrates that the model is jointly consistent with high average transition rates (Figures 6a and 6b) and high job stability for most workers (Figure 6d).Appendix Figure D.9 shows in addition the cross-sectional distributions of employer tenure and the number of employers over a worker's career. The model compares very favorably to the data for both distributions. In particular, it accounts for a large fraction of short-term jobs

but also for the substantial share of jobs with more than 10, 20, and even 30 years of tenure.[12]

We also demonstrate the importance of heterogeneity in employment stability in generating such a comprehensive match between model and data. Appendix Figure E.2 contrasts the match of Figure 6d to a model without heterogeneity in employment stability. Strikingly, mean job tenure in such a model is bounded by three years and the 75^{th} percentile at age 55 stays below 5 years relative to 20 years in the data. This comparison highlights that a model with representative-worker mobility falls severely short of capturing the large empirical heterogeneity of employment and earnings dynamics.

Figure 7: Earnings and wealth



(a) Earnings (b) Wealth-to-income ratio

Notes: Panel (a) and (b) show the mean of log earnings and the mean wealth-to-income ratio, respectively, both normalized to 0 at age 20. Wealth-to-income ratios are calculated as the end-of-year assets divided by yearly income. In all panels, the blue lines are the model profiles, while the red dots show the estimated empirical profiles from the PSID data.

Figure 7 turns to the life-cycle dynamics of earnings and wealth. Looking at the life-cycle profile of mean log earnings in Figure 7a, we find that the model matches the steep increase in earnings after labor market entry and the flattening out after age 40. It closely matches the large average increase of roughly 100 log points over the life cycle but shows slightly less concavity in comparison to its empirical counterpart. Figure 7b shows the life-cycle profile of the wealth-to-income ratio as a measure of wealth accumulation. Again, we find a close fit between model and data. Wealth-to-income ratios in model and data rise from slightly above 0 at age 20 to approximately 4 at age 55. Hence, the model is jointly consistent with average earnings growth and wealth accumulation over the life cycle.

[12]For most of the paper, we abstain from a cross-sectional comparison as it requires taking a stand on the age distribution in the model. We compare, if possible, age-specific model moments to the data that are independent of the specific age structure. If we have to, we assume a uniform age distribution to aggregate model results.

In Appendix D.3.1, we further show that the model also matches life-cycle inequality profiles, specifically, the overall increase in the variance of log earnings and the increase in consumption inequality. In Appendix D.3.3, we further provide a detailed analysis on further dimensions of individual earnings dynamics. We first demonstrate that the model is consistent with standard estimates for the process of earnings using a permanent-transitory decomposition, as in Meghir and Pistaferri (2004), Blundell et al. (2008), or Heathcote et al. (2010). We also corroborate the finding from Hubmer (2018) that the distribution of earnings growth in a life-cycle labor market model is consistent with the empirically observed distribution documented by Guvenen et al. (2021a). We also decompose earnings growth over the life cycle and demonstrate that our decomposition is consistent with the results in Topel and Ward (1992) on early career wage growth. The joint consistency of the model with these facts lends support to the calibration of the human capital and wage processes as the two dimensions underlying life-cycle earnings dynamics in the model. Finally, Appendix D.3.4 demonstrates that the consumption-saving and earnings dynamics of our model result in a joint distribution of earnings and wealth consistent with the data and also demonstrates that the model matches the wealth mobility dynamics over the life cycle in the PSID data. The fit of the joint dynamics lends further support to the economic mechanisms underlying our model framework, its endogenous earnings dynamics, and wealth accumulation decisions.

4 Consequences of heterogeneity in employment stability

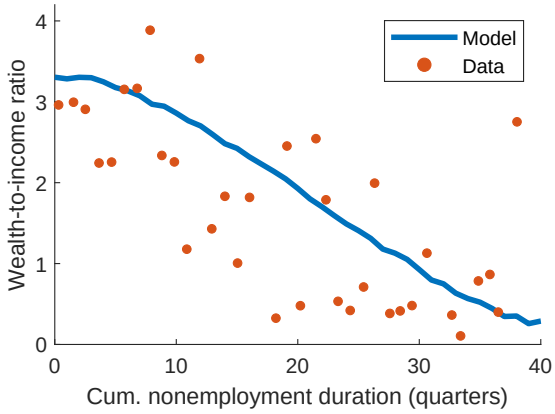
This section explores how heterogeneity in employment stability shapes earnings growth and wealth accumulation. We compare model predictions for stable and unstable careers to the empirical evidence from Section 2.2 on life-cycle earnings and wealth dynamics. Because these differences are not targeted in the calibration, they provide an independent assessment of the model's mechanism. We then use the structural model to decompose the channels through which employment stability influences earnings growth and wealth accumulation.

4.1 Heterogeneity in earnings growth and wealth accumulation

We begin by examining the relationship between employment stability and wealth at the end of working life. In the PSID data, accumulated non-employment duration is strongly negatively correlated with wealth. Figure 8 shows that the model reproduces this relationship by plotting wealth-to-income ratios against non-employment duration in both the data and simulated model outcomes. The implied slope in the model is -0.05, which lies within the empirical range reported in Table 1 of -0.07 to -0.03.

In the second step, Figure 9 compares the life-cycle profiles of earnings and wealth-to-income ratios for stable and unstable careers in the model and in the data. Simulated employment histories

Figure 8: Wealth-to-income ratio



Notes: The figure shows the wealth-to-income ratios at age 50 against the accumulated non-employment duration in the model and in the data. The blue line shows the model profile, while the red dots show the empirical wealth-to-income ratios from the PSID data.

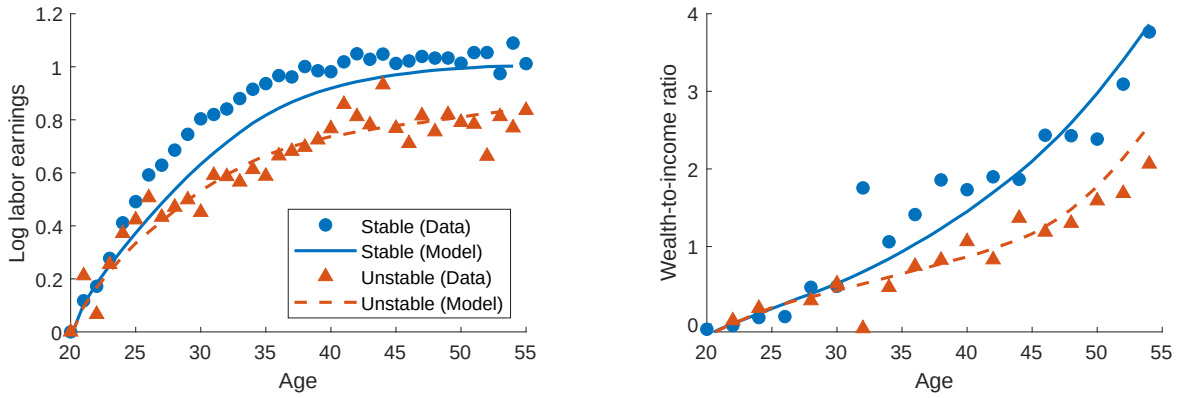
are classified using the same procedure as in the empirical analysis. The model closely reproduces the observed relationship between employment stability, earnings growth, and wealth accumulation.

Figure 9a shows that earnings profiles diverge shortly after labor market entry. The gap widens until mid-career and remains persistent thereafter. By age 50, workers in stable careers earn approximately 20 log points more than those in unstable careers in both model and data. This difference would be larger when measured in total income, since the comparison based on labor earnings excludes non-employed workers who receive only transfer income.

Figure 9b reports the corresponding wealth-to-income ratios. The model replicates the gradual divergence observed in the data: wealth accumulation begins to separate around age 30 and continues to diverge over the remaining working life. By age 55, workers in stable careers have accumulated roughly one and a half additional years of income relative to workers in unstable careers.

Taken together, the model accounts for the empirical relationship between employment stability, earnings growth, and wealth accumulation. We now use the structural framework to decompose these differences and quantify the underlying mechanisms.

Figure 9: Earnings and wealth-to-income ratio



(a) Earnings (b) Wealth-to-income ratio

Notes: Panel (a) shows the log of labor earnings of employed workers in the data and in the model, normalized to 0 at age 20. Panel (b) shows the wealth-to-income ratio in the data and in the model, normalized to 0 at age 24. The solid lines are the model profiles, while the dots show the estimated empirical profile from the data. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. Data are from the Panel Study of Income Dynamics.

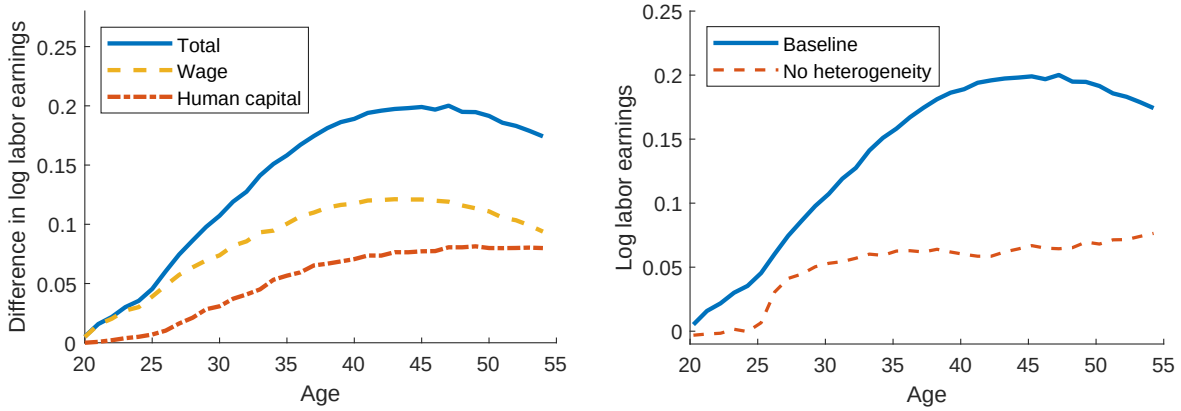
4.2 Decomposing differences in earnings and wealth accumulation dynamics

As a first step, Figure 10a decomposes the earnings gap between stable and unstable careers into wage and human capital components. The solid blue line plots the total difference in life-cycle earnings from Figure 9a. The dashed lines isolate the contributions from differences in wages and in human capital. Both components diverge over the life cycle. Early in the career, the earnings gap is primarily driven by wage differences. Over time, differences in human capital become increasingly important; by age 55, roughly half of the total earnings gap is attributable to human capital.

Figure 10b contrasts this benchmark model with a version that eliminates heterogeneity in employment stability. Without such heterogeneity, the model fails to generate a large and increasing earnings gap over the life cycle. Introducing heterogeneity in employment stability is therefore essential to account for the empirical patterns documented in Section 2.2.

We next turn to differences in wealth accumulation. Employment stability affects asset dynamics through three channels. First, higher employment risk raises precautionary savings: workers accumulate a buffer stock of assets while employed to insure against future non-employment. We refer to this channel as the *precautionary saving effect*. Second, stable careers generate stronger income growth through greater human capital accumulation. This *income growth effect* induces the life-cycle saving motive: workers with rising earnings accumulate assets in anticipation of lower income in retirement and smooth income over the life cycle. Conversely, when earnings growth is weak

Figure 10: Earnings difference decomposition



(a) Wage and human capital

(b) With and without heterogeneity

Notes: The figure shows the decomposition of the difference in log of labor earnings between the stable and unstable careers in the model. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. The solid line shows the total difference in log labor earnings. The dashed-dotted line and the dashed line show the difference in log labor earnings arising due to differences in wages and human capital, respectively.

and a career does not take off, the life-cycle saving motive is limited. Third, non-employment spells directly reduce assets: displaced workers experience lower income and draw down previously accumulated buffer stocks to smooth consumption. We refer to this mechanism as the *non-employment effect*.

The precautionary saving and non-employment effects together describe buffer-stock dynamics: assets are built up during employment and depleted during non-employment. These forces largely offset each other over time and stabilize consumption in response to short-run income fluctuations. In contrast, sustained wealth accumulation over the life cycle is driven by the income growth effect associated with stable careers.

To quantify the three channels, we conduct three counterfactual experiments.

First, to isolate the *precautionary saving effect*, we assign identical life-cycle wage and human capital profiles to workers in stable and unstable careers. Workers continue to face different separation rates, but we impose that separation risk is not realized. Differences in saving behavior therefore arise solely from differences in employment risk, shutting down both income growth and realized non-employment effects.

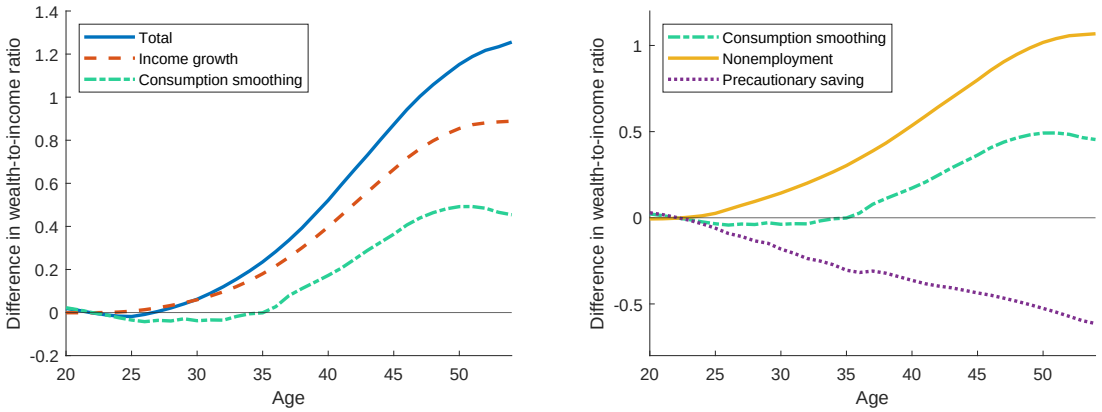
Second, to isolate the *income growth effect*, we eliminate heterogeneity in separation rates by assigning the same age-dependent job-separation profile to all workers. At the same time, we fix wage and human capital paths at their baseline levels for each group. As a result, the two groups

differ only in their life-cycle earnings profiles, while facing identical employment risk. Any difference in wealth accumulation therefore reflects the impact of differential income growth.

Third, to measure the *non-employment effect*, we retain baseline heterogeneity in wages, human capital, and separation rates but eliminate again the realization of non-employment risk. We compute the resulting wealth-to-income ratios for stable and unstable careers and compare them to the baseline model. The difference isolates the asset decumulation associated with realized non-employment spells.

We refer to the sum of the precautionary saving and non-employment effects as the *consumption smoothing effect*, as it captures the role of income risk and buffer-stock dynamics, in contrast to the income growth effect arising from differential earnings trajectories.

Figure 11: Decomposition of difference in wealth-to-income ratio



(a) Income and consumption smoothing

(b) Unemployment and buffer stock

Notes: This figure shows the decomposition of the difference in wealth-to-income ratio between the stable and the unstable careers. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 54.

Figure 11 presents the decomposition of the differences in wealth-to-income ratios between stable and unstable careers over the life cycle. The dominant contribution arises from the income growth effect. As shown in Figure 11a, this component increases steadily with age and accounts for more than 60% of the total wealth difference by the end of working life. The remaining share reflects the consumption smoothing effect.[13]

Figure 11b further separates the consumption smoothing effect into precautionary saving and non-employment components. Because workers in unstable careers face higher employment risk, they accumulate larger precautionary buffers while employed. This generates a negative wealth

[13]The decomposition is not strictly additive, but the fact that the income growth and consumption smoothing components approximately sum to the total gap indicates, however, that the decomposition is quantitatively informative.

gap early in the life cycle, when the precautionary savings motive is the strongest (Gourinchas and Parker, 2002; Cagetti, 2003). In contrast, realized non-employment spells lead to asset decumulation and thus a positive contribution to the wealth gap relative to stable careers. Since unstable careers experience non-employment more frequently, this non-employment effect is positive and increases over time. Early in working life, the precautionary component dominates, rendering the overall consumption smoothing effect negative. Over time, repeated drawdowns during non-employment offset precautionary accumulation, implying that buffer-stock dynamics do not generate sustained wealth growth and that part of the assets accumulated for life-cycle reasons are used to smooth consumption during working life.

4.3 Worker-type heterogeneity

The baseline model emphasizes heterogeneity in job stability and abstracts from direct effects of worker skills on separation rates. In this section, we introduce two empirically disciplined extensions that also incorporate worker heterogeneity in separation rates. First, we allow for *dynamic* worker heterogeneity by letting the job-specific separation rate λ depend on a worker's human capital h. Second, we consider a model extension with *fixed* worker heterogeneity by introducing worker types that permanently differ in their average separation rates.

Across both extensions, the key prediction of the model regarding the differences in wealth accumulation between stable and unstable careers remains quantitatively very similar to the baseline and, if anything, is marginally more pronounced in the dynamic worker case. We therefore conclude that incorporating empirically disciplined worker heterogeneity does not materially affect our main finding on the importance of employment stability for earnings growth and wealth accumulation.

We implement the dynamic worker heterogeneity specification by allowing separation rates λ to depend directly on human capital h, so that separation rates become $\lambda(h)$. As human capital evolves over the life cycle, this introduces *dynamic* worker heterogeneity. The remainder of the model is unchanged; the corresponding Bellman equations are reported in Appendix E.2.

