

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

Moritz Kuhn[†]

Leanne Nam[‡]

Gašper Ploj[§]

June 19, 2026

Abstract

Labor markets feature large heterogeneity in employment stability: some careers provide lifetime employment, while others involve frequent transitions in and out of work. While this heterogeneity shapes earnings dynamics and labor market risk, its implications for household saving behavior remain poorly understood. We document two new empirical facts. More stable careers (i) exhibit steeper life-cycle earnings growth and (ii) accumulate significantly more wealth per dollar of income, even within narrowly defined worker groups.

To interpret these facts, we develop a life-cycle search-and-saving model with heterogeneous employment stability, job-to-job mobility, endogenous human capital accumulation, and incomplete markets. The model matches both life-cycle earnings and wealth dynamics across careers. Our central finding is that heterogeneity in employment stability reshapes the nature of saving. Stable careers generate sustained earnings growth with a focus on life-cycle saving, while unstable careers are characterized by precautionary, buffer-stock behavior that limits long-run wealth accumulation. Quantitatively, differential earnings growth accounts for about 60% of the wealth gap across careers.

These microeconomic differences have important macroeconomic implications. We demonstrate that heterogeneity in employment stability amplifies wealth inequality and substantially increases the macroeconomic consumption response to unemployment shocks.

Keywords: Employment risk, employment stability, consumption-saving behavior

JEL: J64, E21, E24

*We thank Christian Bayer, Richard Blundell, Thomas Hintermaier, Philip Jung, Per Krusell, Iouri 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. 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

¹Worker characteristics (e.g., education) also contribute to heterogeneity in employment stability. Our focus is on within-worker-group heterogeneity, which we document to be quantitatively as large as the overall dispersion

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 several untargeted empirical facts. It replicates patterns in consumption inequality documented by Aguiar and Hurst (2013) (see Appendix D.3.1), earnings losses following job displacement as in Jacobson et al. (1993) (Section 5.2), and key features of earnings dynamics and the distribution of annual earnings growth (Appendix D.3.3) documented by Topel and Ward (1992); Blundell et al. (2008); Guvenen et al. (2021a). The model’s empirical success demonstrates that it provides an informative quantitative framework to study the source of the empirical facts.

The search-and-saving model explains the tight empirical relationship between employment stability, earnings growth, and wealth accumulation as follows: 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 the associated large life-cycle variation in income. By contrast, along an unstable career path, careers do not take off: repeated job loss and employment interruptions limit human capital accumulation and dampen earnings growth. With little life-cycle income variation to smooth, the life-cycle saving motive is weak. Instead, savings are driven by a precautionary motive, workers build up assets while employed and draw them down during non-employment spells, giving rise to buffer-stock dynamics.

Hence, the heterogeneity of career paths gives rise to distinct saving regimes. Workers in stable careers behave like *Modigliani savers*, accumulating assets in response to stable life-cycle earnings growth. In contrast, workers in unstable careers are caught in a *Sisyphus cycle*: they repeatedly accumulate precautionary savings and decumulate them during non-employment. If this cycle persists over the life cycle, long-run wealth accumulation remains limited despite ongoing saving efforts. Taken together, heterogeneity in employment stability thereby jointly determines earnings risk and earnings growth and shapes the household saving problem. It shifts the dominant saving motive from life-cycle saving in stable careers to precautionary saving in unstable careers, leading to systematic differences in the nature of wealth accumulation.

We demonstrate that heterogeneity in employment stability also has important macroeconomic consequences. As it shapes savings behavior, it will, in turn, shape the distribution of wealth, and amplify 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 optimal design of insurance policies.

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.

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 the unstable career paths observed in the data. 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 and is 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 and 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 28% 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 part of the literature that focuses on macroeconomic dynamics, as in Krusell et al. (2010), with microeconomic behavior, as in Lise (2012).

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

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. His amplification mechanism results from a shifting portfolio choice. 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.

Hall and Kudlyak (2022) document substantial heterogeneity in worker flows between employment, job search, and non-participation, highlighting that labor market dynamics are shaped by large differences in transition rates across individuals. Their evidence shows that a relatively small share of workers accounts for a disproportionate amount of labor market turnover, pointing to persistent heterogeneity in employment stability. Relatedly, Ahn et al. (2023) emphasizes the role of labor market duality, distinguishing between workers in stable employment relationships and those in more precarious jobs characterized by frequent transitions across labor market states. Our framework is closely related to these findings, as it incorporates heterogeneity in employment stability across workers, which generates differential exposure to job loss and drives dispersion in employment trajectories over the life cycle.

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

³Heterogeneity 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.

heterogeneity is key to account for life-cycle earnings dynamics. Additional evidence for heterogeneity in employment 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 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 a sophisticated earnings process with nonlinear persistence and conditional skewness, showing that the consequences of earnings shocks depend on earnings histories. They also find that 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 and microfound 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 in a first step 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 second part of the paper. 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.3, we provide corroborating evidence for both steps based on independent high-quality data from the *Survey of Consumer Finances* (SCF).

The PSID is one of the longest-running household panels 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 the effect of this change on our measurement and demonstrates the robustness of our results.

2.1 Heterogeneity in employment stability

To quantify heterogeneity in employment stability, Figure 1 presents two complementary representations of the distribution of non-employment duration: the histogram of accumulated life-cycle non-employment across careers and the corresponding Lorenz curve. Panel 1a shows the distribution of non-employment duration at age 50, accumulated between ages 25 and 50. Strikingly, more than two out of ten workers experience no spell of non-employment over their prime-age working life. Around half of all workers accumulate between 1 and 10 quarters of non-employment, whereas

the remaining 30% accumulate more than 10 quarters. A non-employment duration of 10 quarters corresponds to a lifetime non-employment rate of 8%.

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



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.


Panel 1b shows the corresponding Lorenz curves as a standard inequality measure, summarizing how unequally non-employment duration is distributed across careers. Looking first at all workers (green line with squares), we find that around 30% of workers account for more than 70% of total non-employment duration. At the same time, we find again about 20% of careers with zero non-employment share, highlighting that labor market mobility is concentrated in a relatively small group of careers. To contrast this inequality with a setting of homogeneous employment stability, we plot as a 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 lies much closer to the identity line. In this case, the most unstable 30% of careers account for only 40% of accumulated non-employment duration, implying that when only realization risk is present, non-employment is much more evenly distributed across careers.

Besides heterogeneity in job stability, worker heterogeneity could also be an important driver of differences in non-employment duration. The arguably most important dimension of worker heterogeneity is educational attainment. College-educated workers accumulate on average about 5 quarters of non-employment over the life cycle, whereas non-college workers experience roughly 50% more. Figure 1b shows, as a dotted red line with circles, the Lorenz curve of non-employment

duration for college graduates and, as a dashed blue line with triangles, the corresponding curve for non-college workers. Strikingly, despite substantial differences in average non-employment duration, within-group heterogeneity in non-employment duration is nearly identical across the two groups. Hence, even among workers with very different average employment stability, the dispersion of employment stability within groups is remarkably similar. In our structural model, we therefore focus on employment stability arising from heterogeneity in job stability, which is consistent with the large within-worker-group dispersion observed in the data, and abstract from between-worker heterogeneity.

In Appendix B, we provide additional evidence supporting the view that heterogeneity in non-employment largely reflects heterogeneity in job stability. First, we show that following displacement from stable jobs, employment histories exhibit persistently higher subsequent non-employment, consistent with workers transitioning into less stable jobs. Second, while non-employment duration is strongly persistent over the life cycle on average, this persistence largely disappears within the subgroup of workers experiencing non-employment during prime-age working life, which is inconsistent with heterogeneity coming from fixed worker types. Third, we document substantial variation in job destruction rates across employers by firm age, with differences of up to a factor of three. Finally, we show that risk attitudes, as an additional dimension of worker heterogeneity, are not meaningfully related to non-employment duration in the PSID. Taken together, these findings point to job heterogeneity as a key driver of heterogeneity in non-employment duration.⁴

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.

Another approach to document heterogeneity in employment stability is to look at the tenure

⁴In Section 4.3, we discuss results from model extensions that incorporate empirically disciplined worker heterogeneity and show that our key model results on differences in earnings growth and wealth accumulation are robust.

distribution (Hall, 1982; Jung and Kuhn, 2018).⁵ Figure 2 shows life-cycle profiles for tenure in the PSID data.⁶ We find that the tenure profile is positively correlated with age. Looking at the mean, 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% 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.⁷

2.2 Employment stability and earnings dynamics

What are the implications of the observed heterogeneity in employment stability for earnings and wealth accumulation? To address this question, we exploit the panel structure of the PSID, which jointly tracks labor market histories, earnings, and wealth over time. We classify careers into stable and unstable groups based on cumulative non-employment duration between ages 25 and 50. Specifically, careers in the top 25% of the distribution of cumulative non-employment at age 50 are classified as unstable, while the remaining careers are classified as stable. We use this grouping to study the relationship between earnings growth and stable and unstable careers.

As this grouping conditions on the ex-post realization of employment histories, we demonstrate in Appendix C.1 that our findings are robust to an ex-ante classification based only on employment histories observed up to age 30. This alternative sorting yields very similar group assignments: only 10% of careers classified as stable at age 30 are reclassified as unstable by age 50.⁸

Figure 3 shows the life-cycle profiles of log labor income for workers with stable and unstable employment, where we remove initial level differences to focus on differences in earnings growth. Strikingly, we observe a large divergence in earnings growth between stable and unstable careers over the life cycle. By age 50, average earnings in the stable group have grown by about 20 log points more than in the unstable group.⁹ This finding speaks against the common practice of treating average life-cycle earnings growth and earnings risk as independent. Instead, our results

⁵Appendix Figure C.2 shows the strong correlation between the different employment stability measures in the PSID data.


⁶Employer tenure is widely observed in cross sectional datasets and therefore provides an attractive variable to explore the relationship between wealth accumulation and employment stability. We use tenure as an alternative measure of employment stability when studying the relationship with wealth accumulation in the SCF data (Appendix C.3).

⁷Note that differently to non-employment duration tenure also ends in case of a job-to-job transition.

⁸In Appendix A.1, we also address a potential measurement concern arising from the biennial sampling of the PSID. We find that misclassification of careers is quantitatively negligible, with a rate below 2%.

⁹As we will study household wealth accumulation, one potential concern is the difference between individual and household income. We demonstrate in Appendix A.2 that accounting for household-level income, if anything, amplifies the observed individual earnings differences.

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.


suggest that more frequent non-employment not only leads to income losses during non-employment spells, but also to persistently lower earnings growth over the life cycle.

One potential explanation is that repeated job loss generates a sequence of large and persistent negative shocks, leading to lower earnings over time. However, such an explanation faces the challenge that the existing empirical evidence on large and persistent earnings losses from displacement conditions on displacement from stable employment (Jacobson et al., 1993; Davis and von Wachter, 2011). Put differently, workers in stable employment experience large and persistent earnings losses upon job loss, but tend to transition afterward into unstable employment, which is typically associated with less stable jobs (Appendix B). We return to this fact below when discussing the macroeconomic implications of heterogeneity in employment stability, where we demonstrate that earnings losses from unstable employment are only transitory.

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 the wealth accumulation of workers with different employment trajectories over time. Figure 4 first examines the relationship between accumulated non-employment duration, as documented in Figure 1b, and wealth-to-income ratios. We focus on wealth-to-income ratios to abstract from differences in income levels and isolate variation in wealth accumulation, given the earnings differences documented in Figure 3.

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



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.

Figure 4a shows a strong negative relationship between wealth-to-income ratios and non-employment duration at age 50. Qualitatively, the negative slope implies that workers with more stable employment histories hold more wealth per dollar of income. In other words, workers with more stable employment histories accumulate more wealth. Quantitatively, the magnitude of this relationship is economically meaningful. The slope implies that four additional years of non-employment are, on average, associated with a reduction in wealth equivalent to roughly one year of income.¹⁰

One concern is that the relationship between non-employment and wealth accumulation may reflect fixed worker characteristics, most prominently education (Bartscher et al., 2020). To address this concern, Table 1 reports several regression specifications of wealth-to-income ratios on non-employment duration. Column (1) presents the baseline specification corresponding to Figure 4a. The estimate implies that one year of non-employment, corresponding to four quarters, is associated with a 0.3 lower wealth-to-income ratio.

In column (2), we add education as an additional control and allow the relationship between non-employment duration and wealth-to-income ratios to differ by educational attainment. Higher educational attainment, captured by a college dummy, is associated with higher wealth-to-income

¹⁰In line with this finding, Iacono and Ranaldi (2020) report a negative correlation between wealth and unemployment in Norwegian data.

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

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

Notes: The dependent variable is the wealth-to-income ratio. The sample consists of workers at age 50. Column (1) shows the regression results of the 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. 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

ratios. At the same time, the coefficient on non-employment duration remains negative and statistically significant. The interaction between non-employment duration and education, however, is not statistically significant, indicating that the slope does not differ significantly between college and non-college workers. Hence, the negative relationship between non-employment duration and wealth accumulation is present within worker groups and is not solely driven by differences between college and non-college workers.


Columns (3) and (4) extend the specifications from columns (1) and (2) by adding controls for sex, race, region, marital status, family size, and the industry with the longest employment duration of the household head. With these additional controls, the estimated relationship between non-employment duration and wealth-to-income ratios declines slightly but remains highly statistically significant. The interaction between non-employment duration and education remains statistically insignificant, indicating again that the slope does not differ significantly between worker groups.

Figure 4b turns from worker characteristics to the role of job heterogeneity. Having shown that observable worker characteristics do not account for the relationship between non-employment duration and wealth accumulation, we next examine whether the same relationship appears across occupations. Focusing on occupations allows us to analyze one observable dimension of job characteristics that may shape both household wealth and non-employment outcomes. To do so, we estimate regressions of the household wealth-to-income ratio and cumulative non-employment du-

ration on a full set of occupation fixed effects, controlling for a rich set of individual characteristics, including education, gender, race, region, marital status, and family composition. By conditioning on these observable characteristics, the estimated occupation effects capture systematic differences across jobs rather than compositional differences across workers.

Figure 4b plots, for each occupation, the estimated 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 associated with substantially lower wealth-to-income ratios, whereas occupations with more stable employment are associated with higher wealth-to-income ratios. Notably, the absolute slope of this relationship is 0.14, which is larger than in the raw data and in the specifications controlling for worker observables in Table 1. This strong negative relationship across occupations provides further support for the role of job heterogeneity in shaping employment stability and its connection to wealth accumulation.

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.

As a final step, we examine how differences in wealth accumulation unfold over the life cycle. Figure 5 shows the life-cycle profiles of wealth-to-income ratios for stable and unstable careers, paralleling the earnings profiles in Figure 3. As with earnings, we observe a persistent and widening gap in wealth-to-income ratios over the life cycle. By age 50, workers with stable careers have accumulated wealth equivalent to roughly one and a half years of income more than workers with unstable careers, corresponding to about 75% higher wealth per dollar of income. These differences in wealth accumulation build up gradually over the life cycle, beginning at age 25 and becoming particularly pronounced after age 30, when the life-cycle savings motive becomes most relevant

Table 2: Wealth-to-income ratio and non-employment

Wealth-to-income ratio	(1)	(2)
Non-employment (in quarters)	−0.05*** (0.01)	−0.05*** (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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

(Gourinchas and Parker, 2002; Cagetti, 2003).

The panel dimension of the PSID allows us to further complement these life-cycle profiles with regression evidence that controls for worker fixed effects as a source of differences in wealth accumulation. Table 2 reports results from two regression specifications. In both cases, we regress wealth-to-income ratios on non-employment duration, controlling for age, industry, and worker fixed effects, thereby accounting for time-invariant worker differences. One important worker characteristic potentially related to both employment stability and wealth accumulation is risk aversion. In the second specification, we therefore add an interaction between age and a measure of risk aversion, allowing wealth accumulation patterns to differ across workers with different degrees of risk aversion. Appendix Figure B.4 uses the same measure and demonstrates that risk aversion is not systematically related to differences in accumulated non-employment at age 50.

As in the cross-sectional regressions above, we find a highly significant negative relationship between non-employment and wealth accumulation. The magnitude is very similar to the estimated coefficients at age 50 in Table 1. The final specification, which allows life-cycle wealth profiles to differ by risk aversion, further suggests that sorting on risk preferences is not a major driver of the relationship between employment stability and wealth accumulation.