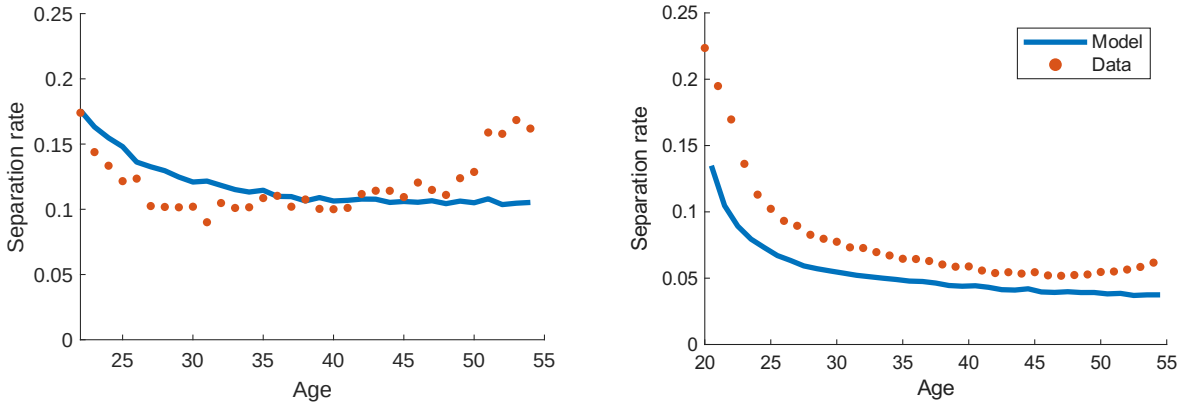
To discipline the effect of human capital on separation rates, we follow Jung and Kuhn (2018), who build on the identification strategy of Dustmann and Meghir (2005) to quantify the contribution of worker heterogeneity for the decline of separation rates over the life cycle. The key idea is to use the difference between separation rates of newly hired workers and those of the average worker to isolate the role of human capital accumulation relative to the job ladder. We flexibly parameterize the human-capital component and calibrate it to match the life-cycle profile of separation rates for newly hired workers, using like for average rates the data from Jung and Kuhn (2018). This yields the following calibrated human-capital component:

$$\lambda(h) = \lambda_l \cdot \max\{(2.5 - 1.5h + 0.1h^2), 0.5\},$$

where λ_l is drawn from the baseline distribution $\{\lambda_l\}_{l=1}^{L}$ and h denotes individual human capital.

Figure 12a shows that the model provides a good fit to the observed decline in separation rates for newly hired workers. Quantitatively, the human-capital component accounts for about one-third of the decline between ages 20 and 40. All remaining parameters are recalibrated as in Section 3.1.

Figure 12: Separation rate



(a) Dynamic worker types

(b) Fixed worker types

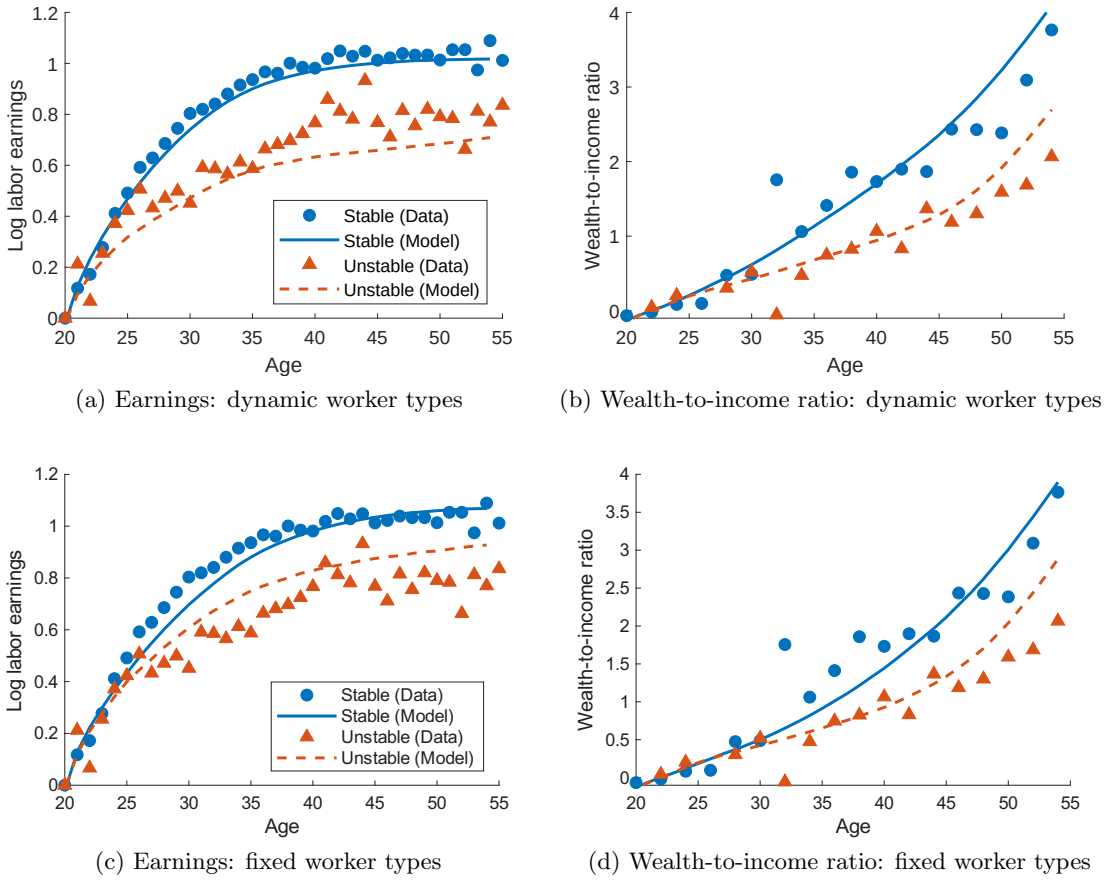
Notes: Panel (a) shows the model fit for the separation rate of workers with tenure less than one year (dynamic worker type). Panel (b) shows the life-cycle profile of mean separation rate in the model with 40% lower separation rate (fixed worker type). The solid lines are the model profiles, while the dots show the estimated empirical profiles from the Panel Study of Income Dynamics.

In the fixed worker heterogeneity specification, we lower the average separation rate to capture salient differences across education groups in the data. Specifically, college-educated workers exhibit separation rates that are about 40% lower than those of non-college workers. At the same time, as documented in Section 2, within-group dispersion in non-employment is very similar across these groups.

We implement this by shifting the marginal distribution of separation rates. Specifically, we reduce the endpoints of the support $[\underline{\lambda}, \overline{\lambda}]$ by 40% while keeping its dispersion unchanged. This preserves within-group heterogeneity while lowering the overall level of separation rates. Figure 12 shows that this modification lowers separation rates uniformly over the life cycle. All other parameters remain fixed.

Figure 13 reports the implications for earnings and wealth accumulation. In the model with dynamic worker heterogeneity, employment stability and human capital accumulation are positively correlated: workers in stable careers accumulate more human capital, which further reduces separation risk, while unstable careers are associated with persistently higher separation rates. As a result, earnings differences between stable and unstable careers become slightly larger, reaching

Figure 13: Earnings and wealth-to-income ratio with worker heterogeneity



(a) Earnings: dynamic worker types

(b) Wealth-to-income ratio: dynamic worker types

(c) Earnings: fixed worker types

(d) Wealth-to-income ratio: fixed worker types

Notes: Panel (a) and (c) show the log of labor earnings of employed workers in the data and in the model, normalized to 0 at age 20. Panel (b) and (d) show the wealth-to-income ratio in the data and in the model, normalized to 0 at age 24. The two top panels show the dynamic worker type model the bottom row the fixed worker type model. See text for details. The solid lines are the model profile, while the dots show the estimated empirical profile from the data. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. Data are from the Panel Study of Income Dynamics.

about 30 log points by age 50.

For wealth accumulation, two counteracting forces emerge. Higher earnings in stable careers increase the wealth gap, while lower earnings in unstable careers reduce wealth accumulation relative to income. On net, the wealth-to-income gap changes only marginally. Importantly, the model continues to match the life-cycle divergence in earnings and wealth closely, and the quantitative differences remain very similar to the baseline.

In the model with fixed worker heterogeneity, differences in earnings and wealth are slightly

attenuated but remain substantial. By the end of working life, stable careers are associated with about 15 log points higher earnings and roughly one additional year of income in wealth. Hence, even within a group facing uniformly lower separation risk, heterogeneity in employment stability continues to generate large differences in wealth accumulation.

Overall, these results show that introducing empirically disciplined worker-type heterogeneity has only modest quantitative effects. While it affects average separation rates (Figure 12), the central mechanism linking employment stability to earnings growth and wealth accumulation remains robust.

5 Heterogeneity in employment stability and the macroeconomy

This section studies the aggregate consequences of heterogeneity in employment stability for wealth inequality and consumption dynamics. While the previous sections established that employment stability shapes individual earnings and wealth accumulation over the life cycle, we now quantify how this heterogeneity translates into macroeconomic outcomes.

We focus on two dimensions. First, we assess how heterogeneity in employment stability contributes to the cross-sectional distribution of wealth. Second, we analyze how it shapes the aggregate consumption response to labor market shocks. Across both topics, a central insight emerges: heterogeneity in employment stability is not only important at the worker level but also has first-order implications for macroeconomic dynamics by affecting both the distribution of wealth and the aggregate response to shocks.

5.1 Effects on wealth inequality

Heterogeneity in individual life-cycle wealth accumulation translates directly into differences in the cross-sectional distribution of wealth. To quantify the contribution of employment stability to wealth inequality, we compare the baseline model with heterogeneous employment stability to a version that eliminates such heterogeneity conditional on age, thereby imposing representative-worker labor market dynamics within each age group.

We compare the wealth distribution in the models with and without heterogeneity. Incorporating the empirically observed dispersion in employment stability, as in our baseline model, increases the 90-10 wealth ratio by 26% relative to the case with representative-worker labor market dynamics. Hence, heterogeneity in employment stability amplifies cross-sectional wealth inequality.

To benchmark the magnitude of this effect, we turn to the SCF. To align the data with the model, we restrict the sample to households with non-negative wealth and compute wealth-to-income ratios across SCF waves. In the SCF data, the 90-10 wealth ratio increased by 32% between 1998 and

2022. Over the same period, overall wealth inequality, as measured by the Gini coefficient, rose from 0.80 to 0.83 (Kuhn and Ríos-Rull, 2025). While the increase in overall wealth inequality is primarily driven by rising concentration at the top, our model provides a complementary mechanism operating at the lower end of the wealth distribution. In particular, heterogeneity in employment stability significantly amplifies wealth inequality among low-income workers. The consumption-smoothing dynamics of workers with unstable employment histories therefore provide an additional explanation for low saving rates among poorer households, consistent with the evidence in Dynan et al. (2004). More broadly, our results highlight that models with representative-worker labor market dynamics understate the role of labor market risk for wealth inequality and miss an important channel through which employment instability shapes savings behavior.

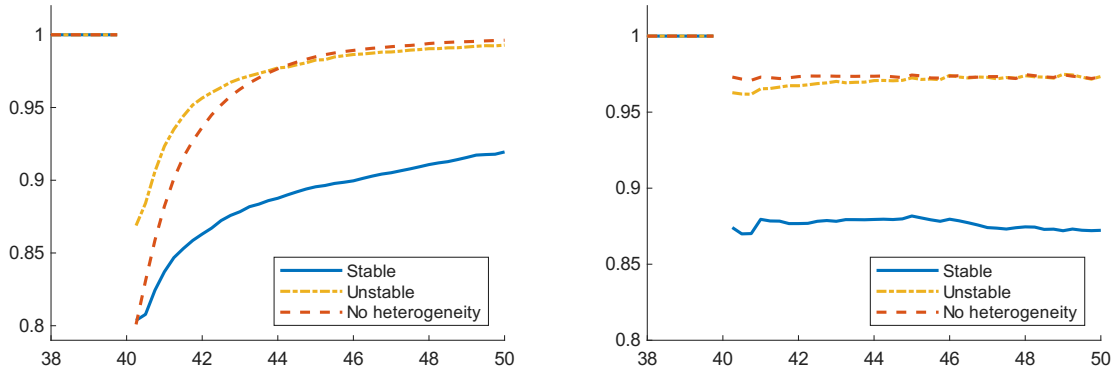
5.2 Consumption dynamics in recessions

As a second macroeconomic implication of heterogeneity in employment stability, we study the consumption response to recessions. It is well documented that workers experience large and persistent earnings losses following job displacement (Jacobson et al., 1993). The structural literature has shown that heterogeneity in employment stability is key to generating such persistent earnings losses while remaining consistent with high average mobility rates (Jung and Kuhn, 2016; Jarosch, 2023). However, this literature had to remain silent on the corresponding implications for consumption dynamics. An exception is Larkin (2019), who shows that portfolio reallocation by workers when jobs become more stable can contribute to a larger consumption drop as, for example, in the aftermath of the Financial Crisis. Instead of looking at portfolio choice, our analysis focuses directly on consumption dynamics and asks how they depend on pre-displacement employment stability. We contrast these responses across stable and unstable careers and to those in a model with representative-worker labor market dynamics that still forms the benchmark.

Figure 14 shows how employment stability shapes the consequences of job loss by comparing workers in stable and unstable employment relationships at age 40. Except for employment stability, these workers are otherwise identical. On impact, job loss leads to earnings declines of 12 percent for workers from stable jobs and 18 percent for workers from unstable jobs (Figure 14a). While these initial losses are similar in magnitude, subsequent dynamics differ markedly. For workers from unstable jobs, earnings recover quickly and exhibit near-complete mean reversion within five years. This reflects the high turnover in unstable employment: even workers who initially avoid displacement are likely to lose their jobs subsequently, so that differences between job losers and stayers dissipate rapidly.

In contrast, job loss from stable employment leads to highly persistent earnings losses. Absent displacement, these workers would have most likely remained employed persistently, so that their counterfactual earnings path absent displacement is persistently much higher. As a result, job loss

Figure 14: Effects of displacement by employment stability



(a) Earnings

(b) Consumption

Notes: This figure shows the evolution of earnings and consumption of workers who become unemployed at age 40 relative to the control group. Workers with low-separation risk are employed in jobs belonging to the bottom quartile of jobs by separation rate at the time of displacement. Workers with high-separation risk are employed in jobs belonging to the top quartile of jobs by separation rate at the time of displacement.

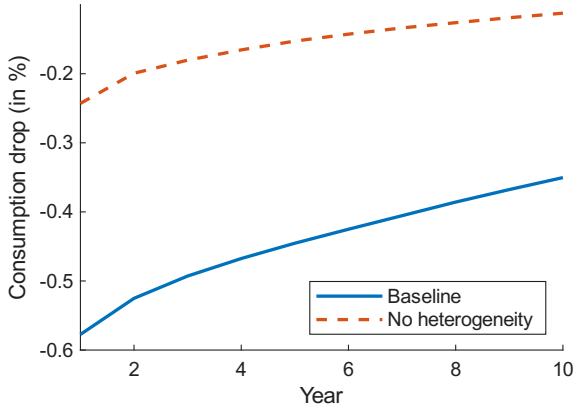
from stable employment generates a large and persistent deviation from this path. More broadly, this highlights that models without heterogeneity in employment stability generate excessive mean reversion and therefore fail to capture persistent earnings dynamics (Jung and Kuhn, 2018).

These differences in earnings dynamics translate directly into heterogeneous consumption responses. Workers adjust consumption based on expectations about both the level and persistence of future income. Workers who lose unstable jobs anticipate a relatively transitory earnings loss and smooth consumption using precautionary savings (Figure 14b). By contrast, workers who lose stable jobs face large and persistent earnings losses. Their wealth allows them to smooth only the transitory component, leading to a sustained permanent decline in consumption of about 13 percent. The larger drop in consumption relative to earnings reflects additional employment losses, which are more pronounced for workers initially in stable jobs. This mechanism directly links to the wealth dynamics in the model: workers in unstable careers hold lower wealth and rely on buffer-stock saving, which facilitates short-run smoothing but limits their ability to absorb persistent shocks.

To assess the aggregate implications, we simulate a deep recession by exogenously displacing 5 percent of employed workers, consistent with the increase in unemployment observed during the Financial Crisis, the COVID recession, and the early 1980s recession. Figure 15 compares the aggregate consumption response in the baseline model with heterogeneity in employment stability to a model with representative-worker labor market dynamics.

Heterogeneity in employment stability strongly amplifies the aggregate consumption response.

Figure 15: Consumption after macroeconomic shock



Notes: This figure shows the consumption drop in response to a macroeconomic shock that increases the non-employment rate by 5 percentage points in year 1. The solid line shows the baseline model with heterogeneity in employment stability. The dashed line shows the model without such heterogeneity, where separation rates are homogeneous conditional on age.

If job losses are concentrated among workers in unstable jobs, aggregate consumption falls by 20 bp. In contrast, if job losses occur among workers in stable jobs, consumption declines by about 60 bp almost three times as large.

These results highlight that not only the magnitude but also the composition of job loss matters for aggregate consumption dynamics. In particular, recessions that disproportionately destroy stable jobs generate more severe and persistent consumption declines. This can be especially relevant in episodes of structural change characterized by sectoral reallocation or the decline of traditionally stable industries (Jaimovich and Siu, 2020; Howes, 2022). More generally, models with representative-worker labor market dynamics understate both the size and persistence of consumption responses, as they abstract from heterogeneity in employment stability that shapes income risk, wealth accumulation, and consumption behavior. A key implication is that not only the magnitude but also the composition of labor market shocks matters for aggregate dynamics. Recessions that disproportionately affect stable employment relationships lead to larger and more persistent declines in consumption than those concentrated among unstable jobs. More broadly, models that abstract from heterogeneity in employment stability understate both wealth inequality and the sensitivity of aggregate consumption to labor market fluctuations, as they miss an important source of variation in income risk and savings behavior.

Taken together, the results in this section highlight that heterogeneity in employment stability plays a central role for macroeconomic outcomes. It contributes to higher wealth inequality by generating persistent differences in earnings growth and savings behavior, particularly at the lower end of the wealth distribution. At the same time, it amplifies aggregate consumption responses to

labor market shocks by creating heterogeneity in the persistence of income losses and the ability of households to smooth consumption.

6 Conclusions

Our analysis started from the observation of large heterogeneity of employment stability in the U.S. labor market. Using data from the PSID, we demonstrate that differences in employment stability are systematically related to wealth accumulation, with workers in more stable jobs being wealthier. We propose a model framework that combines a frictional life-cycle labor market model with an incomplete markets consumption-saving model. We demonstrate that the model is consistent with a wide range of empirical facts on earnings, income, and wealth dynamics. Using the structural model, we explore at the microeconomic level the consequences of differences in employment stability for earnings, consumption, and wealth.

We find a strongly diverging pattern of earnings and wealth for workers with stable and unstable employment over the life cycle. By the end of prime-age working life, workers with stable employment have 20% higher labor earnings and have accumulated additional wealth equivalent to one and a half years of income. Employment instability leads to less income growth, less consumption, and less wealth. The consequences of employment stability for life-cycle dynamics stem from two sources. First, lower employment stability leads to less income growth, from less human capital accumulation and lower wages. Second, lower employment stability requires more precautionary savings to smooth consumption over time, thereby depressing consumption and wealth accumulation. By contrast, starting working life in a lifetime job is associated with stable employment and persistently higher consumption and wealth.

References

Acabbi, E. M., A. Alati, and L. Mazzone (2024): "Human Capital Ladders, Cyclical Sorting, and Hysteresis," *Available at SSRN 4068858*.

Aguiar, M. and E. Hurst (2013): "Deconstructing life cycle expenditure," *Journal of Political Economy*, 121, 437–492.

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109, 659–684.

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85, 693–734.

Barnette, J. (2020): "Wealth After Job Displacement," .

Bartscher, A. K., M. Kuhn, and M. Schularick (2020): "The college wealth divide: Education and inequality in America, 1956-2016," *Available at SSRN 3587685*.

Bayer, C. and M. Kuhn (2018): "Which Ladder to Climb? Wages of workers by job, plant, and education," .

Bewley, T. (undated): "Interest bearing money and the equilibrium stock of capital," Tech. rep., mimeo.

Bilal, A. (2019): "The Geography of Unemployment," .

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.

Borella, M., M. De Nardi, M. Pak, N. Russo, and F. Yang (2023): "The Importance of Modeling Income Taxes over Time: U.S. Reforms and Outcomes," *Journal of the European Economic Association*, 21, 2237–2286.

Bricker, J., L. J. Dettling, A. Henriques, J. W. Hsu, L. Jacobs, K. B. Moore, S. Pack, J. Sabelhaus, J. Thompson, and R. A. Windle (2017): "Changes in US family finances from 2013 to 2016: Evidence from the Survey of Consumer Finances," *Fed. Res. Bull.*, 103, 1.

Burdett, K. and D. T. Mortensen (1998): "Wage differentials, employer size, and unemployment," *International Economic Review*, 257–273.

Cagetti, M. (2003): "Wealth Accumulation over the Life Cycle and Precautionary Savings," *Journal of Business & Economic Statistics*, 21, 339–353.

Cajner, T., I. Güner, and T. Mukoyama (2020): "Gross Worker Flows over the Life Cycle," .

Caplin, A., V. Gregory, E. Lee, S. Leth-Petersen, and J. Sæverud (2023): "Subjective Earnings Risk," Working Paper 31019, National Bureau of Economic Research.

Couch, K. A. and D. W. Placzek (2010): "Earnings losses of displaced workers revisited," *American Economic Review*, 100, 572–89.

Daly, M., D. Hryshko, and I. Manovskii (2016): "Improving the measurement of earnings dynamics," Tech. rep., National Bureau of Economic Research.

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

De Nardi, M., G. Fella, and G. Paz-Pardo (2019): "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 18, 890–926.

Deaton, A. and C. Paxson (1994): "Intertemporal choice and inequality," *Journal of political economy*, 102, 437–467.

Díaz-Giménez, J., A. Glover, and J.-V. Ríos-Rull (2011): "Facts on the distributions of earnings, income, and wealth in the United States: 2007 update," *Federal Reserve Bank of Minneapolis Quarterly Review*, 34, 2–31.

Dustmann, C. and C. Meghir (2005): "Wages, Experience and Seniority," *The Review of Economic Studies*, 72, 77–108.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004): "Do the rich save more?" *Journal of political economy*, 112, 397–444.

Gourinchas, P.-O. and J. A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89.

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

——— (2021b): "What do data on millions of US workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89, 2303–2339.

Hagedorn, M. and I. Manovskii (2010): "Search frictions and wage dispersion," *Manuscript, University of Zurich and University of Pennsylvania*.

Hall, R. E. (1982): "The Importance of Lifetime Jobs in the US Economy," *American Economic Review*, 72, 716–724.

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006," *Review of Economic dynamics*, 13, 15–51.

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

Howes, C. (2022): "Why does structural change accelerate in recessions? The credit reallocation channel," *Journal of Financial Economics*, 144, 933–952.

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

Huggett, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control*, 17, 953–969.

Iacono, R. and M. Ranaldi (2020): "The wage curve across the wealth distribution," *Economics Letters*, 196, 109580.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993): "Earnings losses of displaced workers," *The American economic review*, 685–709.

Jaimovich, N. and H. E. Siu (2020): "Job Polarization and Jobless Recoveries," *The Review of Economics and Statistics*, 102, 129–147.

Jarosch, G. (2023): "Searching for job security and the consequences of job loss," *Econometrica*, 91, 903–942.

Jung, P., P. Korfmann, and E. Preugschat (2023): "Optimal regional labor market policies," *European Economic Review*, 152, 104318.

Jung, P. and M. Kuhn (2016): "Earnings Losses and Labor Mobility over the Lifecycle," .

——— (2018): "Earnings Losses and Labor Mobility Over the Life Cycle," *Journal of the European Economic Association*, jvy014–jvy014.

Kaas, L., E. Lalé, and N. Siassi (2023): "Job Ladder and Wealth Dynamics in General Equilibrium," Tech. rep.

Karahan, F. and S. Ozkan (2013): "On the persistence of income shocks over the life cycle: Evidence, theory, and implications," *Review of Economic Dynamics*, 16, 452 – 476.

Kimball, M. S., C. R. Sahm, and M. D. Shapiro (2009): "Risk Preferences in the PSID: Individual Imputations and Family Covariation," *American Economic Review*, 99, 363–68.

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

Krebs, T., M. Kuhn, and M. L. Wright (2015): "Human capital risk, contract enforcement, and the macroeconomy," *American Economic Review*, 105, 3223–72.

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

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

Kudlyak, M. and F. Lange (2014): "Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories," *IZA DP*.

Kuhn, M. (2013): "Recursive Equilibria in an Aiyagari-Style Economy with Permanent Income Shocks," *International Economic Review*, 54, 807–835, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12018.

Kuhn, M., I. Manovskii, and X. Qiu (2026 forthcoming): "The geography of job creation and job destruction," *Journal of Monetary Economics*.

Kuhn, M. and J.-V. Ríos-Rull (2016): "2013 Update on the US earnings, income, and wealth distributional facts: A View from Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 37, 2–73.

Kuhn, M. and J.-V. Ríos-Rull (2025): "Income and Wealth Inequality in the United States: An Update Including the 2022 Wave," *Quarterly Review*, 45.

Kuhn, M., M. Schularick, and U. Steins (forthcoming): "Income and Wealth Inequality in America, 1949-2016," *Journal of Political Economy*.

Larkin, K. P. (2019): "Job Risk, Separation Shocks and Household Asset Allocation," in *2019 Meeting Papers*, Society for Economic Dynamics, 1058.

Lise, J. (2012): "On-the-Job Search and Precautionary Savings," *The Review of Economic Studies*, 80, 1086–1113.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432–67.

Meghir, C. and L. Pistaferri (2004): "Income variance dynamics and heterogeneity," *Econometrica*, 72, 1–32.

Michelacci, C. and H. Ruffo (2015): "Optimal Life Cycle Unemployment Insurance," *American Economic Review*, 105, 816–59.

Morchio, I. (2020): "Work Histories and Lifetime Unemployment," *International Economic Review*, 61, 321–350.

Mortensen, D. T. and C. A. Pissarides (1994): "Job creation and job destruction in the theory of unemployment," *The review of economic studies*, 61, 397–415.

Nam, L. (2022): "Optimal progressive pension systems in a life-cycle model with heterogeneity in job stability," .

OECD (2015): *Pensions at a Glance 2015*.

Pfeffer, F. T., R. F. Schoeni, A. Kennickell, and P. Andreski (2016): "Measuring wealth and wealth inequality: Comparing two US surveys," *Journal of economic and social measurement*, 41, 103–120.

Pinheiro, R. and L. Visschers (2015a): "Unemployment risk and wage differentials," *Journal of Economic Theory*, 157, 397–424.

——— (2015b): "Unemployment risk and wage differentials," *Journal of Economic Theory*, 157, 397–424.

Primiceri, G. E. and T. Van Rens (2009): "Heterogeneous life-cycle profiles, income risk and consumption inequality," *Journal of monetary Economics*, 56, 20–39.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

Stevens, A. H. (1997): "Persistent effects of job displacement: The importance of multiple job losses," *Journal of Labor Economics*, 15, 165–188.

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

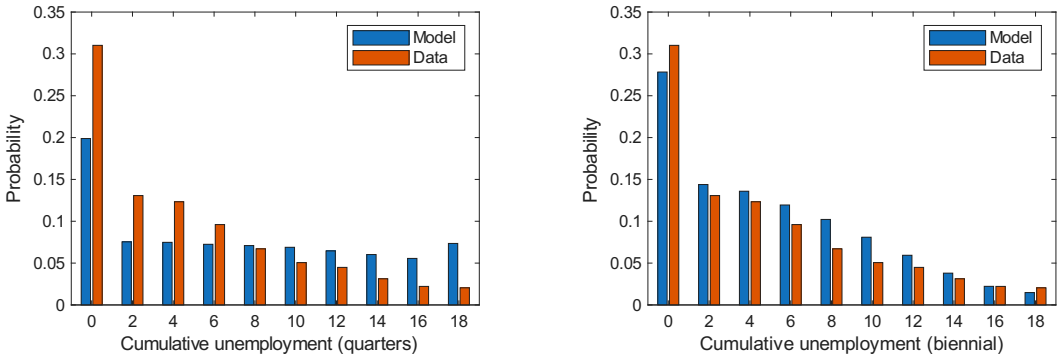
A Appendix

A.1 PSID sample

The PSID provides annual survey data through 1997 and switches to a biennial frequency starting in 1999. This change introduces a potential bias in our construction of accumulated non-employment duration, as we no longer observe individuals' complete labor market histories. As a result, some individuals may be misclassified as having stable or unstable careers, since these classifications rely on observed labor market histories.

Using the model from Section 3, we assess the extent of misclassification that would arise if labor market outcomes were observed only biennially. Specifically, when assigning workers to stable and unstable career categories at the end of prime-age working life, we construct accumulated non-employment duration using observations from every second year. We then compute the share of workers who are misclassified under this alternative measure and adjust for the proportions of the PSID sample observed before and after 1999. We find that the share of mismatched workers is approximately 2%.

Figure A.1: Accumulated non-employment duration over the life cycle



(a) Before correction

(b) After correction

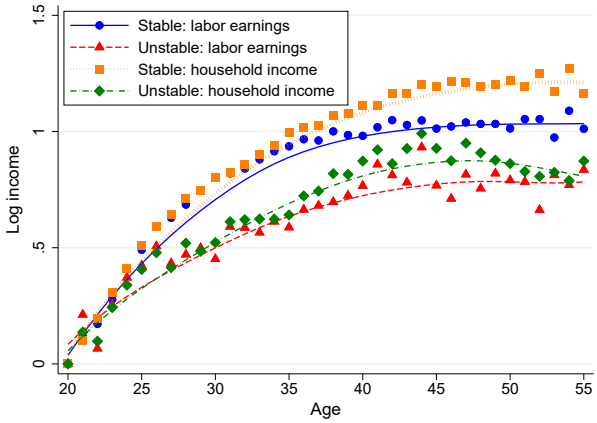
Notes: Panel (a) shows the accumulated non-employment duration (in quarters) of workers at the end of prime-age working life in the model and in the PSID data. Panel (b) shows the distribution in the model when households are observed in very second year.

In Figure A.1, we compare the distribution of accumulated non-employment duration at the end of prime-age working life in the model and in the data. While the model-generated distribution of non-employment duration is more evenly spread than in the data, constructing accumulated non-employment duration using observations from every second year allows the model to closely match the empirical distribution.

A.2 Life-cycle profile of household income

In Section 2, we show that the life-cycle profile of labor earnings exhibits lower earnings growth for workers with stable employment. Since labor earnings here refer to those of the household head in the PSID data, this measure may miss an important source of insurance at the household level provided by other family members. We therefore repeat the analysis using total pre-tax income of all household members, including transfers.

Figure A.2: Log income



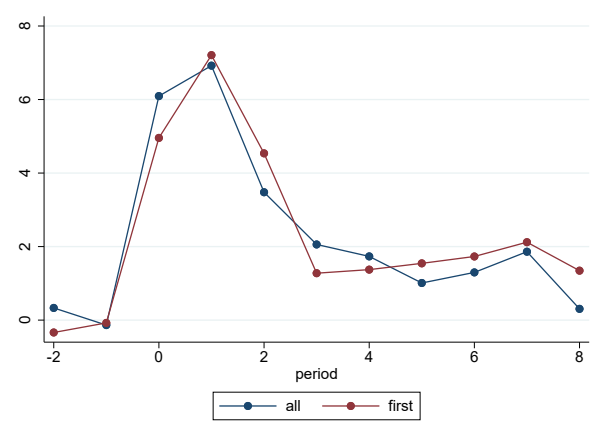
Notes: The figure shows the life-cycle profiles of log labor income and log total household income for stable and unstable groups of workers in terms of employment stability. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 50. The profiles of log income are normalized to zero at age 20. Data are from the Panel Study of Income Dynamics.

Figure A.2 plots the life-cycle profiles of log total household income alongside the profiles of log labor earnings from Figure 3. While household income grows faster than the labor earnings of the household head for both groups, the gap between stable and unstable careers is even more pronounced: at age 50, workers with stable employment experience about 0.35 log points higher income growth. If workers with unstable careers would receive substantial insurance from other household members, differences in income growth between stable and unstable workers should shrink when moving from individual labor earnings to total household income. Instead, the gap becomes even larger, indicating that additional household income does not offset the weaker earnings growth of unstable workers. This points to limited within-household risk sharing for those with unstable employment histories.

B Employment stability: worker vs. job component

This appendix provides supporting evidence for the interpretation of employment stability as a job component. We examine employment dynamics around the termination of jobs and document subsequent labor market outcomes and analyze the differences in job destruction rates across firms.

Figure B.1: Non-employment rate after displacement



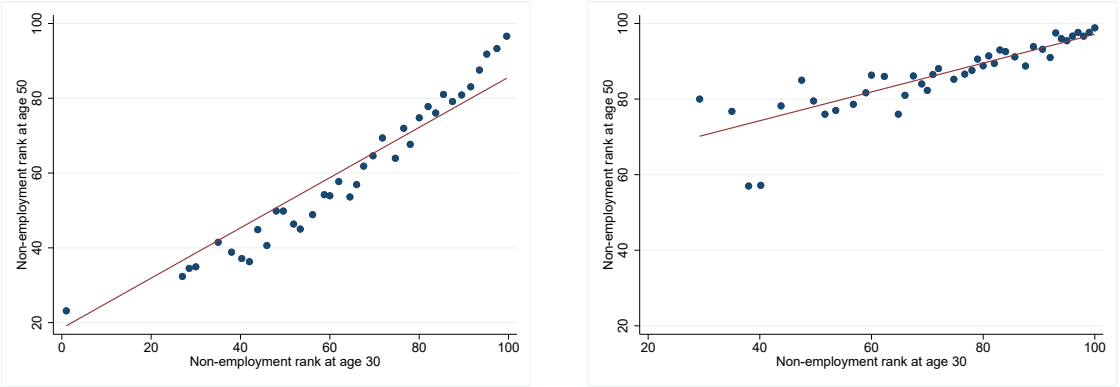
Notes: This figure shows the non-employment rate after a displacement event in period 0.

Figure B.1 plots non-employment rates in an event-time window around the end of a stable job. Period 0 denotes the separation event, and the horizontal axis measures time relative to job loss. The figure compares two groups: all stable jobs and first stable jobs.