Appendix C.3 provides additional evidence on the relationship between employment stability and wealth accumulation using high-quality wealth data from the SCF. Although the SCF is cross-sectional, it contains information on workers' employment histories, including employer tenure and the number of employers over the working life.¹¹ Using these variables as alternative measures of employment stability, we find that wealth-to-income ratios are strongly increasing in employer tenure and decreasing in the number of employers. These findings from independent data closely

¹¹Appendix Figure C.6 shows that, in the PSID data, employer tenure and accumulated non-employment duration are strongly correlated.

corroborate the PSID results, indicating that the relationship between employment stability and wealth accumulation is robust across datasets and across alternative measures of employment stability.

In summary, our empirical results document that employment stability is highly heterogeneous across careers and that dispersion within worker groups is as large as in the overall population, despite differences in average employment stability across worker groups. A novel empirical finding is that this heterogeneity is systematically related to both earnings growth and wealth accumulation, generating economically meaningful differences across careers. Our regression analysis further shows that these differences persist within worker types, including for both college-educated and non-college workers. These findings raise the key question of how employment stability jointly shapes earnings growth and wealth accumulation. They also point to a tension in a standard precautionary-saving view: workers with unstable careers face greater non-employment risk, yet accumulate less wealth per dollar of income over the life cycle. In the next section, we develop a novel search-and-saving model with heterogeneity in employment stability to examine how these differences shape earnings growth, savings motives, and wealth accumulation.

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 featuring heterogeneity in employment stability and endogenous human capital investment. In line with the empirical evidence, we focus on within-group heterogeneity in analyzing the relationship between employment stability, earnings growth, and wealth accumulation. We interpret this heterogeneity as arising from differences in job stability, consistent with a large empirical and structural literature (Hall, 1982; Pinheiro and Visschers, 2015b; Jung and Kuhn, 2018; Jarosch, 2023; Larkin, 2019). In Appendix E.1, we describe a version of this model without heterogeneity in employment stability, which serves as a benchmark throughout the analysis.

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, which lasts for T^W periods. At the end of the working phase, workers move to the transition phase, which is of stochastic length with expected duration T^T . In the end, workers leave the transition phase for the retirement phase, which lasts for T^R periods. In each period before retirement, 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, and the wage w captures the rental rate of human capital on the current job. We discretize wages and separation rates on 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 . Differences in λ capture 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

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

where r denotes the risk-free rate on the economy's single risk-free asset and y denotes current-period 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, $y(w_j, h_j, n) = b w_j h_j$, where b denotes the replacement rate, w_j is the wage on the last job, and h_j corresponds to the human capital stock when last employed. 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 $\max\{w_{k-1}, w_1\}$. During retirement, agents receive social security benefits proportional to their stock of human capital prior to retirement times the economy-wide average wage, $y(w_j, h_j) = s \bar{w}_j h_j$, 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 job-specific separation 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. We assume that human capital remains constant during non-employment and that only employed workers can invest in their human capital. At the production stage, employed agents receive earnings equal to the job's wage rate times the worker's stock of human capital, while 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, consisting of a wage rate w and a separation probability λ , are drawn for both employed and non-employed workers from the same joint distribution $f(w, \lambda)$ (Pinheiro and Visschers, 2015b; Jarosch, 2023; Larkin, 2019). An agent who receives a job offer decides whether to accept or reject it. If the agent accepts the offer, she will be employed in the new job at the beginning of the next period. If the agent rejects the offer, she remains non-employed (or employed in her current job), and there is no recall of previous job offers.

At the investment stage, only employed workers have the opportunity to invest in their human capital. Effort provision for human capital accumulation is a choice $t \in [0, 1]$ (training). Disutility from effort enters utility additively separable as a quadratic cost κt^2 . Non-employed agents do not have the opportunity to accumulate human capital. If employed agents do not exert effort, their human capital remains constant at level h until the next period.¹² 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 belong to 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 who exerts effort t accumulates human capital with probability $p_H(t, j)$, reaching level h^+ . We allow for age dependence of $p_H(t, j)$. The law of motion for human capital when 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 extends Jung and Kuhn (2018) by 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 . We relegate details to Appendix D.1. Upon reaching the retirement phase, workers leave the labor market and receive 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 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

¹²Although 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.

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,

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

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 stochastic human capital accumulation occurs at the beginning of the production stage:

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

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

$$V_e^P(a, w, \lambda, h, j) = \max_{\{c, a' \geq 0\}} u_j(c) + \beta \left(\pi_e V_e^S(a', w, \lambda, h, j) + (1 - \pi_e) V_e(a', w, \lambda, h, j + 1) \right) \quad (4)$$

s.t. $c = (1 + r)a + y(w, h, e) - a'$,

where V_e^P denotes the employed agent's value function at the production stage, V_e^S denotes the value function at the search stage, and V_e denotes the value function at the beginning of the next period. The time discount factor is denoted by β , and the probability of receiving a job offer is π_e . Job offers are drawn from the distribution $f(w, \lambda)$, so that the value function of an employed worker at the search stage is given by

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

where N_w is the number of wage realizations and N_λ is the number of separation rate realizations in the support of the offer distribution. The bundle $\{w_s, \lambda_k\}$ denotes an outside job offer. The value function at the search stage represents the expected continuation value across job offers, taking into account the agent's optimal acceptance decision.

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

$$V_n^P(a, w, h, j) = \max_{\{c, a' \geq 0\}} u_j(c) + \beta \left(\pi_n V_n^S(a', w, h, j) + (1 - \pi_n) V_n(a', w^-, h, j + 1) \right) \quad (6)$$

s.t. $c = (1 + r)a + y(w, h, n) - a'$,

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

The value function of a non-employed worker at the search stage captures, as in the case of an employed agent, the expected continuation value across job offers, taking into account the agent's optimal acceptance decision. It is given by

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

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

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

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 one model period to correspond to one quarter and assume that the utility function over 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.¹³ Human capital takes on discrete values $h_{i,t} \in \{h_1, \dots, h_{N_h}\}$. The human capital grid comprises N_h states that we set equidistant in log space between $h_1 = 1$ and $h_{N_h} = 6.5$. We allow the probability $p_H(t)$ to be age dependent and potentially declining with age at rate ρ ,

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

where t denotes effort provision and $\bar{p}_H$ the baseline level. As discussed below, this specification, together with the wage process, allows the model to match 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 the replacement rate in non-employment to 0.4, as in Shimer (2005), and the replacement rate in retirement to 0.45, in line with OECD estimates for the mean net pension replacement rate in the United States (OECD, 2015). We set the length of the working phase T^W to 35 years, the expected duration of the transition phase T^T to 10 years, and the retirement phase T^R to 20 years. Labor market entry occurs at age 20. We focus our discussion of model predictions on working life, i.e., between ages 20 and 55.

We denote the support of wages by $[\underline{w}, \bar{w}]$ and that of job stability by $[1 - \bar{\lambda}, 1 - \underline{\lambda}]$. We set $N_w = 5$, $\underline{w} = 1$, and $\bar{w} = 1.85$, in line with the empirical support of mean log earnings, and use grid points that are equidistant in logs. For job stability $1 - \lambda$, we set $N_\lambda = 10$, $\bar{\lambda} = 0.35$, and $\underline{\lambda} = 0.006$, the latter corresponding to 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.¹⁴ While the grid points are set exogenously, we calibrate the distribution over these grid points. For the job-offer distribution $f(w, \lambda)$, we assume that the marginal distributions of wages and job stability ($1 - \lambda$) follow truncated exponential distributions. 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 given by $f(w^*) = (1 - \exp(-\psi_w))^{-1} \psi_w \exp(-\psi_w w^*)$, where ψ_w determines the shape of the distribution. The density of standardized job stability $1 - \lambda^*$ is defined analogously with shape parameter ψ_λ . To obtain the joint distribution, we parameterize the dependence between the marginals using a copula C_θ , where θ governs the correlation between w^* and $1 - \lambda^*$.

We calibrate parameters within the model using a simulated method of moments that minimizes the difference between model moments and empirical moments (see Appendix D.2 for details). For the empirical moments, we use the life-cycle profiles of mean (log) earnings, labor market

¹³We employ the OECD equivalence scale, assigning 1 to the first adult, 0.5 to each additional adult, and 0.3 to each child.

¹⁴Specifically, we set the second grid point at $\lambda_2 = 0.05$ and the remaining grid points according to the nonlinear rule $\lambda_j = \underline{\lambda} + \left(\frac{j-1}{N_\lambda-1}\right)^{0.6} \times (\bar{\lambda} - \underline{\lambda})$.

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.¹⁵ In Appendix D.2, we describe the implementation details and also offer an intuitive discussion of how the empirical life-cycle profiles help identify the model’s 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 3 presents the model parameters together with their estimated values.

Table 3: Estimated parameters

Parameter	Value	Description
β	0.994	Quarterly discount factor
κ	0.651	Utility cost of effort
π_e	0.430	Probability of a job offer when employed
π_n	0.850	Probability of a job offer when nonemployed
ψ_w	0.538	Marginal distribution of w^*
ψ_λ	0.504	Marginal distribution of $1 - \lambda^*$
θ	0.474	Joint distribution of w^* and $1 - \lambda^*$
$\bar{p}_H$	0.071	Human capital upgrading probability
ρ	0.984	Decay rate of human capital upgrading probability

The calibrated value of the quarterly discount factor β corresponds to an annualized value of 0.976, which lies well within the range of conventional values in the macroeconomic literature. The utility cost parameter κ implies moderate costs of effort provision: measured as lifetime consumption-equivalent variation, these costs amount to 0.33% during the first ten years of working life and decline to less than one-tenth of a percent during the last ten years. Turning to labor market parameters, the model implies that job-offer arrival rates in non-employment π_n are roughly twice as high as those in employment π_e , consistent with high quarterly job-finding rates. At the same time, employed workers continue to receive job offers at a substantial rate. This difference in contact rates is both 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 different wage and job-stability levels in the offer distribution. The implied distribution is economically plausible: roughly one-third of job offers correspond to the lowest wage, whereas around one in twelve offers correspond to the highest wage. Similarly, less

¹⁵We 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.


than one in twenty jobs are highly stable lifetime jobs, while nearly one in seven offers correspond to the least stable job type. The copula parameter θ implies a positive correlation between wages and job stability. This correlation is economically meaningful. For example, if wages and job stability were independent, the probability of a low-wage, least-stable job would be 3.1%; with the estimated correlation, this probability increases to 6.8%. Conversely, the probability of a high-wage, highly stable job is only 0.9%, indicating that such jobs are hard to find. This positive association between wages and job stability is consistent with empirical evidence showing that higher-wage jobs tend to be more stable (Jung and Kuhn, 2018). Figure D.7a illustrates the estimated joint distribution of wages and separation rates, as well as the marginal distributions of separation rates across wage levels. The distribution is asymmetric, with most probability mass concentrated on low-wage, unstable jobs. Figure D.7b further shows that the distribution of separation rates for low-wage jobs first-order stochastically dominates that for high-wage jobs. Finally, the human capital parameter $\bar{p}_H$ implies that, for a labor market entrant, maximum effort provision ($t = 1$) yields a 25% probability of career progression in the first year. The decay parameter ρ implies that after 10 years in the labor market, the same effort yields an 15% probability of progression. Taken together, these parameter values are economically reasonable and consistent with existing empirical evidence. We now show that they also allow the model to match both targeted and untargeted features of the data.

3.2 Model fit

Figure 6 shows the empirical life-cycle profiles for separation, job-to-job, and job-finding rates together with their targeted model counterparts. Starting with separation rates in Figure 6a, the model closely matches their evolution over the life cycle: both model and data exhibit a sharp decline up to age 30 and a continued gradual decline between ages 30 and 50. The model accounts for this pattern through mobility along the job ladder, as workers transition over time into more stable and better-paying jobs. Heterogeneity in job stability is a key model feature in matching this life-cycle decline. As jobs differ in separation rates, workers can leave unstable jobs and move to more stable jobs over the life cycle. By contrast, a model with fixed worker heterogeneity in separation rates does not generate such a declining life-cycle profile, as, by construction, fixed worker heterogeneity does not change with age. Hence, the observed decline in separation rates over the life cycle provides further evidence of the importance of heterogeneity in employment stability at the job level. In Section 4.3, we study a model extension in which human capital affects separation rates. We discipline its contribution to the decline in separation rates over the life cycle following Jung and Kuhn (2018).

Figure 6b shows that the model also matches the life-cycle profile of job-to-job transition rates well, with only a slightly steeper decline between ages 20 and 30 compared to the data. Job-finding

Figure 6: Transition rates and 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.

rates in Figure 6c are matched closely in both level and trend and display little life-cycle variation, in line with the data. Figure 6d presents the life-cycle profiles of the mean, median, and 75th percentile of the tenure distribution. The model closely replicates both the increase in tenure over the life cycle and the substantial dispersion in job stability across workers. Importantly, the model simultaneously matches high average transition rates early in life (Figures 6a and 6b) and high job stability for a large share of workers later in life (Figure 6d). This joint fit is a key success of the model and reflects its ability to generate substantial heterogeneity in employment stability across careers.

Appendix Figure D.9 provides additional validation using cross-sectional job stability measures from the SCF, showing the distributions of employer tenure and the number of employers over the life

cycle. The model matches both distributions closely, accounting for the large mass of short-duration jobs as well as the substantial share of long-term employment relationships with tenure exceeding 10, 20, and even 30 years.¹⁶ We also demonstrate that heterogeneity in employment stability is essential for achieving this fit. Appendix Figure E.1 contrasts the tenure profiles in Figure 6d with those from a model without heterogeneity in employment stability. In the case without heterogeneity, mean job tenure is bounded at around three years, and the 75th percentile at age 55 remains below five years, compared to roughly 20 years in the data. This stark difference highlights that models with representative-worker labor market dynamics are unable to generate the observed dispersion and persistence in employment relationships. Only heterogeneity in employment stability allows the model to match both the level and the life-cycle evolution of employment dynamics observed in the data.

Figure 7: Earnings and wealth



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 first at the life-cycle profile of mean log earnings in Figure 7a, the model closely matches the steep increase in earnings after labor market entry and the subsequent flattening after age 40. It replicates the large average increase of roughly 100 log points over the life cycle, although it exhibits slightly less concavity than in the data. Figure 7b presents the life-cycle profile of the wealth-to-income ratio as a measure of wealth accumulation. Again, the model closely matches the data. Wealth-to-income ratios in both model and data rise from slightly above zero at age 20 in the data (zero in the model) to

¹⁶For most of the paper, we abstain from cross-sectional comparisons, as they require taking a stand on the age distribution in the model. Whenever possible, we compare age-specific model moments to the data. When aggregation is required, we assume a uniform age distribution.

approximately 4 at age 55. Hence, the model is jointly matches average earnings growth and wealth accumulation over the life cycle.

In Appendix D.3.1, we show that the model also matches life-cycle inequality profiles, in particular the increase in the variance of log earnings and the rise in consumption inequality. Appendix D.3.3 provides a more detailed analysis of individual earnings dynamics by comparing model-simulated earnings data with established empirical descriptions of earnings dynamics. We first demonstrate that the model is consistent with standard estimates of the earnings process based on permanent-transitory decompositions, as in Meghir and Pistaferri (2004), Blundell et al. (2008), and Heathcote et al. (2010). We also corroborate the finding of Hubmer (2018) that life-cycle labor market models can generate earnings growth distributions consistent with the empirical evidence documented by Guvenen et al. (2021a). In addition, we decompose earnings growth over the life cycle and show that our decomposition aligns with the evidence in Topel and Ward (1992) on early-career wage growth. The joint consistency of the model with these facts supports the calibration of the human capital and wage processes as the key drivers of life-cycle earnings dynamics.

Finally, Appendix D.3.4 shows that the model’s consumption-saving and earnings dynamics generate a joint distribution of earnings and wealth that is consistent with the data and captures wealth mobility over the life cycle in the PSID. This joint fit provides further support for the economic mechanisms underlying the model, in particular, its endogenous earnings dynamics and wealth accumulation decisions.

4 Consequences of heterogeneity in employment stability


This section examines how heterogeneity in employment stability shapes earnings growth and wealth accumulation. We compare the model predictions for life-cycle earnings and wealth profiles across stable and unstable careers to their empirical counterparts documented in Section 2.2. Importantly, these differences are not targeted in the calibration and therefore provide an independent assessment of the model’s mechanism. After establishing that the model closely matches the observed heterogeneity across career paths, we use the structural framework to decompose the channels through which employment stability affects 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. Figure 4 from the empirical analysis showed a strong negative relationship between accumulated non-employment duration and wealth-to-income ratios. Figure 8 shows that the model reproduces this relationship closely 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 well within the empirical range reported in Table 1, from -0.07 to -0.03 .

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.

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 are classified into stable and unstable careers using the same procedure as in the empirical analysis. These profiles provide a joint assessment of the model's ability to capture how employment stability shapes both earnings dynamics and wealth accumulation over the life cycle.