Several patterns emerge. First, non-employment increases sharply at the time of separation, reflecting the mechanical transition out of employment. More importantly, non-employment remains persistently elevated in the periods following job loss. This persistence indicates that the end of a stable job is associated with lasting disruptions to workers' employment trajectories rather than a brief interruption. Second, the pattern is qualitatively similar when focusing only on first separations. The magnitude of post-separation non-employment is comparable, suggesting that the consequences of losing a stable match are not limited to later-career jobs but are also present earlier in the employment history.

Next, we analyze the persistence of non-employment over the life cycle. If differences in non-employment duration across workers were driven solely by worker fixed effects, we would expect a perfect correlation between non-employment accumulated at young ages and non-employment accumulated later in life for all workers. Figure B.2 plots the rank of accumulated non-employment between ages 25 and 30 against the rank of accumulated non-employment at the end of prime working age. Figure B.2 shows a strong positive correlation: workers who experience frequent non-employment at young ages also tend to accumulate substantial non-employment by the end of

Figure B.2: Persistence of non-employment duration



(a) All

(b) Workers with non-employment duration ¿ 2 years between age 30 and 48

Notes: Panel (a) plots the rank in accumulated non-employment duration at young ages (between 25 and 30) against the rank in accumulated non-employment duration at the end of prime working life (up to age 52). Panel (b) shows the same relationship for workers who accumulate more than two years of non-employment between ages 30 and 48.

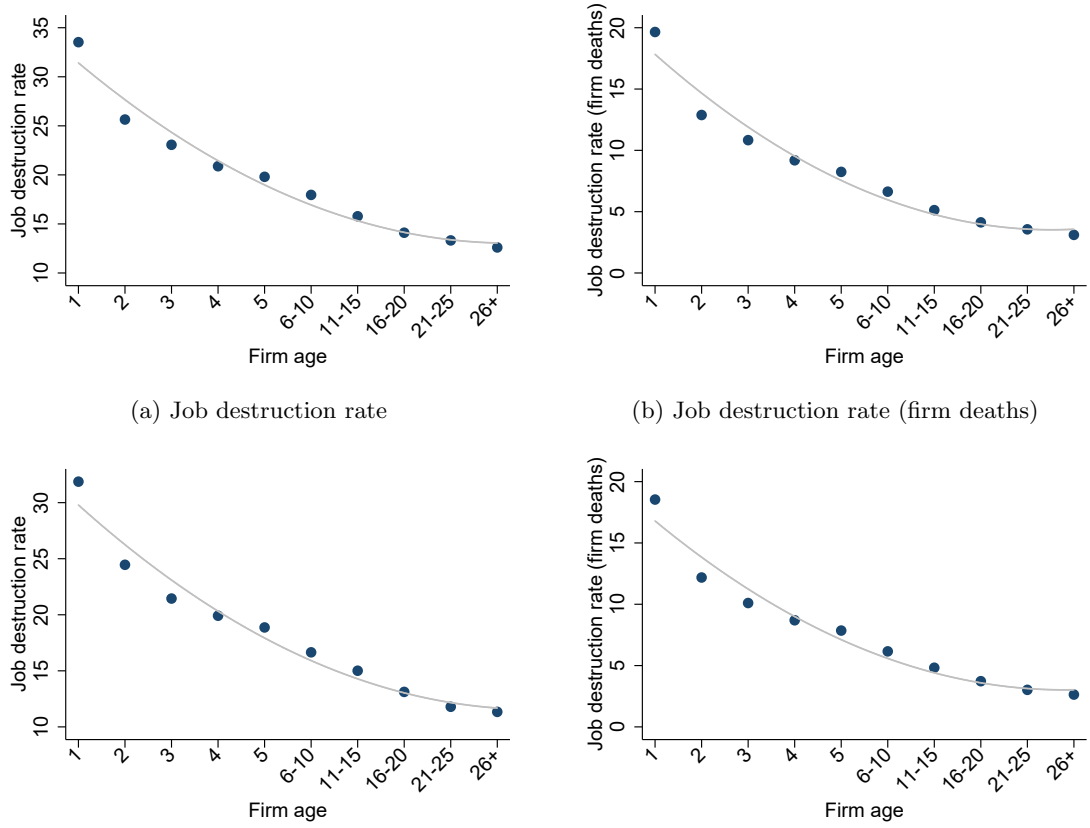
prime age.

However, this relationship weakens when focusing on workers with substantial non-employment spells during mid-career. Figure B.2 shows that the correlation declines markedly for workers who accumulate more than two years of non-employment between ages 31 and 48. This finding suggests that persistent differences in non-employment are not driven solely by worker fixed effects, but also reflect shocks and life-cycle dynamics. It further supports the view that employment stability is shaped to a large extent by differences across jobs rather than by permanent worker characteristics alone.

In the following, we further explore the potential sources of heterogeneity in employment stability. We show in Figure B.3 an evidence from the BDS on heterogeneity in job loss probabilities across employers of different age. We consider two definitions of job loss: total job destruction rate (Figure B.3a) and job loss due to firm closure (Figure B.3b). We remove year and industry fixed effects in both cases.

In Figure B.3a, we observe large heterogeneity in job loss across employers, with the least stable employer having job loss rates that are twice as large as the most stable employers. Such differences in observed job loss could be the result of differences in job-to-job transitions or worker quits. Figure B.3b therefore considers the more restrictive definition of job loss where we consider only job loss due to firm closure. For this case, we find the differences in job stability to be even larger, with the least and the most stable employers differing by a factor of four. Such large

Figure B.3: Heterogeneity in job destruction rates by firm age



(a) Job destruction rate

(b) Job destruction rate (firm deaths)

(c) Job destruction rate

(d) Job destruction rate (firm deaths)

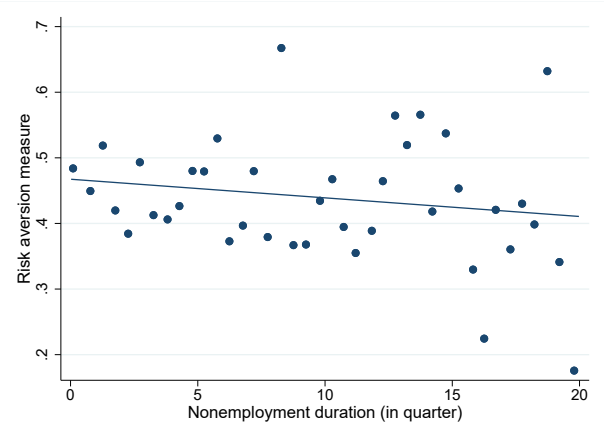
Notes: Panel (a) and (c) show the relationship between job destruction rate and firm age from the BDS. Panel (b) and (d) show the relationship between the job destruction rate due to firm deaths and firm age. Job destruction rates are computed as the number of jobs destroyed over the last 12 months divided by average employment, where the denominator is computed as the average of employment for periods t and $t-1$. In Panel (a) and (b), we control for year and industry fixed effects. In Panel (c) and (d), we control for year and MSA fixed effects.

differences in the probability of job loss across employers are also supported by existing research. Larkin (2019) documents large heterogeneity in separation rates into unemployment in U.S. CPS data, and Jarosch (2023) documents such heterogeneity across German employers. Figure B.3c and Figure B.3d shows that controlling for year and MSA fixed effects yields very similar results and heterogeneity in job destruction rates persists. [14]

To examine the relationship between wealth and accumulated non-employment duration, we also account for workers' risk aversion, as differences in risk tolerance may influence employment

[14]The BDS data do not provide publicly available data where industry and geographical breakdown is available in the same file.

Figure B.4: Non-employment duration and risk aversion



Notes: This figure shows a binned scatter plot of accumulated non-employment duration of workers at age 50 and the risk aversion measure.

stability choices. We measure risk tolerance using PSID data, following (Kimball et al., 2009). Figure B.4 shows no evident correlation between accumulated non-employment duration and the measure of risk aversion among workers at age 50.

Table B.1: Wealth-to-income ratio and non-employment

Wealth-to-income ratio	(1)	(2)
Non-employment (in quarters)	-0.05^{***}	-0.05^{***}
	(0.01)	(0.01)
Interaction of age and risk aversion	No	Yes
Observations	16594	15439
R^2	0.43	0.43

Notes: The dependent variable is wealth-to-income ratio. The sample consists of workers between age 25 and 50. Column (1) shows the regression results of wealth-to-income ratio on non-employment duration of workers in quarters, controlling for age, industry, and worker-fixed effects. Column (2) includes the interaction effect of age and risk aversion measure. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.1 presents a regression analysis of wealth-to-income ratios on workers' accumulated non-employment duration. Although this specification uses all observations rather than focusing only on workers at the end of their prime-age careers, we still find a significant negative association between wealth-to-income ratios and non-employment duration. Controlling for worker fixed effects yields an estimate of -0.05. Including risk aversion produces very similar results, suggesting that differences in risk preferences play only a minor role in explaining the differences in wealth accumulation across

workers.

Taken together, these findings support the view that employment stability reflects an important job-specific component. Following displacement, workers face persistently higher non-employment risk, indicating that the loss of a stable match has lasting consequences. Moreover, workers' positions in the employment-stability distribution evolve over the life cycle as shocks accumulate. We also document substantial heterogeneity in job loss risk across employers—ranging from a twofold difference using overall separations to a fourfold difference when focusing exclusively on firm closures—highlighting pronounced variation in job stability across firms. These results further underscore that employment stability is not solely a fixed worker effect and reinforce the interpretation of employment stability as a characteristic of jobs.

C Employment stability and wealth accumulation

C.1 PSID data

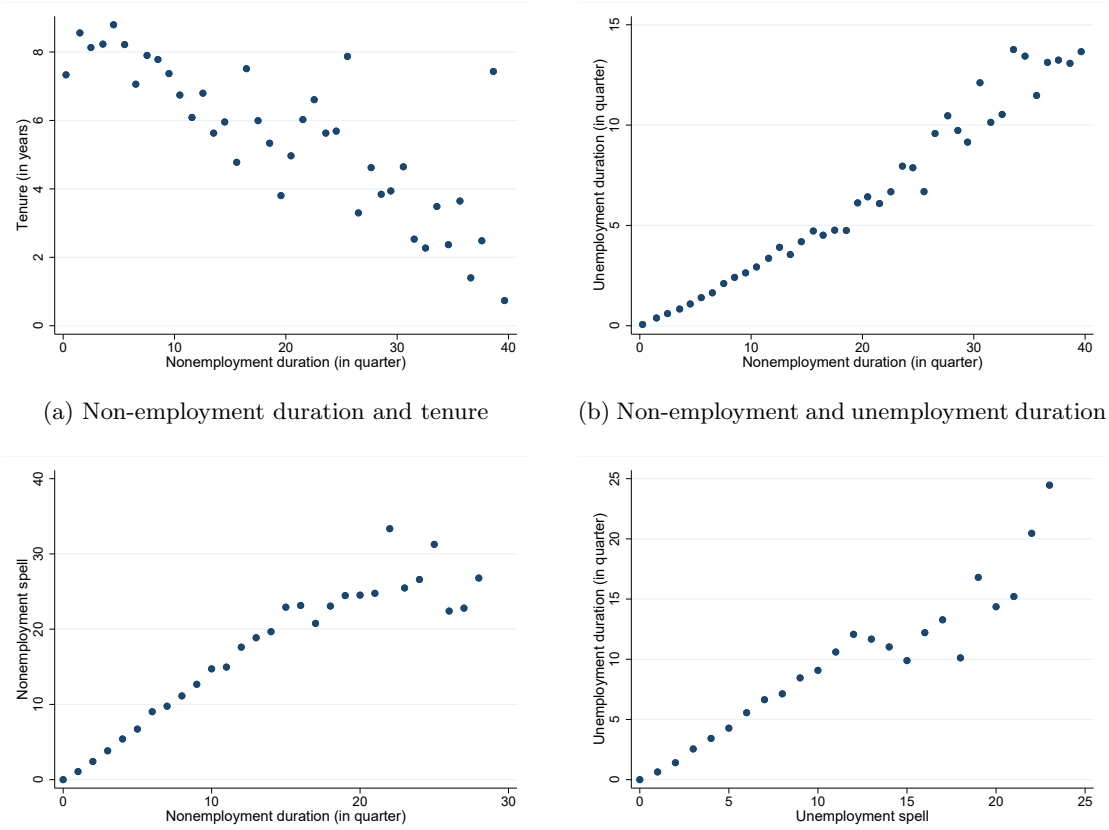
This section presents a series of robustness checks for the relationship between wealth accumulation and employment stability documented in Section 2. In particular, we consider alternative measures of employment stability, including unemployment duration and spells, non-employment spells, and tenure. We then reassess the relationship with wealth-to-income ratios to examine whether the previously documented negative correlation persists across these alternative measures.

Figure C.1 reports the correlations among the different measures of employment stability. We find a robust positive correlation between non-employment and unemployment duration, between non-employment duration and the number of non-employment spells, and between unemployment duration and the number of unemployment spells. In contrast, non-employment duration and tenure are negatively correlated. Taken together, these findings support the use of non-employment duration as a meaningful measure of employment stability.

Figure C.2 shows a binned scatter plot of wealth-to-income ratios against non-employment of workers at age 50 after controlling for education and industry. Even after adding the additional controls, there is a significant negative relationship between wealth-to-income ratios and the non-employment duration.

We next examine the relationship between wealth-to-income ratios and alternative measures of employment stability in Figure C.3. Using unemployment duration, the number of unemployment spells, and the number of non-employment spells, we continue to find a robust negative correlation with wealth-to-income ratios. The relationship also persists when lifetime income, rather than current household income, is used to construct the wealth-to-income ratio.

Figure C.1: Different measures for employment stability



(a) Non-employment duration and tenure

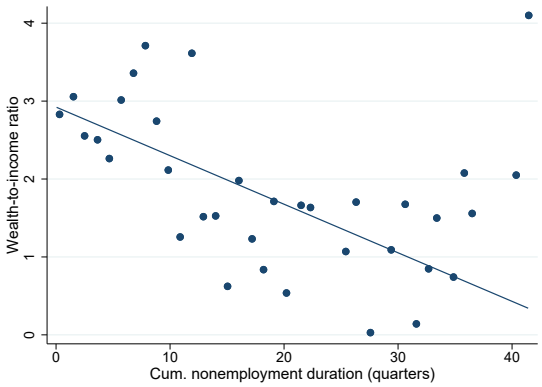
(b) Non-employment and unemployment duration

(c) Non-employment duration and spell

(d) Unemployment duration and spell

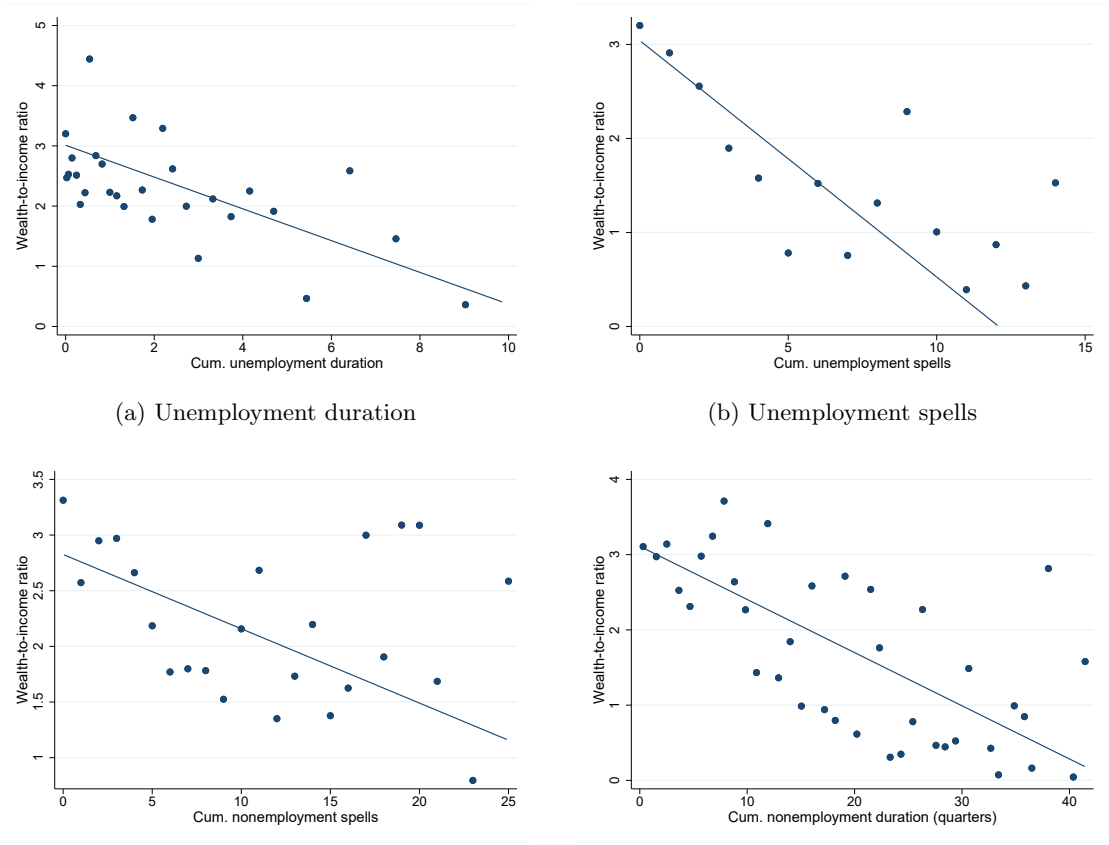
Notes: This figure shows binned scatter plots of different measures for employment stability. Panel (a) shows tenure plotted against non-employment duration. Panel (b) displays the correlation between unemployment and non-employment duration. Panel (c) and (d) show the correlation between duration and spell for non-employment and unemployment, respectively.