Figure 9a shows that earnings profiles of stable and unstable careers begin to diverge shortly after labor market entry. The gap widens until mid-career and remains persistent thereafter. By age 50, workers with stable careers earn approximately 20 log points more than those in unstable careers, both in the model and in the data.¹⁷

Figure 9b shows the corresponding wealth-to-income ratios for stable and unstable careers. The model replicates the gradual divergence observed in the data: wealth accumulation paths begin to separate around age 30 and continue 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 closely accounts for average employment dynamics, earnings growth, and wealth accumulation. More importantly, it also captures heterogeneity in labor market histories and reproduces the empirical relationship between employment stability, earnings growth, and wealth accumulation across stable and unstable careers. These results provide strong support for

¹⁷Consistent with the empirical measurement, we only consider labor earnings of employed workers. The difference would be larger when measured for total income, as we currently exclude non-employed workers who receive only transfer income. Appendix Figure A.2 provides support for this prediction.

Figure 9: Earnings and 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 the end of prime age. Data are from the Panel Study of Income Dynamics.

the model's underlying mechanism. We now use the structural framework to decompose these differences and quantify the underlying channels.

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, while the dashed lines isolate the contributions from wages and human capital. Both components diverge over the life cycle. Early in the career, the earnings gap is primarily driven by job search and the resulting differences in wages. As workers find jobs that offer stable career paths, differences in human capital become increasingly important. By age 55, roughly half of the total earnings gap is attributable to human capital accumulation.

Figure 10b contrasts this benchmark model with a version without heterogeneity in employment stability. In the absence of such heterogeneity, the model fails to generate the large and increasing earnings gap over the life cycle that we document in Section 2.2. Instead, the earnings gap is small and largely flat over the life cycle. Notably, the model shows a sharp discontinuity at age 25. This sharp increase arises mechanically from the classification of careers into stable and unstable groups based on realized employment histories after age 25. Without heterogeneity in job stability, there is no persistence in employment dynamics, as all jobs end with the same high probability, and hence

no systematic divergence in earnings across careers occurs. All differences are a consequence of realization risk, which we document to be small in Figure 1.

In contrast, the model with heterogeneity generates persistent differences in employment stability. This persistence arises from workers sorting over the life cycle into more stable and higher-paying jobs, leading to a continuous divergence in both wages and human capital. In summary, heterogeneity in employment stability is essential for accounting for the empirical patterns of earnings divergence observed in the data.

Figure 10: Earnings difference decomposition



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 54. 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.

We next turn to differences in wealth accumulation. Employment stability affects asset dynamics through three channels. First, higher employment risk increases 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 over the life cycle. This *income growth effect* strengthens the life-cycle saving motive: workers with rising earnings accumulate assets in anticipation of lower income in retirement and to smooth consumption over the life cycle. Conversely, when earnings growth is weak and a career does not take off, the life-cycle saving motive is limited. Third, non-employment spells directly lead to asset decumulation, as 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: as-

sets are accumulated during employment and drawn down during non-employment. Over time, these forces tend to offset each other and primarily serve to stabilize consumption in response to short-run income fluctuations. This pattern of repeated accumulation and decumulation corresponds to a *Sisyphus cycle* of wealth accumulation, in which workers build up precautionary savings and draw them down during non-employment spells, limiting long-run wealth accumulation. In contrast, sustained wealth accumulation over the life cycle is driven by the income growth effect associated with stable careers. With strong and predictable income growth, workers behave as *Modigliani savers*, accumulating assets to smooth income over the life cycle. Hence, employment stability determines which saving regime dominates: unstable careers are governed by the Sisyphus cycle, whereas stable careers follow Modigliani-type life-cycle saving. Thus, employment stability shifts the dominant saving motive from life-cycle saving to precautionary saving.

To quantify the contribution of each channel, we conduct three counterfactual experiments. Each experiment isolates one channel by shutting down the other mechanisms while holding the remaining components fixed.

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 shut down the realization of separation risk. Differences in saving behavior therefore arise solely from differences in employment risk, effectively capturing the precautionary component of the Sisyphus cycle while eliminating both income growth and realized non-employment effects.


Second, to isolate the *income growth effect*, we eliminate job-separation risk from the model. 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 no employment risk. Any difference in wealth accumulation therefore reflects the strength of Modigliani-type life-cycle saving driven by differential income growth.

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

Together, these experiments allow us to attribute differences in wealth accumulation across career paths to distinct mechanisms operating through income risk, income growth, and realized non-employment spells. Importantly, the results quantify how employment stability shifts the dominant saving motive from Modigliani-type life-cycle saving toward precautionary saving and its implied Sisyphus cycle. We refer to the sum of the precautionary saving and non-employment effects as the *consumption smoothing effect*, as it captures buffer-stock dynamics that stabilize short-run income fluctuations, in contrast to the income growth effect, which governs long-run wealth accumulation.

While the decomposition is not exactly additive, Figure 11 shows that it closely accounts for the total difference in wealth accumulation between stable and unstable careers.

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



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.

Figure 11b decomposes the consumption smoothing effect into the 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 gap early in the life cycle, when precautionary motives are 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, meaning that wealth accumulation in unstable careers is higher. Over time, repeated drawdowns during non-employment offset precautionary accumulation, implying that buffer-stock dynamics do not generate sustained wealth growth and that assets are repeatedly used to smooth consumption rather than to accumulate wealth over the life cycle.

4.3 Worker-type heterogeneity

The baseline model emphasizes heterogeneity in job stability and abstracts from direct effects of worker characteristics on separation rates. A natural concern is whether allowing for worker heterogeneity would alter the model’s implications for earnings and wealth accumulation. To address this, we consider two empirically disciplined extensions that incorporate worker heterogeneity in separation rates. First, we introduce *dynamic* worker heterogeneity by allowing job-specific separation rates to depend on human capital. Second, we consider *fixed* worker heterogeneity by introducing worker types that differ permanently in their average separation rates.

Across both extensions, the model’s key prediction for differences in wealth accumulation between stable and unstable careers remains quantitatively very similar to the baseline and, if anything, becomes slightly more pronounced in the dynamic specification. Hence, incorporating empirically disciplined worker heterogeneity does not materially affect the central mechanism of the model: employment stability shapes earnings growth and shifts the dominant saving motive, thereby driving differences in wealth accumulation.

We implement dynamic worker heterogeneity by allowing separation rates to depend on human capital, $\lambda = \lambda(h)$. As human capital evolves over the life cycle, this introduces endogenous, time-varying differences in separation risk across workers. 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). 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 job ladder dynamics.¹⁸ We flexibly parameterize the human-capital component and calibrate it to match the life-cycle profile of separation rates for newly hired workers using the data from Jung and Kuhn (2018). This yields the following calibrated specification:

$$\lambda(h) = \lambda_l \cdot \max\{(2.5 - 1.8h + 0.3h^2), 0.5\},$$


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.

In the fixed worker heterogeneity specification, we capture differences across education groups by

¹⁸Intuitively, newly hired workers accept all offered jobs, so any decline in their separation rates with age reflects human capital accumulation rather than selection into more stable jobs. The average separation rate combines both forces.

Figure 12: Separation rate



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.


lowering the average separation rate. In the data, college-educated workers exhibit separation rates that are about 40% lower than those of non-college workers, while within-group dispersion in non-employment is virtually identical (Section 2).

We implement this by shifting the marginal distribution of separation rates: we reduce the endpoints of the support $[\underline{\lambda}, \bar{\lambda}]$ by 40% while keeping dispersion unchanged. This preserves within-group heterogeneity while lowering the overall level of separation risk. 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 workers in unstable careers face persistently higher separation rates. As a result, earnings differences between stable and unstable careers become slightly larger, reaching about 30 log points by age 50.

For wealth accumulation, two counteracting forces emerge. On the one hand, the stronger increase in wealth relative to earnings in stable careers widens the wealth gap. On the other hand, the stronger decline in earnings relative to wealth in unstable careers dampens the gap in wealth-to-income ratios. On net, the wealth-to-income gap changes only marginally. Focusing on levels of wealth, the ratio of wealth between stable and unstable careers at the end of working life increases by 16% relative to the model without worker heterogeneity. Importantly, the model continues

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



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 the end of prime age. Data are from the Panel Study of Income Dynamics.

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 earnings growth and 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 remains unchanged: employment stability drives earnings growth and shifts the dominant saving motive, thereby generating persistent differences in wealth accumulation.