Figure C.2: Wealth-to-income ratios and non-employment



Notes: This figure shows binned scatter plot of wealth-to-income ratios against non-employment of workers at age 50 after controlling for education and industry.

Figure C.3: Robustness: wealth-to-income ratio and employment stability



(a) Unemployment duration

(b) Unemployment spells

(c) Non-employment spell

(d) Lifetime income

Notes: This figure shows binned scatter plots of wealth-to-income ratios for different measures for employment stability and income. Panel (a) shows this relationship for unemployment duration, Panel (b). Panel (b) displays the correlation between unemployment and non-employment duration. Panel (c) and (d) show the correleation between duration and spell for non-employment and unemployment, respectively.

C.2 SCF data

Using the Survey of Consumer Finances (SCF), we also document the large heterogeneity in job stability in the cross-sectional data of U.S. households. The SCF is a triennial household survey that has become the key resource on the distribution of income and wealth for the United States (see, for example, Kuhn and Ríos-Rull, 2016; Bricker et al., 2017; Kuhn et al., forthcoming). Besides the detailed information on household income and wealth, the SCF also offers information on household members' labor market situation. We pool data across survey waves from 1992 to 2016 and restrict the sample to households with employed household heads of ages 20 to 60.

In addition to the results from the PSID data, we explore the relationship between the labor market situation and wealth accumulation in the SCF. For our analysis, we pool data across survey waves from 1992 to 2016 and restrict the sample to households with employed household heads ages 20 to 60. As our model will abstract from self-employment, we drop households with self-employed household heads and households with extreme wage observations, defined as wages lower than 75 percent of the minimum wage.[15] Additionally, we exclude the top 1 percent of households by wealth and earnings as we do not provide a theory of the very right tail of these distributions, and, similarly, we exclude households in the bottom 1 percent by earnings and households with negative wealth.

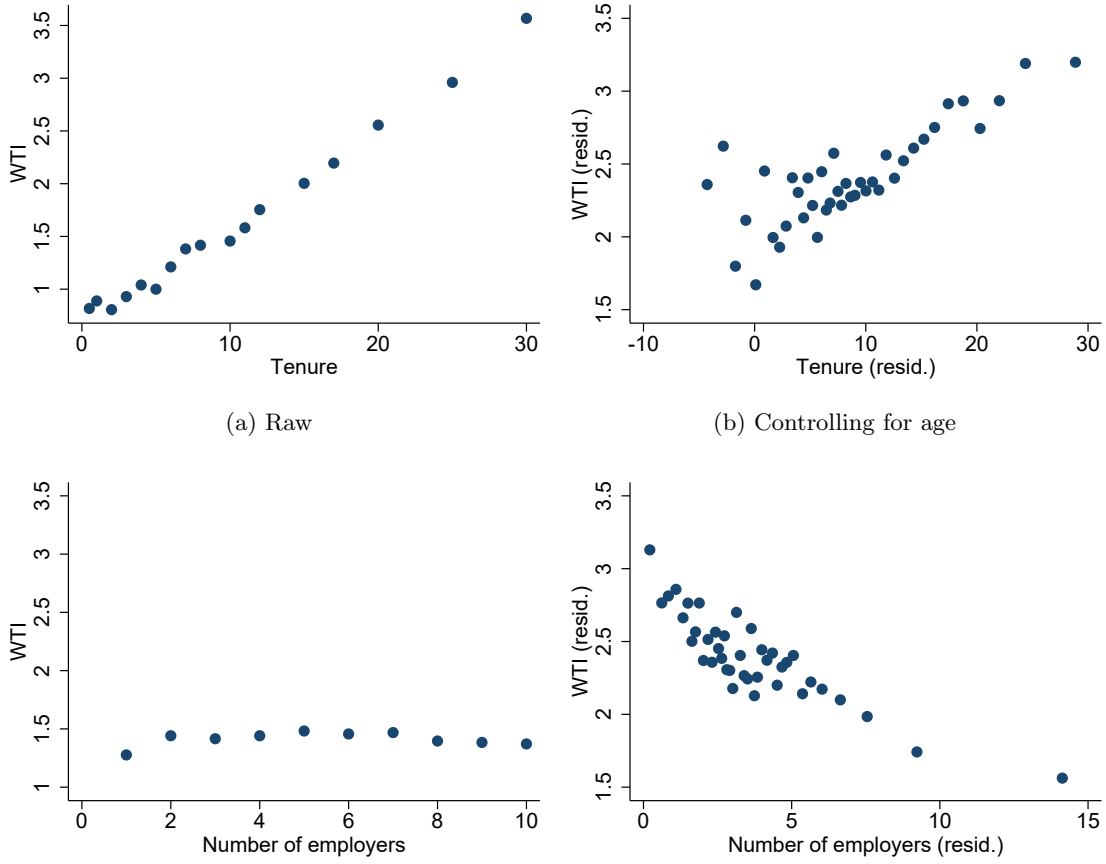
Regarding the construction of variables, we consider household wealth as the difference between household assets and debt. Household income is gross income from all sources including transfers, and earnings is income from wages and salaries.[16] We control for income differences nonparametrically by always considering wealth-to-income ratios. Job stability itself is unobserved in the data, and we only observe retrospectively whether an employer-employee relationship has been stable by looking at a worker's employer tenure or the number of employers during a worker's career.[17] Using these statistics as measures of job stability will lead to measurement error regarding the true level of underlying job stability for two reasons. First, realized tenure can be high despite low job stability due to luck. Second, realized tenure can be low despite the worker having a stable job because the worker might have received a better job opportunity and therefore changed employers. We will therefore interpret the observed correlations from this section through the lens of our structural model in Section 3. In the structural model, we will also consider realized job tenure when mapping the model to the data. We will also have job-to-job mobility as observed in the data so that we can impose consistency in the measurement of job stability between the model and data. The structural model will, in addition, offer us the opportunity to consider meaningful

[15] We rely on individual hours and earnings information in the SCF data to construct wages.

[16] We follow Bricker et al. (2017) and Kuhn and Ríos-Rull (2016) for the construction of these variables.

[17] Employer tenure is defined as the years a person has already been working for his/her current employer. The number of employers a person has worked for is defined as the number of full-time jobs lasting one year or more that a person had over his/her entire career.

Figure C.4: Wealth-to-income ratios, tenure, and number of employers



(a) Raw

(b) Controlling for age

(c) Raw

(d) Controlling for age

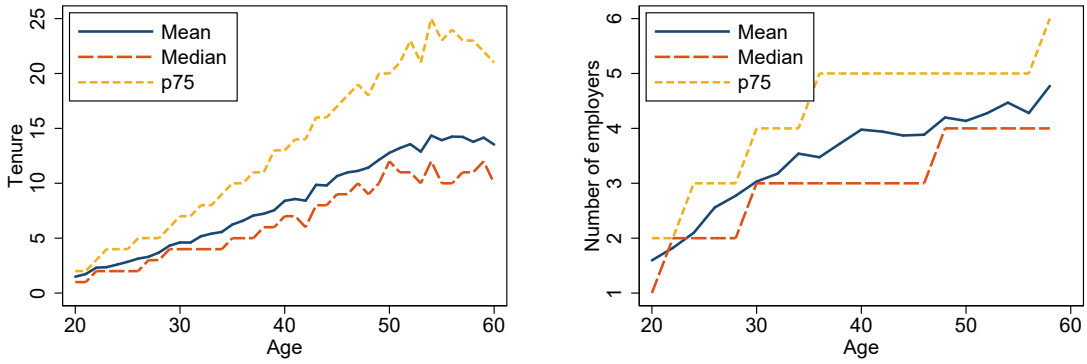
Notes: This figure shows binned scatter plots of wealth-to-income ratios against tenure or number of employers for which a person has worked full-time jobs lasting one year or more. In panels (a) and (c), each dot represents a median wealth-to-income ratio for a given bin. Panels (b) and (d) show binned scatter plots of wealth-to-income ratios against tenure or number of employers after nonparametrically controlling for age. Means have been added back to residualized variables to facilitate interpretation of the scale. Data are from the 1992-2016 waves of the Survey of Consumer Finances. Observations are weighted with SCF sample weights.

counterfactuals to study the causal effect of job stability on wealth accumulation in isolation.

Figure C.5 shows life-cycle profiles for tenure and the number of employers in the SCF data. We find that both profiles are positively correlated with age. Looking at the mean, the median, and the 75th percentile of the tenure distribution in Figure C.5a, we observe a spreading out of the distribution as workers age. As pointed out in Hall (1982), the typical U.S. worker has a stable employment history. At age 60, more than 50 percent of workers have been with their employer for 10 years, and almost a quarter of workers at age 60 have been at the same employer for at least 25

years. The life-cycle profiles for the number of employers in Figure C.5b provide a similar picture. We find that the mean number of employers increases linearly up to age 40 when the growth starts slowing down in the second part of working life. On average, an American worker has worked for four employers at the end of his/her working life.

Figure C.5: Tenure and number of employers over the life cycle



(a) Employer tenure

(b) Number of employers

Notes: Panel (a) shows the life-cycle evolution of the cross-sectional distribution of tenure (in years). Panel (b) shows the life-cycle evolution of the cross-sectional distribution of number of employers for which a person has worked full-time jobs lasting one year or more. Two-year age bins are used in panel (b). Data are from the 1992-2016 waves of the Survey of Consumer Finances. Observations are weighted with SCF sample weights.

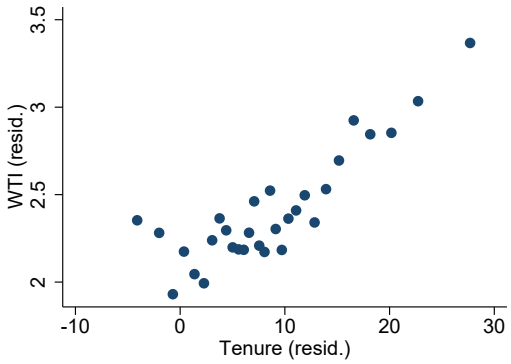
Figure C.4 provides the estimates of the empirical measures of job stability and their relationship to wealth-to-income ratios. In Figure C.4a, we observe an almost linear relationship between tenure and wealth-to-income ratios. This raw correlation could be the result of tenure and wealth-to-income ratios both increasing with age. In Figure C.4b, we therefore show the correlation between wealth-to-income ratios and tenure, controlling nonparametrically for age. In this case, we still find a positive relationship between job stability and wealth accumulation but with a smaller slope than before. Qualitatively, a positive slope implies that per dollar of income, workers in more stable jobs have more wealth, or, in short, workers with more stable jobs are wealthier.

Figure C.4b offers a second interesting observation. While the relationship in the raw data appears almost perfectly linear (Figure C.4a), the relationship turns into a U-shape for low tenures after controlling for age. This U-shape relationship means that workers who have low tenure relative to their age group tend to have higher wealth-to-income ratios. As we will see below, this is a characteristic property of the model when job losers accumulate precautionary wealth before layoffs and then get a negative shock to income, so that wealth-to-income ratios increase. Consumption smoothing of job losers will lower wealth-to-income ratios over time and allow them to converge back to their target wealth-to-income ratio if no further job loss will occur.

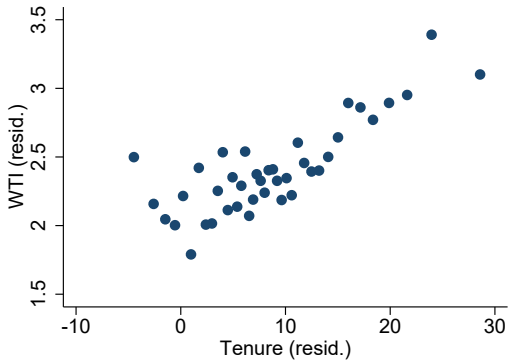
Figures C.4c and C.4d corroborate the previous findings using the total number of employers as an alternative measure of job stability. The effect of age on the relationship can now be seen even more clearly. Before controlling for age (Figure C.4c), there is no apparent relationship between wealth-to-income ratios and the number of employers. After age effects are taken out, we find a declining relationship between the number of employers and wealth-to-income ratios. Following the interpretation that more employers are a consequence of less stable jobs, we find again that job stability (fewer employers) is positively related to wealth-to-income ratios. These results on tenure and number of employers point toward a positive relationship between job stability and wealth accumulation, with workers in more stable jobs being wealthier.

To make sure that the observed relationship between labor market experience and wealth is not driven by demographic characteristics of workers, systematic differences in jobs across industries and occupations, or differences in risk attitudes among workers, we perform further robustness checks where we control for additional observable characteristics of households in the SCF. Figure C.6 shows that our findings are not affected by the inclusion of additional controls. In the first column, we nonparametrically control for age, education, occupation, and industry. The relationship between tenure and wealth-to-income ratios (the top row) remains unaffected and significant. The same holds for the relationship between the number of employers and wealth-to-income ratios (the bottom row). In the second step, we additionally control for differences in risk attitudes of workers by nonparametrically controlling for different levels of risk attitudes, as elicited in the SCF survey. As shown in the second column, the relationship between labor market experience and wealth is not affected by the inclusion of all these additional variables.

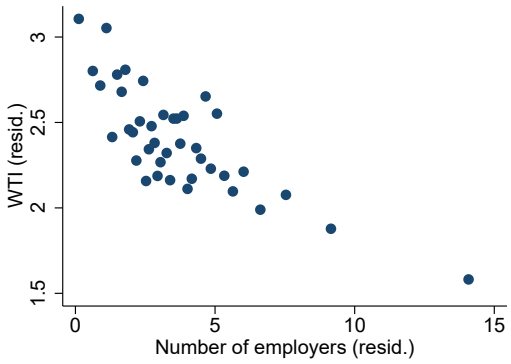
Figure C.6: Wealth-to-income ratios, tenure, and number of employers (with additional controls)



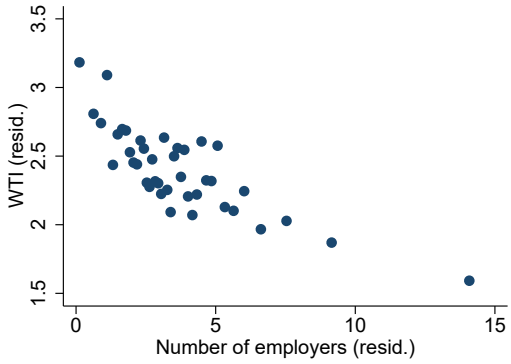
(a) Controls: age, education, occupation, industry



(b) Controls: age, education, occupation, industry, risk attitude



(c) Controls: age, education, occupation, industry



(d) Controls: age, education, occupation, industry, risk attitude

Notes: This figure shows binned scatter plots of wealth-to-income ratios against tenure or number of employers for which a person has worked full-time jobs lasting one year or more. Panels (a) and (c) show binned scatter plots of wealth-to-income ratios against tenure or number of employers after nonparametrically controlling for age, education, occupation and industry. In panels (b) and (d), we additionally nonparametrically control for risk attitudes. Means have been added back to residualized variables to facilitate the interpretation of the scale. Data are from the 1992-2016 waves of the Survey of Consumer Finances. Observations are weighted with SCF sample weights.

D Additional model details

D.1 Value functions for the transition phase

In the transition phase, agents solve a fixed point problem. As a result, value functions do not have any time index. The value functions for the transition phase follow directly the value functions of the working phase. The only difference is that they comprise a probability ψ that at the end of

the period, the worker retires and enters the retirement phase. All decisions are otherwise identical to the working phase.

The value function of an employed worker at the beginning of the transition phase V_e^T is given by the expectations over the employment status as an outcome of the separation stage,

$$V_e^T(a, w, \lambda, h) = \lambda V_n^{T,P}(a, w, h) + (1 - \lambda)V_e^{T,I}(a, w, \lambda, h),$$

where $V_n^{T,P}$ denotes the value function of an unemployed worker at the production state and $V_e^{T,I}$ denotes the value function of an employed worker at the investment stage.

At the investment stage, an employed agent makes a human capital investment decision:

$$V_e^{T,I}(a, w, \lambda, h) = \max_{t\in[0,1]} -\kappa t^2 + p_H(t)V_e^{T,P}(a, w, \lambda, h^+) + (1 - p_H(t))V_e^{T,P}(a, w, \lambda, h).$$

The Bellman equation of an employed agent at the production stage is

$$\begin{aligned} V_e^{T,P}(a, w, \lambda, h) = \max_{\{c,a'\geq 0\}} & u_j(c) + \beta\Big[\psi V_r(a', w, h, j_r = 1) + & (9)\\ & (1 - \psi)\left(\pi_e V_e^{T,S}(a', w, \lambda, h) + (1 - \pi_e)V_e^T(a', w, \lambda, h)\right)\Big] & (10)\\ s.t. \quad & c = (1 + r)a + y(w, h, e) - a', \end{aligned}$$

where V_r denotes the agent's value function in the retirement phase, $V_e^{T,P}$ denotes the employed agent's value function at the production stage, $V_e^{T,S}$ denotes the employed agent's value function at the search stage, and V_e^T denotes the value function of an employed worker at the beginning of the transition phase. The value function of an employed worker at the search stage of the transition phase is

$$V_e^{T,S}(a', w, \lambda, h) = \sum_{s=1}^{N_w}\sum_{k=1}^{N_\lambda} \max\Big\{ \underbrace{V_e^T(a', w, \lambda, h)}_{\text{staying in current job}}, \underbrace{V_e^T(a', w_s, \lambda_k, h)}_{\text{accepting outside offer}} \Big\} f(w_s, \lambda_k),$$

where N_w is the number of wage realizations in the support of the offer distribution and N_λ is the number of realizations for separation rates in the support of the offer distribution.

The value function of a non-employed worker at the production stage is

$$\begin{aligned} V_n^{T,P}(a,w,h) &= \max_{\{c,a'\geq 0\}} u_j(c) + \beta\Big(\psi V_r(a',w,h,j_r=1) + \\ &\quad (1-\psi)\left(\pi_n V_n^{T,S}(a',w,h) + (1-\pi_n)V_n^{T}(a',w^-,h)\right)\Big) \\ s.t. &\quad c = (1+r)a + y(w,h,u) - a'. \end{aligned}$$

For the value function of an unemployed worker at the search stage, we get

$$V_n^{T,S}(a',w,h) = \sum_{s=1}^{N_w}\sum_{k=1}^{N_\lambda} \max\Big\{ \underbrace{V_n^T(a',w^-,h)}_{\text{staying unemployed}}, \underbrace{V_e^T(a',w_s,\lambda_k,h)}_{\text{accepting job offer}} \Big\} f(w_s,\lambda_k).$$

D.2 Model solution and estimation

D.2.1 Solving the model

We solve the model using backward induction and apply on-grid search to solve the consumption-saving and effort choice problem. We discretize the state space for assets, wages, job destruction probability, and human capital. Denoting the asset grid by $\mathcal{A}$, the wage grid by $\mathcal{W}$, the grid for job destruction probabilities by $\mathcal{L}$, and the grid for human capital by $\mathcal{H}$, we construct the state space as the Cartesian product of the separate grids $\mathcal{A}\times\mathcal{W}\times\mathcal{L}\times\mathcal{H} = \{a_1,...,a_{Na}\}\times\{w_1,...,w_{Nw}\}\times\{l_1,...,l_{Nl}\}\times\{h_1,...,h_{Nh}\}$. The upper bounds on the grids are chosen large enough so that they do not constitute a constraint on the optimization problem.

We assume that both wages and job destruction probabilities have after standardization a truncated exponential marginal distribution with support of $[0,1]$.[18] To allow for a possible correlation between both marginal distributions, we construct a joint distribution over standardized wages and job destruction probability $F(w^*,\lambda^*)$ using Frank's copula C_θ, where the value of θ determines the correlation between w^* and λ^*. Finally, we discretize this distribution into bins that correspond to grids for w and λ.

Using these discretized grids and the joint distribution, we store the computed value functions and policy rules as finite-dimensional arrays. Finally, we use these obtained policy rules and randomly generated shocks to simulate life cycles of 200,000 agents.

[18] We standardize the support of wages and job destruction probabilities to allow for an easier numerical implementation of the joint distribution. We discuss the details in Section 3.1.

D.2.2 Parameter estimation

We estimate some of the model parameters using a simulated method of moments. We minimize the sum of squared percentage deviations of the model-implied age profiles from their empirical counterparts. Life-cycle profiles of separation, job-to-job and job-finding rate, tenure (mean, median and 75th percentile), log earnings (mean and variance) and wealth-to-income ratio are used in the estimation. If the parameter vector is denoted θ, then the objective function we minimize is

$$
\begin{aligned}
\min_{\theta} \quad & \sum_{a=21}^{55} \left(\frac{\pi_s(a,\theta) - \hat{\pi}_s(a)}{\hat{\pi}_s(a)} \right)^2 + \sum_{a=21}^{55} \left(\frac{\pi_{eo}(a,\theta) - \hat{\pi}_{eo}(a)}{\hat{\pi}_{eo}(a)} \right)^2 \\
+ \quad & \sum_{a=21}^{55} \left(\frac{\pi_{ne}(a,\theta) - \hat{\pi}_{ne}(a)}{\hat{\pi}_{ne}(a)} \right)^2 + \sum_{a=21}^{55} \left(\frac{t_{mean}(a,\theta) - \hat{t}_{mean}(a)}{\hat{t}_{mean}(a)} \right)^2 \\
+ \quad & \sum_{a=21}^{55} \left(\frac{t_{median}(a,\theta) - \hat{t}_{median}(a)}{\hat{t}_{median}(a)} \right)^2 + \sum_{a=21}^{55} \left(\frac{t_{p75}(a,\theta) - \hat{t}_{p75}(a)}{\hat{t}_{p75}(a)} \right)^2 \\
+ \quad & \sum_{a=21}^{55} \left(\frac{e_{mean}(a,\theta) - \hat{e}_{mean}(a)}{\hat{e}_{mean}(a)} \right)^2 + \sum_{a=25}^{55} \left(\frac{e_{var}(a,\theta) - \hat{e}_{var}(a)}{\hat{e}_{var}(a)} \right)^2 \\
+ \quad & \sum_{a=23}^{55} \left(\frac{wti(a,\theta) - \hat{wti}(a)}{\hat{wti}(a)} \right)^2,
\end{aligned}
$$

where the empirical profiles are denoted with a hat.

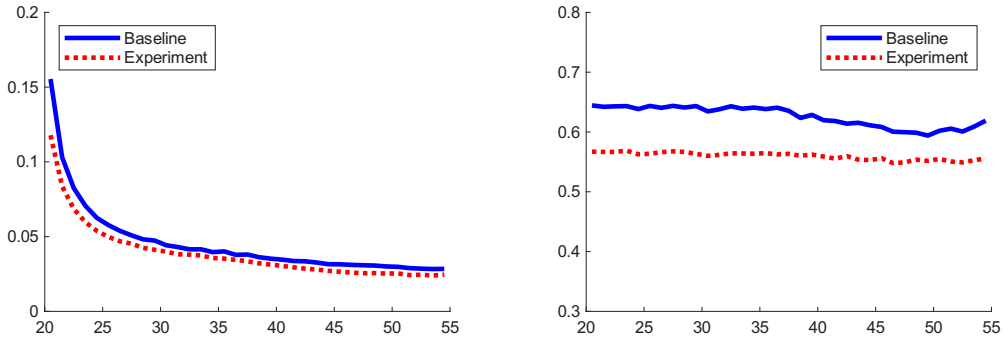
D.2.3 Discussion of identification of model parameters

In the following, I discuss the identification of parameters and provide a detailed mapping between parameter blocks and the empirical moments that discipline them. Because all parameters interact in the model, I clarify the economic channels through which parameters shape observable life-cycle patterns and demonstrate how targeted moments discipline them.

We organize the identification strategy around three core components: job-offer dynamics, the accumulation of human capital, and preferences together with savings behavior.

To begin, the offer arrival rates for employed and unemployed workers, denoted by π_e and π_u, are pinned down using observed life-cycle patterns in job-to-job transitions and job-finding rates. An increase in these arrival rates raises worker mobility and speeds up the process of matching with higher-quality jobs. Consequently, the identification of these parameters relies on the age profiles of employment-to-employment and non-employment-to-employment transitions. To illustrate the role of these parameters, we conduct a comparative statics exercise in which we vary their values and examine the resulting changes in the model's predicted life-cycle transition patterns.

Figure D.1: Change in job-arrival rates π_e and π_u

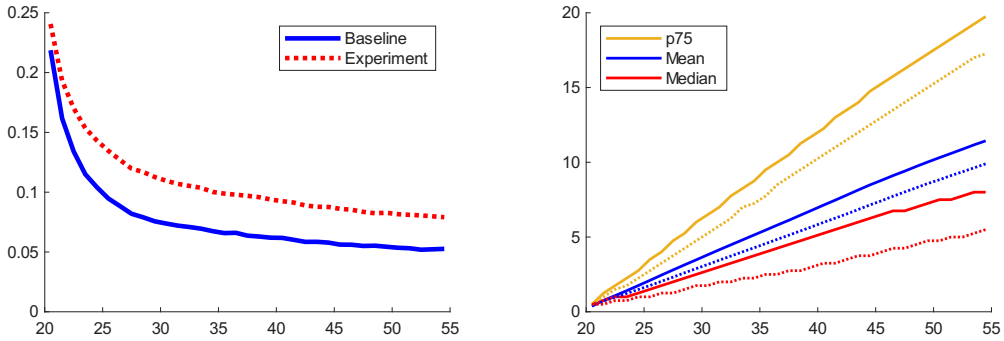


(a) Job-to-job rate to change in π_e (b) Job finding rate to change in π_u

Notes: These figures document the sensitivity of transition rates to parameter perturbations in the job-arrival probabilities π_e (reduced from 0.43 to 0.33) and π_u (reduced from 0.85 to 0.75).

Figure D.1 shows the response of job-to-job transitions and job-finding rates to changes in the job-offer arrival parameters. When we lower the arrival rates for employed and unemployed workers, π_e and π_u, by 10 percentage points, both life-cycle profiles shift downward. This reflects the reduced frequency with which workers receive job offers. The contraction in job-to-job transitions is especially strong among younger workers. Early in the career, mobility is a key mechanism for progressing up the job ladder. With fewer job offers arriving, opportunities to switch jobs decline, leading to a particularly marked reduction in job-to-job transitions at younger ages.

Figure D.2: Change in job-offer distribution shape parameter ψ_λ



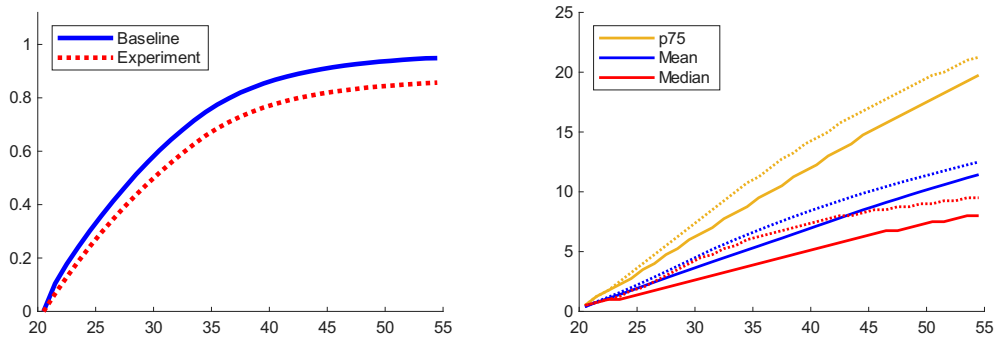
(a) Separation rate (b) Tenure distribution

Notes: These figures document the sensitivity of separation rate and tenure distribution to a parameter perturbation in the job-offer distribution shape parameter ψ_λ (increased from 0.50 to 0.8).

Next, we identify the joint job-offer distribution, $f(\lambda, w)$, by matching the joint dynamics of

wages, tenure, and worker flows. In particular, the marginal distribution of job stability, governed by ψ_λ, is pinned down using the life-cycle pattern of separation rates together with the level and dispersion of tenure. This parameter determines the extent to which job offers are concentrated in relatively unstable versus highly stable matches. Figure D.2 illustrates the implications of shifting this distribution toward fewer highly stable jobs. When stable matches are relatively scarce, the sorting process into durable employment relationships becomes more gradual. Early in their careers, workers disproportionately accept positions with higher separation risk and consequently experience repeated job loss before transitioning into more stable matches. This mechanism raises aggregate separation rates and depresses average tenure, particularly at younger ages. Over the life cycle, the slower reallocation into stable jobs also postpones the build-up of long-tenure matches, resulting in a flatter tenure profile.

Figure D.3: Change in job-offer distribution shape parameter ψ_w

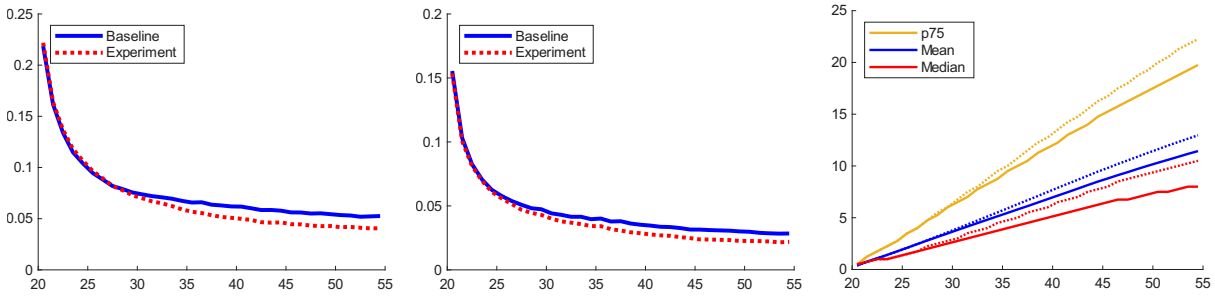


(a) Earnings (b) Tenure distribution

Notes: These figures document the sensitivity of separation rate and tenure distribution to a parameter perturbation in the job-offer distribution shape parameter ψ_w (increased from 0.54 to 3.28).

The parameter ψ_w, governs the marginal distribution of wage offers, and in particular the frequency with which high-wage opportunities arrive. We identify this parameter by targeting the slope and curvature of average wage growth over the life cycle, along with the dispersion of tenure at younger ages. From an economic perspective, ψ_w determines how easily workers can move up the job ladder through job-to-job transitions. Figure D.3 illustrates the effects of a lower arrival rate of high-wage offers. When such opportunities become less common, upward mobility declines. Because young workers depend heavily on job-to-job moves to generate rapid wage increases, reduced wage upgrading translates into more modest earnings growth early in the career. In particular, wage growth between ages 20 and 30 becomes noticeably flatter. Moreover, slower progression toward higher-paying jobs postpones sorting into more productive and potentially more stable matches, thereby influencing the dispersion of tenure at the beginning of the working life.

Figure D.4: Change parameter θ for correlation between wages and separation rates



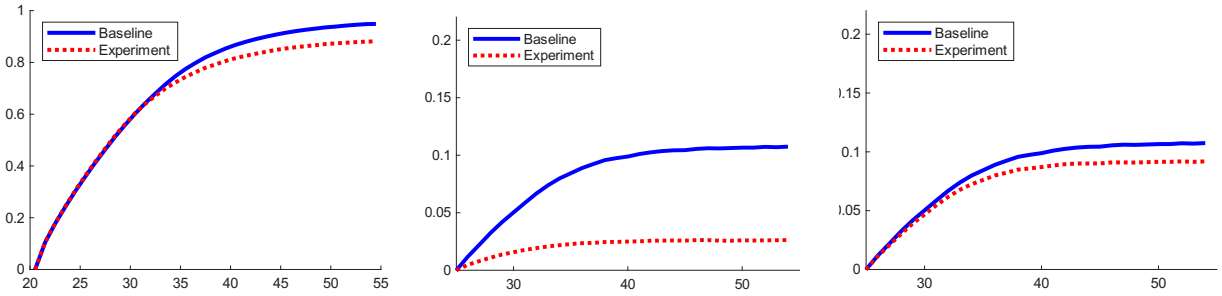
(a) Separation rate (b) Job-to-job rate (c) Tenure distribution

Notes: These figures document the sensitivity of separation rate, job-to-job rate, and tenure distribution to a parameter perturbation in the correlation θ between wages and separation rates (increased from 0.49 to 0.74).

The parameter θ, which governs the correlation between wages and separation risk, is identified by matching the joint life-cycle patterns of job-to-job transitions, separation rates, and tenure dispersion at older ages. Economically, θ captures the extent to which higher-paying jobs are also more stable. As illustrated in Figure D.4, a stronger positive correlation implies that workers who move up the wage ladder simultaneously transition into jobs with lower separation risk. This strengthens the sorting process over the life cycle: high-wage workers become increasingly concentrated in stable matches, while low-wage workers remain disproportionately in less stable employment. As a result, mobility and separation rates decline more closely in tandem with age. This enhanced sorting mechanism has direct implications for tenure dynamics. When wages and job stability are strongly positively correlated (that is, when θ is high), workers who successfully sort into high-paying jobs tend to remain in those positions for extended periods. Consequently, tenure becomes increasingly concentrated at higher levels toward the end of prime age, compressing the lower tail of the tenure distribution. The combined patterns in mobility, separations, and tenure dispersion therefore discipline the magnitude of θ in the model.

Third, the utility cost of human capital investment, κ, is identified using the late-career flattening of mean earnings, when wage growth is primarily driven by human capital accumulation rather than job mobility (see Figure D.5a). An increase in κ raises the cost of effort devoted to skill acquisition, leading workers to reduce their investment. Consequently, earnings growth slows, and the wage profile becomes flatter at older ages, when human capital accumulation is a key determinant of income dynamics. The parameters governing the evolution of human capital, ρ and $\overline{p}_H$, are disciplined by the life-cycle behavior of average earnings and earnings dispersion. Faster skill upgrading increases within-cohort inequality over time, as workers accumulate human capital at heterogeneous rates (Figures D.5b and D.5c). A decline in the average probability of skill

Figure D.5: Change in parameters governing human capital κ, $\overline{p}_H$, and ρ

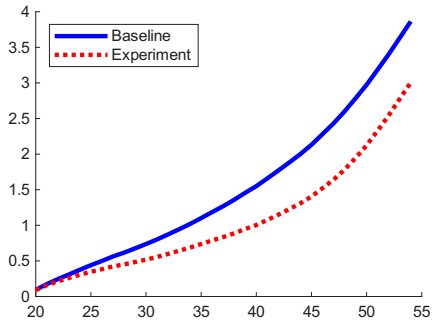


(a) Log earnings to a change in κ (b) Variance of log earnings to a change in $\overline{p}_H$ (c) Variance of log earnings to a change in ρ

Notes: These figures document the sensitivity of log earnings and variance of log earnings to parameter perturbations in the parameters governing human capital accumulation κ (increased from 0.40 to 0.54), $\overline{p}_H$ (reduced from 0.071 to 0.035), and ρ (reduced from 0.984 to 0.982).

upgrading, $\overline{p}_H$, lowers both the level and the variance of earnings across all ages. In contrast, a reduction in the persistence parameter, ρ, which shapes the age profile of upgrading probabilities, flattens the life-cycle pattern of earnings dispersion and mainly compresses inequality toward the end of prime age.