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 individual level but also has first-order implications for the macroeconomy.

5.1 Effects on wealth inequality

Heterogeneity in 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 and imposes representative-worker labor market dynamics.

We focus on the 90–10 ratio of the wealth distribution to capture the divergent wealth accumulation dynamics of the most stable and unstable careers. Workers with lifetime jobs on stable career paths accumulate substantial wealth over the life cycle as they smooth largely predictable income growth. In contrast, workers in unstable careers accumulate little wealth, as they are tied to buffer-stock dynamics that repeatedly build up and deplete assets, limiting long-run wealth accumulation. Consistent with this mechanism, introducing the empirically observed dispersion in employment stability increases the 90–10 wealth ratio by 28% relative to the representative-worker benchmark.

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. In the SCF, the 90–10 wealth ratio increased by 32% between 1998 and 2022, while the Gini coefficient rose from 0.80 to 0.83 (Kuhn and Ríos-Rull, 2025).¹⁹ While the rise in overall inequality is largely driven by the upper tail,

¹⁹Our model is not designed to match the high level of wealth inequality in the data, as it abstracts from mechanisms such as entrepreneurship, return heterogeneity, portfolio choice, or non-homothetic preferences that are important for explaining wealth concentration at the top.

our results highlight a complementary mechanism operating at the bottom of the distribution. In particular, heterogeneity in employment stability significantly increases wealth inequality among low-wealth households by limiting long-run wealth accumulation in unstable careers. More broadly, our findings show 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 saving 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; Davis and von Wachter, 2011). The structural literature has shown that heterogeneity in employment stability is key to generating such persistent earnings losses while remaining consistent with high average labor market transition 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 larger consumption declines, for example in the aftermath of the Financial Crisis. In contrast, our analysis focuses directly on earnings and consumption dynamics and studies how they depend on pre-displacement employment stability. We compare outcomes across three cases: workers in stable careers, workers in unstable careers, and a model with representative-worker labor market dynamics.

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 18% for workers from stable jobs and 12% for workers from unstable jobs (Figure 14a). While these initial losses are of similar 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 separate in subsequent periods, so that differences between displaced and non-displaced workers dissipate rapidly.

In contrast, job loss from stable employment leads to large and highly persistent earnings losses. Absent displacement, these workers would have remained employed in stable jobs with high probability, implying a persistently higher counterfactual earnings path. Job loss therefore generates a large and persistent deviation from this path.

The model without heterogeneity provides a useful benchmark. We again consider workers with the same levels of assets, human capital, and wages at age 40. Separation rates now depend only on

Figure 14: Effects of displacement by employment stability



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.

age and are calibrated to match the high average separation rates observed in the data. This model generates earnings dynamics that closely resemble those of unstable employment, with strong mean reversion. Although initial losses are comparable to those following displacement from stable jobs, they quickly converge to the path observed for unstable careers. In the absence of heterogeneity in employment stability, there is no persistent component in employment relationships and therefore no persistent divergence in post-displacement earnings.

These results highlight that models without heterogeneity in employment stability, despite matching average mobility rates, generate too much mean reversion relative to the data and fail to account for the persistent earnings losses observed in the data (Jung and Kuhn, 2018).²⁰

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 relatively transitory earnings losses and smooth consumption using precautionary savings (Figure 14b).²¹ Consistent with similar earnings paths, the consumption response in the model without heterogeneity closely aligns with that following job loss from unstable employment. By contrast, workers who lose stable jobs adjust consumption more strongly and persistently to accommodate their large and persistent decline in earnings. While


²⁰A simple back-of-the-envelope calculation illustrates this point: if workers separate into non-employment at a rate of 5% per quarter, this implies roughly a 20% annual probability of falling off the job ladder, leading to an earnings process with strong mean reversion.

²¹We consider a life-cycle model so that even temporary earnings losses reduce total lifetime income and therefore lead to a persistent decline in consumption.

their wealth allows them to smooth the transitory component in the first three years, the sustained decline in earnings translates into a permanent consumption drop of about 13% (Kuhn, 2013).

To summarize, the model without heterogeneity in employment stability does not lie between the stable and unstable dynamics but instead largely mimics the earnings and consumption responses associated with unstable job loss, thereby missing the large and persistent consequences of job loss from stable employment. As a result, it fails to capture the large welfare costs of job displacement that have been widely documented and are considered an important component of large costs of business cycles (Krebs, 2007).

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.

To assess the aggregate implications of these differences, we simulate a deep recession by exogenously displacing 5% 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 to that in a model without heterogeneity and representative-worker labor market dynamics.

Heterogeneity in employment stability strongly amplifies the aggregate consumption response. In the model without heterogeneity, aggregate consumption falls by about 20 basis points. In contrast, in the model with heterogeneity, featuring the empirically large share of stable jobs, the consumption decline reaches about 60 basis points, almost three times as large. For the 2025 U.S. economy, this difference corresponds to an additional drop in household consumption of roughly 80 billion dollars.

These results highlight that not only the size of the unemployment shock matters for aggregate consumption dynamics, but also its composition. In particular, recessions that disproportionately destroy stable jobs generate larger and more persistent consumption declines. This mechanism is

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 magnitude and persistence of consumption responses, as they abstract from heterogeneity in employment stability that shapes income risk, wealth accumulation, and consumption behavior.

Taken together, the results in this section underscore that heterogeneity in employment stability is a central determinant of macroeconomic outcomes. It contributes to higher wealth inequality by generating persistent differences in earnings growth and saving behavior, particularly at the lower end of the income and wealth distribution. At the same time, it amplifies aggregate consumption responses to labor market shocks by introducing heterogeneity in the persistence of income losses and in households’ ability to smooth consumption.

6 Conclusions

Labor markets are characterized by large heterogeneity in employment stability. This paper shows that this heterogeneity is a first-order determinant of earnings dynamics, saving behavior, and macroeconomic outcomes. Using PSID data, we document that more stable careers are associated with significantly higher earnings growth and substantially greater wealth accumulation, even within narrowly defined worker groups. These patterns point to a tight empirical link between employment stability, earnings dynamics, and wealth accumulation that is not captured by existing models.

To interpret these facts, we develop a life-cycle search-and-saving model that combines heterogeneous employment stability, endogenous human capital accumulation, and incomplete markets. The model jointly matches life-cycle profiles of earnings and wealth, as well as their systematic relationship with employment stability. This unified framework allows us to study how heterogeneity of employment stability shapes both income dynamics and household financial behavior.

Our central contribution is to show that employment stability fundamentally reshapes the nature of saving. Stable careers generate sustained earnings growth and give rise to a dominant life-cycle saving motive. In contrast, unstable careers feature limited earnings growth and frequent employment interruptions, so that saving is dominated by precautionary, buffer-stock behavior. This distinction gives rise to a novel characterization of saving regimes: stable careers follow Modigliani-type life-cycle saving, whereas unstable careers are governed by a Sisyphus cycle of repeated accumulation and decumulation. Quantitatively, differences in earnings growth account for the majority of the wealth gap across careers, while consumption-smoothing dynamics play a secondary but important role.

We further show that this heterogeneity has quantitatively important macroeconomic consequences.

First, it amplifies wealth inequality by limiting wealth accumulation at the lower end of the distribution, increasing the 90–10 wealth ratio by 28% relative to a representative-worker benchmark. Second, it amplifies aggregate consumption responses to labor market shocks. Following an unemployment spike, consumption declines are almost three times larger in the model with heterogeneity, reflecting the persistent income losses associated with job loss from stable employment. These results highlight that not only the level but also the persistence of labor market risk is central for aggregate dynamics.

Taken together, our findings establish heterogeneity in employment stability as a key missing ingredient in standard macroeconomic models. Models with representative-worker labor market dynamics abstract from this dimension and therefore understate both wealth inequality and the macroeconomic costs of recessions. Incorporating heterogeneity in employment stability as in our search-and-saving model provides a tractable and empirically disciplined way to jointly account for earnings dynamics, saving behavior, and their aggregate consequences.

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.
- AHN, H. J., B. HOBIJN, AND A. ŞAHİN (2023): “The dual US labor market uncovered,” Tech. rep., National Bureau of Economic Research.
- 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.
- BRENDLER, P., M. KUHN, AND U. STEINS (2024): “To have or not to have: Understanding wealth inequality,” Tech. rep., CEPR Discussion Papers.
- 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.
- HALL, R. E. AND M. KUDLYAK (2022): “Churn and stability: The remarkable heterogeneity of flows among employment, job search, and non-market activities in the us population,” Tech. rep., Technical report, Federal Reserve Bank of San Francisco.
- 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. ŞAHİN (2017): “Gross worker flows over the business cycle,” *American Economic Review*, 107, 3447–76.
- KRUSELL, P., T. MUKOYAMA, AND A. SAHİN (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 PSID data and additional evidence

This section provides additional evidence on the robustness of our empirical analysis. Section A.1 studies whether the switch from annual to biennial sampling in the PSID affects our measure of employment stability based on accumulated non-employment duration. We show that any resulting misclassification of stable and unstable careers is negligible. Section A.2 switches the analysis on income growth from individual earnings to total household income and demonstrates that the differences in income growth across stable and unstable careers become, if anything, more pronounced.

A.1 PSID sample

The PSID provides annual survey data until 1997 and switches to a biennial frequency starting in 1999. This change raises the concern that incomplete observation of labor market histories could lead to misclassification of careers based on accumulated non-employment duration. In particular, missing intermediate observations may cause some workers to be incorrectly classified as having stable or unstable careers.

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



Notes: This figure shows the distribution of accumulated non-employment duration (in quarters) at the end of prime-age working life in the model. It compares the case in which all workers are fully observed to a specification in which workers are only partially observed, with a fraction of observations occurring biennially to match the sampling structure of the PSID.

To quantify the magnitude of this issue, we use the model from Section 3 to construct a model-based estimate of misclassification. We simulate complete labor market histories and then replicate the PSID sampling structure by observing a share of outcomes only every second year. Based on these partially observed histories, we reclassify workers into stable and unstable careers at the end of prime-age working life and compare these classifications to those based on the full histories.


Accounting for the relative weight of pre- and post-1999 observations in the PSID, we find that the implied misclassification rate is approximately 2%.

Figure A.1 illustrates the effect of sampling frequency on the distribution of accumulated non-employment duration. We compare the distribution under full observation to a specification in which workers are only partially observed, with a fraction of observations occurring biennially to match the sampling structure of the PSID. Accounting for the biennial sampling structure slightly increases the mass in the lower part of the distribution, but the overall shape remains very similar. This suggests that the change in sampling frequency in the PSID does not materially distort the measurement of employment stability. Consistent with the low misclassification rate, these findings confirm that our classification of stable and unstable careers is robust to differences in sampling frequency.

A.2 Life-cycle profile of household income

In Section 2, we document that workers in stable careers exhibit stronger life-cycle earnings growth than those in unstable careers, based on earnings of the household head. A potential concern is that this measure may understate effective income dynamics at the household level, as it excludes income from other household members and government transfers that could provide insurance. To address this concern, we repeat the analysis using total pre-tax household income, which includes income from all household members as well as 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 baseline profiles

of log earnings from Figure 3. While household income grows faster than individual earnings for both groups, the divergence between stable and unstable careers becomes even more pronounced. By age 50, workers in stable careers exhibit about 35 log points higher income growth over the life cycle.

If additional household income provided substantial insurance, the gap between stable and unstable careers should narrow when moving from individual earnings to total household income. Instead, the gap widens. This indicates that income from other household members and transfers does not offset weaker earnings growth in unstable careers. More broadly, these results do not provide support for within-household insurance for workers with unstable employment histories.


B Employment stability: worker vs. job component

This appendix provides additional evidence on the sources of heterogeneity in employment stability, with a particular focus on the role of job heterogeneity. We present four complementary pieces of evidence. First, we show that workers who lose stable jobs experience persistently higher non-employment rates relative to a control group, inconsistent with employment stability being a fixed worker characteristic. Second, we examine the persistence of employment stability over the life cycle and show that early-career stability does not predict long-run outcomes once workers experience intermittent non-employment spells. Third, we document large differences in job destruction rates across firms of different ages using data from the Business Dynamics Statistics, providing direct evidence of job-level heterogeneity in separation risk. Fourth, we address the concern that workers may sort into stable or unstable careers based on risk preferences and show that risk tolerance is not systematically related to accumulated non-employment duration. Taken together, these results provide consistent evidence that heterogeneity in employment stability is driven to an important extent by differences across jobs rather than by fixed worker characteristics.

Figure B.1 plots non-employment rates in an event-time window around the end of a job with at least 2 years of tenure. Period 0 denotes the separation event, and the horizontal axis measures time relative to job loss. We compare two groups: all separations and first separations of workers' a career.

Two key findings emerge. First, non-employment rises sharply at separation and remains persistently elevated thereafter. This persistence indicates that the end of a stable job leads to a lasting deterioration in employment outcomes rather than a temporary interruption followed by a return to stable employment. This pattern is difficult to reconcile with employment stability being driven solely by fixed worker characteristics, which would predict a rapid return to stable employment after displacement. Second, the dynamics are very similar when focusing only on first separations. The magnitude and persistence of post-displacement non-employment are comparable across the

Figure B.1: Non-employment rate after displacement




Notes: This figure shows the non-employment rate after a displacement event in period 0 for workers with a previous job tenure of at least two years.

two cases. This suggests that the consequences of losing a stable job are not driven by accumulated human capital or late-career effects, but instead reflect the loss of a stable job match. Taken together, these results support the interpretation that employment stability is, to an important extent, a property of jobs rather than workers.

Next, we examine the persistence of employment stability over the life cycle. If differences in non-employment duration were driven solely by fixed worker characteristics, we would expect a near-perfect correlation between non-employment accumulated early in life and non-employment accumulated later on. In that case, intermittent non-employment episodes during the working life should have little effect on employment stability later in life.

To study the persistence of employment stability over a worker's career, Figure B.2 first 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 life for all workers. Panel B.2 shows a strong positive relationship: workers with high non-employment at young ages tend to remain in the upper part of the distribution later in life. While this pattern is consistent with a worker fixed effect in employment stability, the relationship weakens substantially once we focus on workers who experience intermittent non-employment during mid-career. This is shown in Panel B.2, where the correlation declines markedly among workers who accumulate two or more quarters of non-employment between ages 31 and 52. Hence, workers who initially appear to have stable careers often move into unstable trajectories following adverse employment realizations. This breakdown of persistence is inconsistent with a model in which employment stability is solely determined by fixed worker types. Instead, it points to an important role for shocks and job-level heterogeneity in shaping employment trajectories over the life cycle. These results further support the interpretation that

Figure B.2: Persistence of non-employment duration



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 quarters of non-employment between ages 30 and 52.

employment stability is not a permanent worker attribute.


As a third step, we provide direct evidence on the job-level sources of heterogeneity in employment stability. Figure B.3 uses data from the Business Dynamics Statistics (BDS) to document substantial heterogeneity in job destruction rates across employers of different age. We consider two measures of job destruction: the total job destruction rate (Figure B.3a) and job destruction due to firm closure (Figure B.3b). In both cases, we remove year and industry fixed effects (Panels (a) and (b)) and, alternatively, year and MSA fixed effects (Panels (c) and (d)).²²

Figure B.3a shows substantial heterogeneity in job destruction across employers. Young firms exhibit job destruction rates that are roughly twice as high as those of older firms. While total job destruction may reflect both worker-initiated separations and firm-side shocks, Figure B.3b isolates job loss due to firm closure. For this measure, differences across firms are even more pronounced: job destruction rates at the youngest firms are up to four times higher than at the oldest firms. Figure B.3c and Figure B.3d show that controlling for year and MSA fixed effects yields very similar results, indicating that this heterogeneity is not driven by local labor market conditions but reflects differences across employers.

These findings provide direct evidence of large differences in job-level separation risk across employers. Importantly, because job destruction due to firm closure is driven by firm-level shocks, it is orthogonal to worker-specific characteristics. This implies that workers are exposed to funda-

²²The BDS data do not provide publicly available data where industry and geographical breakdown is available in the same file.

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




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.

mentally different employment risks depending on the firms they match with. Taken together with the displacement and persistence evidence above, these results support the view that heterogeneity in employment stability arises to an important extent from job-level differences rather than solely from fixed worker characteristics.