Figure D.6: Change in time discount parameter β



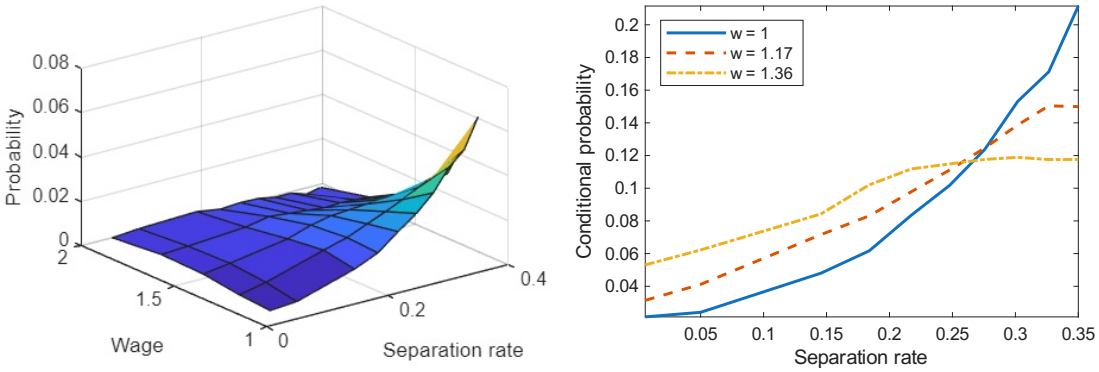
Notes: This figure documents the sensitivity of wealth-to-income ratio to parameter perturbation in the time discount factor (reduced from 0.975 to 0.97).

The discount factor, β, is identified by matching the life-cycle profile of the wealth-to-income ratio, which reflects the accumulation of assets over time. A lower value of β implies greater impatience, as workers attach less value to future utility. This shift in intertemporal preferences increases the relative importance of current consumption and reduces incentives to save for retirement. As a consequence, asset holdings decline at all ages, leading to a downward shift in the life-cycle wealth profile. Figure D.6 illustrates this adjustment.

D.2.4 The estimated job-offer distribution

Panel D.7a of Figure D.7 shows the estimated job-offer distribution for wages and separation rates, which is asymmetric with most of the probability mass concentrated at low-wage, unstable jobs. Additionally, we also find that wages and separation rates are negatively correlated, implying that high-wage jobs have a low separation rate and low-wage jobs have high separation rates. Panel D.7b additionally shows the conditional distribution of separation rates for different levels of wages. The distribution of separation rates in low-wage jobs first-order stochastically dominates the distribution of separation rates in high-wage jobs.

Figure D.7: Job-offer distribution



(a) Joint distribution

(b) Conditional distributions

Notes: Panel (a) shows the estimated job-offer distribution over wages and separation rates used in the numerical implementation. Panel (b) shows the conditional distribution of separation rates for different levels of wage.

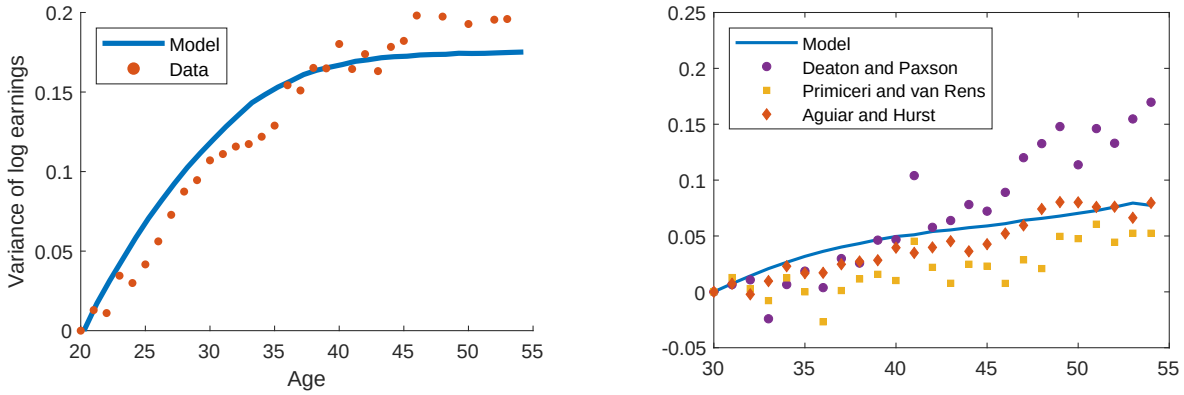
D.3 Additional results on model fit

D.3.1 Earnings and consumption inequality

In the next step, we explore the model's ability to account for the life-cycle pattern of earnings variance. Figure D.8a shows the life-cycle increase in log earnings variance and the close match between model and data. We see the typical almost linear increase in the variance by age in both the model and the data. Labor market search models oftentimes struggle to account for this increase, as discussed in Lise (2012), Jung and Kuhn (2018), or Hubmer (2018). Augmenting models of job search with differences in human capital accumulation provides one way to account for the observed increase (Jung and Kuhn, 2018; Hubmer, 2018). While we build on this approach, we further refine it by endogenizing the human capital accumulation decision. As we discuss in Section

4.1, the endogenous human capital accumulation is key for the question of this paper because it provides a mechanism to transform transitory differences in search outcomes into persistent earnings differences.

Figure D.8: Earnings and consumption inequality



(a) Earnings variance

(b) Consumption variance

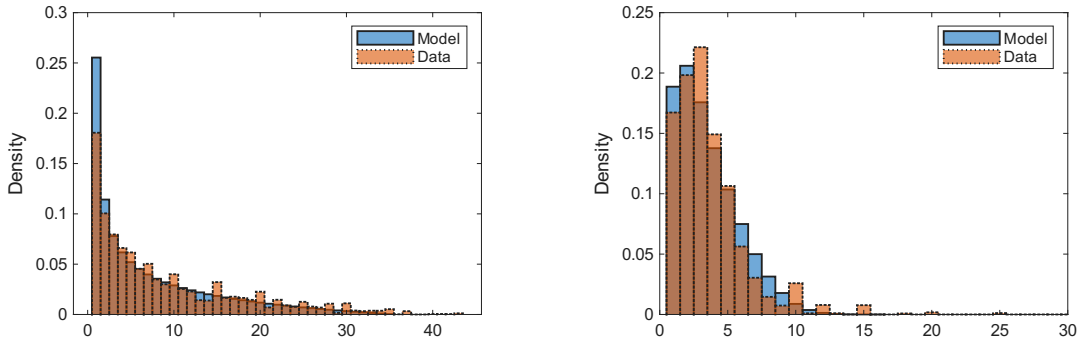
Notes: Panel (a) shows the variance of log earnings in the PSID data, normalized to 0 at age 20. The blue line is the model profile, while the red dots show the estimated empirical profile from the PSID data. Panel (b) shows the variance of log consumption from Deaton and Paxson (1994), Primiceri and Van Rens (2009), and Aguiar and Hurst (2013), normalized to 0 at age 30.

Figure D.8b demonstrates that the model also aligns with the empirical estimates of the life cycle increase in the variance of log consumption. Empirical estimates of consumption variance differ across studies (Deaton and Paxson, 1994; Primiceri and Van Rens, 2009; Aguiar and Hurst, 2013), and the model falls in the middle of the range of existing estimates. Comparing Figures D.8 and D.8b, we also note that the increase in the variance of consumption is roughly one-third lower than that for earnings. Hence, consumption is partly insulated from earnings dynamics.

D.3.2 Cross-sectional distributions of tenure and the number of employers

Figure D.9 shows the cross-sectional distribution of tenure and the number of employers for which a worker has worked for at least one year during her working life. We combine all workers and show the corresponding distribution using histograms. When pooling data from the model, we assume that each age group has the same share in the pooled sample.

Figure D.9: Cross-sectional distribution of tenure and number of employers



(a) Tenure distribution

(b) Number of employers

Notes: This figure shows the distribution of tenure and the number of employers from the SCF and the model when all ages are pooled together. Red bars are the SCF data; blue bars are the model equivalent. In line with the SCF design, only employment spells with a duration of at least a year are used in the simulated data.

D.3.3 Life-cycle earnings dynamics

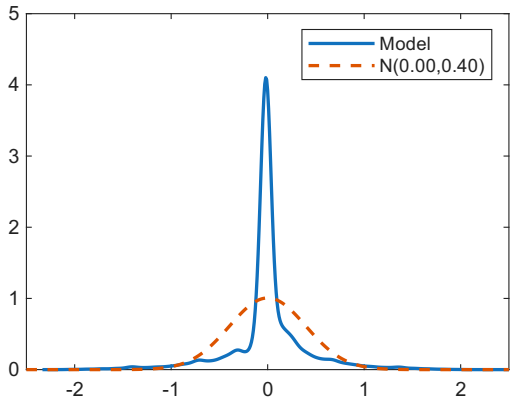
Results presented in Section 3.2 demonstrate that the model matches the life-cycle profiles for means and variances. Here we provide additional evidence that the model also provides a good fit along other dimensions. To explore the fit for earnings dynamics, we compare how the model-implied earnings dynamics align with statistical representations of earnings processes as typically estimated in applied work and used to parametrize exogenous earnings dynamics in consumption-saving models. Such a description of earnings dynamics by a reduced-form statistical representation allows for a straightforward comparison of earnings dynamics between model and data. First, we perform a standard decomposition of earnings dynamics into a permanent and transitory component and estimate the variances of the innovation terms (Meghir and Pistaferri, 2004; Blundell et al., 2008; Heathcote et al., 2010). Second, as emphasized in Guvenen et al. (2021a), we look at the higher moments of the distribution of earnings growth. Third, we decompose earnings growth into contributions from human capital accumulation and job switching and demonstrate that the model aligns with the evidence on early career wage growth by Topel and Ward (1992). We close with a discussion of the model predictions for frictional wage dispersion.

In the first case, we estimate the variance of the permanent component of earnings dynamics using simulated earnings series from the model aggregated to an annual frequency. We apply the identification approach, as in Blundell et al. (2008), to the simulated data. We estimate a variance of the permanent component of 0.025 that falls well within the range of empirical estimates. Blundell et al. (2008) estimate time-varying variances of the permanent component ranging from 0.01 to

0.03 for the period from 1980 to 1990.[19] Empirical estimates for the variance of transitory shocks are harder to compare as they also comprise the contribution from measurement error that is likely substantial in the data, so it is not surprising that our finding is that the empirical estimates for the variance of the transitory shocks (0.03-0.05) are substantially larger than the model-implied estimate (0.016). We interpret the difference as the contribution from measurement error but also unmodeled earnings components such as bonuses and overtime pay.

In a second step, we consider the findings by Guvenen et al. (2021a), who emphasize that earnings growth rates are not normally distributed but exhibit large negative skewness and high excess kurtosis. As has been demonstrated by Hubmer (2018), these patterns can be well explained by a life-cycle version of the job ladder model with a human capital process. The same also applies to our model. Figure D.10 shows the distribution of one-year earnings growth rates in the model with a superimposed normal distribution that has the same standard deviation. Earnings changes in our model are left-skewed and strongly clustered around 0.

Figure D.10: One-year earnings changes



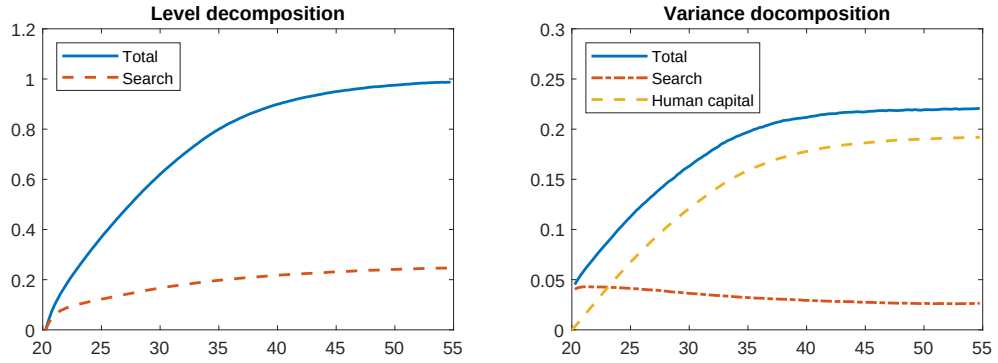
Notes: Figure shows the estimated kernel density of the model-based one-year earnings changes superimposed on Gaussian densities with the same standard deviation. Earnings growth net of average age effect shown. Kernel density estimation with a bandwidth of 0.05 used.

Last, we use the model to decompose life-cycle earnings growth and the increase in the variance into a component from search for higher wages and a component from human capital accumulation. Figure D.11 shows the decomposition for mean and variance and highlights human capital accumulation as the key driver of life-cycle earnings dynamics. This decomposition aligns well with empirical evidence. Topel and Ward (1992) provide estimates for the contribution of employer switching to wage growth after labor market entry. They find that search for better-paying em-

[19] Heathcote et al. (2010) provide a detailed discussion of different estimation approaches. We use the estimation as a reduced-form description of earnings dynamics without requiring the process to be the true underlying process or estimates to be unbiased. See Daly et al. (2016) for further discussion.

ployers accounts for about one-third of wage growth within the first ten years in the labor market. Looking at the decomposition of mean earnings in Figure D.11a, we find, in line with their results, that climbing the wage ladder is an important driver of early career wage growth. Between ages 20 and 30, it accounts for roughly one-third of earnings growth, close to the Topel and Ward (1992) estimate. After age 30, wage growth from job search flattens out, in line with a slowdown in employer switching (Figure 6). Human capital investment accounts for almost the entire increase in earnings once most workers have found stable jobs.

Figure D.11: Decomposing earnings dynamics over the life cycle



(a) Mean (b) Variance

Notes: Panel (a) shows the profile of mean of log earnings and the contribution of the wage component to the growth of earnings over the life cycle. Panel (b) shows the contribution of human capital dispersion and wage dispersion to the overall earnings dispersion over the life cycle.

For the increase in the variance, we find a similar decomposition. At age 20, all workers start from the same level of human capital, so differences in entry wages account for all of the dispersion in earnings. Over time, workers climb the wage ladder leaving less well-paid jobs and accept better-paid jobs, which leads to wage compression and contributes negatively to the increase in life-cycle earnings inequality. Workers in well-paid jobs receive fewer opportunities to climb the wage ladder as many jobs offer lower wages, and therefore, these workers are more likely to stay with their current employer. As a result, initial wage differences decrease. This mechanism also highlights the general challenge when trying to account for wage dispersion relying on employer differences alone. As our decomposition shows, differences in human capital accumulation are the driver of rising earnings inequality over the life cycle. The covariance between human capital and wages (not shown) is small but positive and contributes little to earnings dispersion in the model. At age 40, the contribution of the covariance accounts for about 10 percent of the search component. This decomposition with a small contribution from search frictions (frictional wage dispersion) is consistent with results in Hornstein et al. (2011) and Hagedorn and Manovskii (2010) that point toward low levels of frictional wage dispersion in the data. Bayer and Kuhn (2018) decompose the

increase in life-cycle wage dispersion using German administrative data and also find a negligible contribution of employer differences to the life-cycle increase in wage dispersion.

An important observation to make is that the model jointly matches results on earnings dynamics and earnings inequality. Hornstein et al. (2011) show that models with on-the-job search and a homogeneous separation rate consistent with the observed average separation rate are able to match large wage inequality across otherwise identical workers. Wage differences in this model are highly transitory, however, as jobs are on average short-lived. The wage differences stem from a long, stretched wage ladder where workers start low on the ladder, and the long period of advancement on the job ladder spreads out wages, generating large inequality. The current model relies on a different underlying mechanism that increases the persistence of search outcomes, which allows us to be jointly consistent with cross-sectional inequality but also with the persistence of jobs and observed wage dynamics.

D.3.4 Wealth dynamics and the joint distribution of income and wealth

The consumption-saving block of the model follows the large literature on incomplete market models with idiosyncratic income risk (Aiyagari, 1994; Huggett, 1993). At their core, these models provide a mapping from earnings dynamics to wealth accumulation. Our model combines this consumption-saving block with a labor market block endogenizing earnings dynamics. This combination of building blocks suggests that a comparison of the joint distribution of income and wealth and the implied cross-sectional wealth dynamics induced by the model's earnings dynamics are particularly well suited to compare how our model performs in accounting for the data. Unlike panel data on consumption and income dynamics, the approach to compare the joint distribution of income and wealth has the additional advantage that these data are easily observed in datasets such as the SCF.

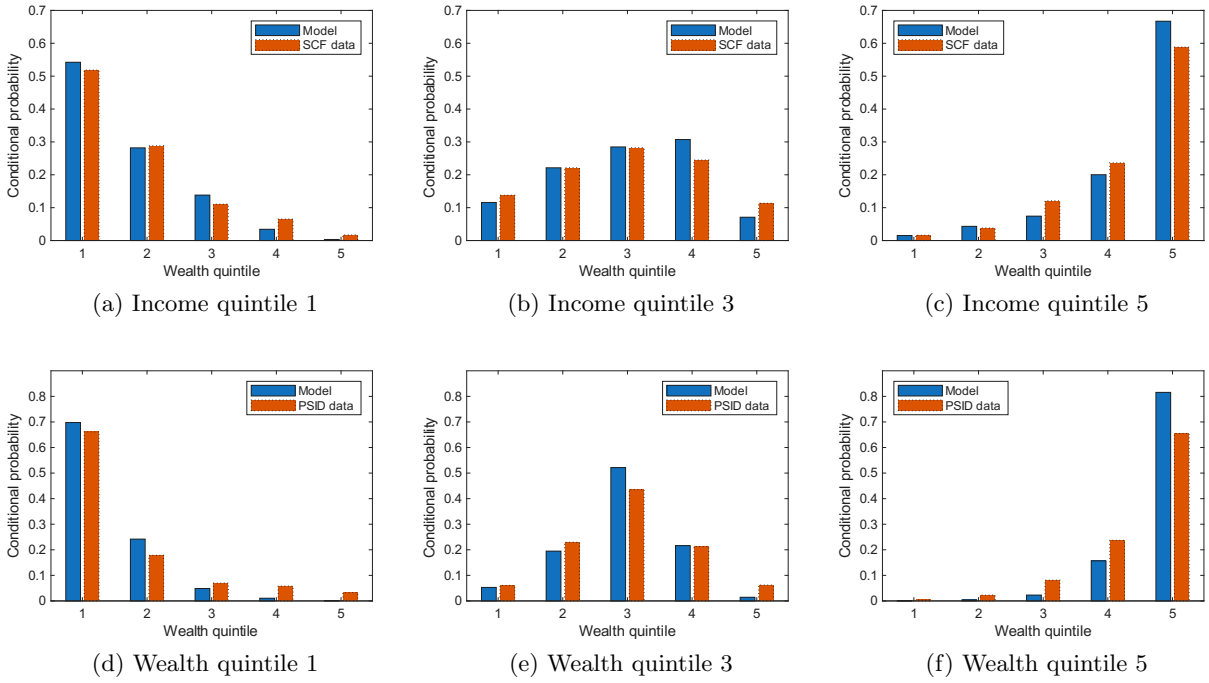
For this comparison, we split households into income and wealth quintiles. We consider the joint distribution by comparing how households are distributed across wealth quintiles conditional on their income quintile. For wealth dynamics, we compare how households move across wealth quintiles over time. Panels (a) through (c) in Figure D.12 show selected conditional distribution functions for wealth by income quintile from model and SCF data.[20] Panels (d) through (f) show the conditional distribution function for starting from a given wealth quintile over all wealth quintiles five periods in the future (i.e., we show rows of a five-step wealth transition matrix).[21] For data on the joint distribution, we rely on the SCF data. We follow previous research (Díaz-Giménez et al., 2011; Kuhn et al., forthcoming) and rely on PSID data from 1984 to 1999 to trace out individual-level wealth dynamics. The repeated cross sections of the SCF data prevent such an

[20] In Table D.1, we look at population shares across all income-wealth cells from model and SCF data and report all conditional PDFs. We report all households and households ages 40 to 50.

[21] We report all conditional probabilities in Table D.2.

analysis.[22]

Figure D.12: Joint distribution of income and wealth and wealth dynamics



(a) Income quintile 1 (b) Income quintile 3 (c) Income quintile 5

(d) Wealth quintile 1 (e) Wealth quintile 3 (f) Wealth quintile 5

Notes: Panels (a) through (c) show the joint distribution of income and wealth of households ages 40 to 50. We split households into quintiles along the income and wealth dimension and show the conditional cumulative distribution functions for wealth by income quintile from model and SCF data. Panels (d) through (f) show the five-year wealth transition probability from the PSID data. We split households into quintiles by wealth in period t and $t+5$ and compute the transition matrix. Household heads of ages 38-42 are used.

Looking at the joint distribution in Figures D.12a to D.12c, we find that the model aligns closely with the SCF data for households ages 40 to 50. We focus on a single age group to alleviate concerns regarding the age structure in the model relative to the age structure in the data. The fit of this untargeted dimension is very good. The most notable difference between the model and data is that too many households from the fifth income quintile are at the bottom of the wealth distribution. The model hence generates too many income-rich but wealth-poor households, but overall the distributions from the model and data align very closely. While income, being a flow, might change quickly, wealth as a stock moves much more slowly. The close fit of the joint distribution therefore suggests that the model-implied wealth dynamics compare favorably to their data counterpart.

[22] Pfeffer et al. (2016) compare SCF and PSID data, concluding that except for the very top of the wealth distribution, the two surveys provide consistent wealth distributions for the vast majority of households.

Looking at wealth dynamics in Figures D.12d to D.12f, we find that the model matches closely the observed wealth dynamics in PSID data. The PSID surveys wealth only every five years during this time period, so we focus on five-year transition probabilities. We observe a high persistence of households' position along the wealth distribution. More than 80 percent of households from the first wealth quintile remain within the two bottom wealth quintiles over a five-year horizon. Similarly at the top, only about 30 percent of households from the top quintile end up in a lower quintile five years later. Comparing the model and data, we find that the model produces slightly too little wealth mobility at the bottom and the top. The middle of the distribution is closely matched. It is important to note that imputation of wealth information in the PSID likely leads to overstating the estimates of wealth mobility in the PSID data.

Table D.1: Joint distribution of income and wealth, ages 40-50

	W0-20	W20-40	W40-60	W60-80	W80-100
			Model		
I0-20	0.11	0.06	0.03	0.01	0.00
I20-40	0.05	0.06	0.05	0.03	0.00
I40-60	0.02	0.04	0.06	0.06	0.01
I60-80	0.01	0.03	0.04	0.06	0.05
I80-100	0.00	0.01	0.01	0.04	0.13
			SCF data		
I0-20	0.10	0.06	0.02	0.01	0.00
I20-40	0.06	0.06	0.05	0.03	0.01
I40-60	0.03	0.04	0.06	0.05	0.02
I60-80	0.01	0.03	0.05	0.06	0.05
I80-100	0.00	0.01	0.02	0.05	0.12

Table D.2: Wealth transition matrix, age 40

	W0-20	W20-40	W40-60	W60-80	W80-100
			Model		
W0-20	0.69	0.24	0.05	0.01	0.00
W20-40	0.22	0.52	0.22	0.04	0.00
W40-60	0.05	0.19	0.52	0.22	0.01
W60-80	0.01	0.05	0.18	0.61	0.16
W80-100	0.00	0.00	0.02	0.16	0.82
			PSID data		
W0-20	0.66	0.18	0.07	0.06	0.03
W20-40	0.26	0.45	0.20	0.06	0.03
W40-60	0.06	0.23	0.44	0.21	0.06
W60-80	0.02	0.12	0.21	0.44	0.22
W80-100	0.00	0.02	0.08	0.24	0.65

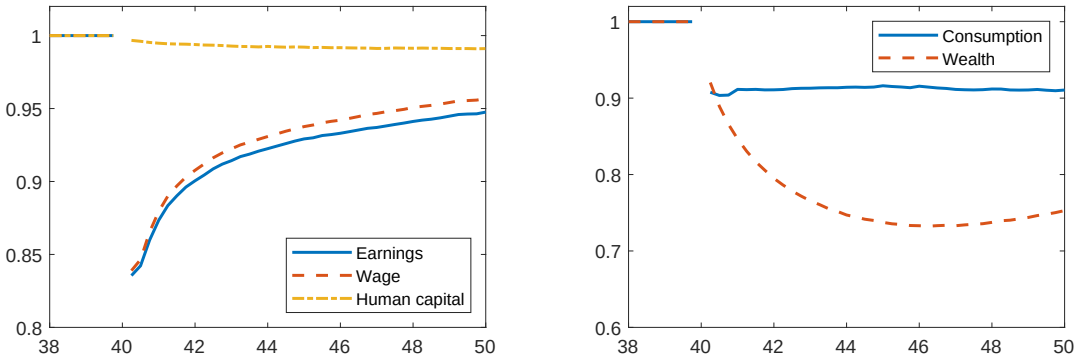
D.3.5 Consequences of Job Displacement

In this section, we provide additional results on the consequences of job displacement that complement the analysis on the macroeconomic consequences of employment stability heterogeneity in Section 5.2. To study the consequences of job displacement, we adapt the approach from the empirical literature on job displacement (Jacobson et al., 1993) and compare identical workers at age 40 where one worker is losing the job while the other worker remains employed and only faces the probability of future job loss. Specifically, we compare a cross section of 40-year-old employed workers to the same group of workers who have been sent (exogenously) into non-employment at age 40.[23] Empirical studies document that such job displacements lead to large and persistent earnings losses for workers (Jacobson et al., 1993; Couch and Placzek, 2010; Davis and von Wachter, 2011), and heterogeneity in employment stability has been identified as a key ingredient in accounting for large and persistent earnings losses in structural models (Jung and Kuhn, 2018; Jarosch, 2023). Models without heterogeneity in employment stability struggle to account for the persistence in earnings losses (Low et al., 2010).

Figure D.13a shows that our baseline model with heterogeneity in employment stability implies large and persistent earnings losses for the average 40-year-old worker. The initial earnings drop in displaced workers amounts to around 15 percent. Over the subsequent five years, displaced workers are able to cut initial earnings losses by half, but there is little further catch-up. The figure

[23]This approach differs from the empirical approach that conditions on pre-displacement tenure and post-displacement employment stability. In the model, we exploit the fact that we can directly implement a displacement event without having to deal with selection effects that are the key concern in the empirical implementation (see Jacobson et al. (1993)).

Figure D.13: Cost of displacement



(a) Earnings and Wages

(b) Consumption and Wealth

Notes: This figure shows the evolution of earnings, consumption, and wealth of workers who become unemployed at age 40 relative to a control group of workers who remain employed. Prior to displacement, both groups are identical.

also shows the evolution of the components of earnings (wage and human capital) to uncover the underlying mechanism of the persistent earnings loss. All of the initial loss in earnings comes from the fact that upon job loss, workers are unlikely to immediately find a well-paying job through off-the-job search. Most of the job offers that workers receive come with low wages and high separation rates. On-the-job search allows workers to catch up by climbing the wage ladder toward better-paying jobs; however, the speed of convergence reduces substantially after the first five years. Looking at the evolution of human capital, we find that job loss has a persistent negative effect on human capital accumulation that builds up dynamically. Two reasons explain the observed divergence. First, workers cannot accumulate human capital while being non-employed directly after the job loss. Second, new jobs are on average less stable when workers start climbing the wage ladder, so that workers will on average spend more time in non-employment limiting their human capital opportunities in the future. Still, we find that, in line with the results in Stevens (1997) and Jung and Kuhn (2018), lower wages account for the largest part of long-run earnings losses of the average worker.

Figure D.13b turns to the consequences of job loss for consumption and wealth. Looking at consumption, we see a sharp (roughly 10 percent) drop in consumption directly on impact. After the one-time persistent shock, consumption dynamics show only a very slight upward trend. These consumption dynamics can again be rationalized by the permanent income hypothesis. Directly upon job loss, agents anticipate that they enter an employment trajectory with lower and more volatile income. Income after job loss will be persistently lower due to lower earnings and lower employment rates, and income will be more volatile because of lower employment stability. As a

consequence of lower permanent income and higher volatility, agents permanently reduce consumption and increase their precautionary savings to smooth consumption in the future. On impact, consumption drops less than income, and this difference shows up directly in wealth dynamics as wealth is used to smooth the transition to the new, lower permanent income level (Kuhn, 2013). Four years after job loss, wealth levels stabilize 25 percent below the level of nondisplaced workers and remain persistently lower in line with recent empirical results in Barnette (2020). Three reasons account for this lower average wealth after job loss. First, income is lower so that the wealth level adjusts, too. Second, the job loss has flattened the life-cycle income profile. Current income is now lower relative to income during the rest of the life cycle, which reduces the need for life-cycle savings that aim at reshuffling age-varying income over time. Third, the lower employment stability after the job loss sets agents on the kind of Sisyphus saving cycle: While cycling through unstable jobs, workers' ability to accumulate wealth is mitigated by the fact that consumption smoothing over repeated spells of non-employment reduces any accumulated savings. These consumption and labor market dynamics also provide intuition for the positive relationship between wealth and non-employment duration implied by the model and observed in the data (Figure 8). Workers who lose their jobs experience significant decreases in income. During the transition, wealth declines, earnings recover so that wealth-to-income ratios fall, and we get the negative relationship between non-employment duration and wealth-to-income ratios, as observed in Figure 8.

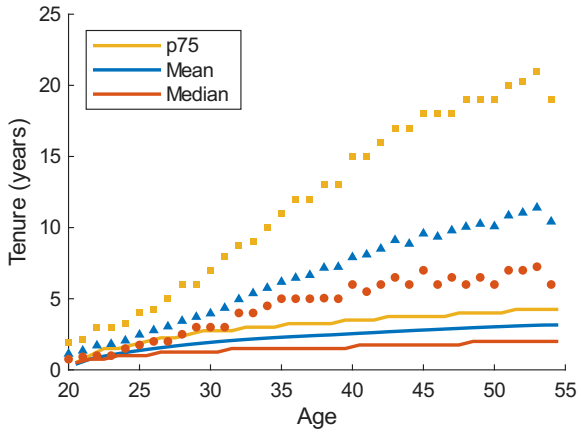
E Model extensions

E.1 A model without heterogeneity in employment stability

This section considers a version of the model without heterogeneity in employment stability. The structure of the model remains unchanged relative to the baseline specification, except that all workers now face the same exogenous, age-dependent job-separation rate. Specifically, the separation rate at each age is set equal to the average separation rate in the baseline model. As a result, the life-cycle profile of separation rates is identical across the two specifications. In this environment, job offers drawn from the job-offer distribution differ only along the wage dimension.

We find that this alternative specification performs substantially worse than the baseline model in matching the dispersion in employment histories of workers. When all workers face the same employment stability, the implied tenure distribution deviates markedly from the data. As shown in Figure E.1, tenure is considerably more compressed, and its life-cycle increase is much weaker than in both the baseline model and the empirical evidence.

Figure E.1: Tenure without heterogeneity in employment stability

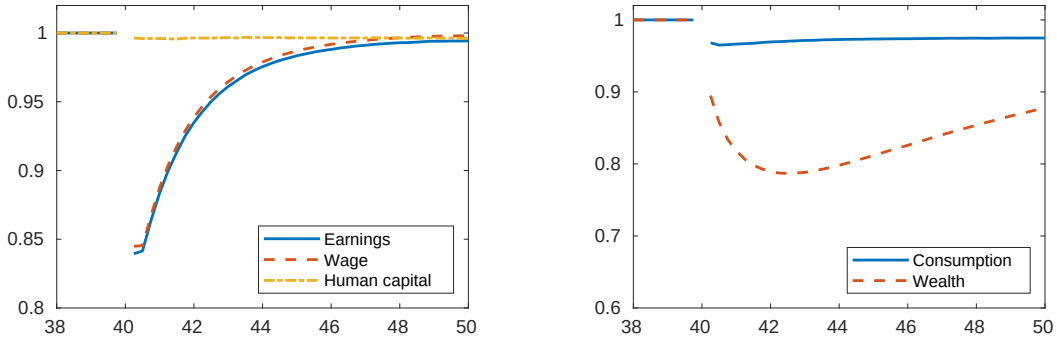


Notes: This figure compares the life-cycle evolution of the distribution of tenure for the model without heterogeneity in separation rates and for the data. The dots display data from the PSID, and the solid lines correspond to the model without heterogeneity in employment stability.

Furthermore, without any cross-sectional heterogeneity in the separation rate, the model also cannot replicate large and persistent earnings losses following displacement. In Figure E.2, we show the cost of displacement for the model without cross-sectional heterogeneity in the separation rate. Contrary to the results from the baseline model in Figure D.13, we find that earnings losses following displacement largely disappear ten years after displacement when jobs do not differ in terms of their job stability. Consistent with less persistent negative effects of non-employment on earnings, we also find that consumption declines substantially less compared to the baseline model.

Job loss in the economy without heterogeneity is largely inconsequential, in contrast to the model with heterogeneity, as discussed in the main part of the paper.

Figure E.2: Cost of displacement without heterogeneity in job stability



(a) Earnings

(b) Consumption and wealth

Notes: This figure shows the evolution of earnings, consumption, and wealth of workers who become unemployed at age 40 relative to the control group. Prior to displacement, both groups are identical.

E.2 A model with dynamic worker types

In the model with dynamic worker types, the separation rate is allowed to depend on a worker's human capital, capturing that employment stability varies over the life cycle and across individuals. Specifically, we assume

$$\lambda(h) = \lambda_l \cdot \max\{(2.5 - 1.5h + 0.1h^2), 0.5\},$$

where λ_l is drawn from the baseline distribution $\{\lambda_l\}_{l=1}^{L}$ and h denotes individual human capital. This modification affects the employed worker's value function (3), which can now be written as

$$V_e(a, w, \lambda, h, j) = \lambda(h)V_n^P(a, w, h, j) + (1 - \lambda(h))V_e^I(a, w, \lambda, h, j).$$

so that the probability of transitioning into non-employment is endogenous to the worker's human capital. All other value functions in Eqs. (3)–(7) remain unchanged.