In a final step, we examine whether differences in employment stability reflect sorting by risk preferences. If workers with lower risk aversion are more willing to accept unstable employment, we would expect a positive relationship between non-employment duration and risk tolerance. In the PSID data, we can measure risk tolerance following Kimball et al. (2009). This measure can be mapped, under a CRRA utility function, into a coefficient of relative risk aversion: a risk tolerance

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



Notes: This figure presents a binned scatter plot of accumulated non-employment duration at age 50 against the risk tolerance measure. A risk tolerance of 0.27 corresponds to a coefficient of relative risk aversion (CRRA) of approximately 6.7.

of 0.27 corresponds to a CRRA of approximately 6.7, and higher risk tolerance corresponds to lower risk aversion. As in the main text, we measure employment stability by accumulated non-employment duration (in quarters) at age 50.

Figure B.4 shows no systematic relationship between accumulated non-employment duration and risk tolerance, even though the risk tolerance measure exhibits substantial variation across individuals. The absence of a positive correlation indicates that workers in unstable careers are not simply more risk tolerant. This finding does not provide support for the idea of sorting by risk preferences as a primary driver of differences in employment stability. We corroborate this finding in the SCF data in Appendix Figure C.7 where in panels (b) and (d) we also control for risk attitude. In Section 2.3, we provide additional regression evidence on the relationship between employment stability and wealth accumulation and show that it is also robust to controlling for worker fixed effects and age provides interacted with risk attitude. Hence, also or evidence on wealth accumulation and employment stability is consistent with the view that it is mainly driven by heterogeneity in employment stability stemming from job heterogeneity.

C Employment stability, earnings growth, and wealth accumulation


This section establishes the robustness of the relationship between employment stability and wealth accumulation along several dimensions. Section C.1 shows that the results are not driven by the classification of careers at the end of working life by using an alternative classification of stable

and unstable careers based solely on information available up to age 30. Section C.2 demonstrates that the findings are robust to alternative measures of employment stability in the PSID. Finally, Section C.3 provides independent validation using SCF data, confirming the relationship between employment stability and wealth accumulation in a separate, high-quality microdata source.

C.1 Ex ante vs. ex post classification of employment histories

In the main analysis, we classify stable and unstable careers based on labor market histories observed between ages 25 and 50. As an alternative, we consider here an ex-ante classification that relies only on information between ages 25 and 30.

Figure C.1: 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 careers consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of non-employment duration at age 30. The wealth-to-income ratios are normalized to zero at age 20. Data are from the Panel Study of Income Dynamics.

Figure C.1 shows the life-cycle profiles of wealth-to-income ratios under this ex-ante classification. The results closely mirror the baseline findings: stable and unstable careers exhibit very different wealth accumulation paths over the life cycle, with a similar gap in wealth-to-income ratios by the end of working life. This robustness reflects the strong persistence of employment stability over the life cycle. Careers classified as stable early on remain overwhelmingly stable: only 10% of workers classified as stable at age 30 are reclassified as unstable by age 50. This high persistence implies that employment stability is largely determined early in the career and supports the interpretation of our baseline classification as capturing durable differences in career trajectories rather than ex-post realizations.

C.2 PSID data: wealth accumulation

This section examines the robustness of the relationship between employment stability and wealth accumulation using alternative measures of employment stability within the PSID sample. In addition to accumulated non-employment duration used in the baseline analysis in Section 2, we consider measures based on unemployment duration and spells, non-employment spells, and tenure. Using different measures, we always find a consistently strong relationship between employment stability and wealth-to-income ratios. We further show that this relationship is robust to using a measure of permanent lifetime income rather than contemporaneous income at age 50.

Figure C.2: Different measures for employment stability




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.

As a first step, we show that the alternative measures all capture heterogeneity in employment stability by reporting in Figure C.2 the correlations across these measures. We find strong and systematic relationships: non-employment duration, unemployment duration, and the number of

non-employment and unemployment spells are all positively correlated, while employer tenure, as expected, is negatively correlated with these measures. The magnitudes of these correlations are economically large, indicating that the different measures rank careers in a very similar way. The strong correlation between unemployment and non-employment measures also provides support for our modeling approach of combining these states for workers closely attached to the labor market, following the argument by Kudlyak and Lange (2014). Taken together, these patterns confirm that the alternative measures consistently capture the same underlying heterogeneity in employment stability.

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




Notes: This figure shows binned scatter plots of wealth-to-income ratios for different measures for employment stability and income. Panel (a) shows the relationship with employer tenure (excluding zero tenure). Panel (b) with unemployment duration. Panels (c) and (d) show the correlation with the number of unemployment and nonemployment spells.

In a second step, we examine the relationship between wealth-to-income ratios and these alternative measures of employment stability in Figure C.3. Using unemployment duration, the number of unemployment and non-employment spells, and employer tenure, we continue to find the robust

positive correlation between stable employment and wealth-to-income ratios documented in Section 2. For tenure, we use contemporaneous tenure at age 50. Unlike accumulated non-employment duration, this measure does not reflect the entire lifetime employment history if a worker has recently switched jobs; we therefore exclude observations with zero tenure. Variation in wealth-to-income ratios across tenure levels closely matches the differences in wealth accumulation between stable and unstable careers at age 50, amounting to about 1.5 years of annual income in Figure 5. In the next section, we further corroborate these findings using SCF data, where we employ tenure as well as the number of employers as measures of employment stability.

In the baseline specification, we construct wealth-to-income ratios using contemporaneous income at age 50. We therefore assess robustness to an alternative income measure based on permanent lifetime income. Specifically, we use average labor income between ages 25 and 50 as permanent income measure when constructing the wealth-to-income ratio. Figure C.4 shows that this alternative measure has a negligible effect on the relationship between employment stability and wealth accumulation, with the negative relationship between employment stability and wealth-to-income ratios remaining virtually unchanged.

Figure C.4: Wealth-to-income ratios with permanent income



Notes: This figure shows binned scatter plot of wealth-to-income ratios against non-employment of workers at age 50. Wealth-to-income ratios are constructed using permanent income rather than contemporaneous income at age 50. Permanent income is constructed as average labor income between age 25 and 50.

C.3 SCF data

In a final step, we use data from the Survey of Consumer Finances (SCF), the primary source on the distribution of income and wealth in the United States (see, for example, Kuhn and Ríos-Rull, 2016; Bricker et al., 2017; Kuhn et al., forthcoming), to provide independent validation of the relationship between employment stability and wealth accumulation. In addition to detailed information on household balance sheets, the SCF contains also detailed information on employment

histories, including tenure and the number of employers. Using PSID data, we have shown that tenure is closely related to accumulated non-employment duration, our baseline measure of employment stability (Figure C.2).²³ The SCF therefore allows us to construct comparable measures of employment stability in an independent dataset and to study the relationship between employment stability and wealth accumulation.

We pool all SCF survey waves from 1992 to 2016 and restrict the sample to households with employed household heads aged 20 to 60. As our model abstracts from self-employment, we drop households with self-employed household heads, as well as households with extreme wage observations, defined as wages below 75% of the minimum wage.²⁴ Additionally, we exclude the top 1% of households by wealth and earnings, as we do not provide a theory of the very right tail of these distributions. Similarly, we exclude households in the bottom 1% by earnings and those with negative wealth. For wealth and income, we follow Bricker et al. (2017) and Kuhn and Ríos-Rull (2016) for the construction of variables: Household wealth is the difference between household assets and debt and household income is gross income from all sources including transfers, and earnings is income from wages and salaries.

Figure C.5 corroborates large heterogeneity in employment stability by showing the life-cycle profiles of tenure and the number of employers in the SCF data. We find that both profiles are positively correlated with age. Considering the mean, median, and 75th percentile of the tenure distribution in Figure C.5a, we observe levels and age profiles that closely match those in the PSID data shown in Figure 2. The life-cycle profiles for the number of employers in Figure C.5b provide a similar picture. The mean number of employers increases linearly up to age 40, after which growth slows in the second part of working life. On average, an American worker has worked for four employers by the end of his or her working life.


Figure C.6 presents the estimates of the empirical measures of employment stability and their relationship with wealth-to-income ratios for a pooled sample of household heads across all ages. In Figure C.6a, we observe an almost linear relationship between tenure and wealth-to-income ratios. This raw correlation, however, could simply reflect the fact that both tenure and wealth-to-income ratios increase with age. In Figure C.6b, we therefore examine the relationship between wealth-to-income ratios and tenure while controlling for age. We continue to find a positive relationship between employment stability and wealth accumulation, albeit with a smaller slope. Quantitatively, the slope is similar, but the support is wider than what we observe when focusing only on 50-year-old workers in the PSID.

Figures C.6c and C.6d corroborate the previous findings using the total number of employers as

²³Employer 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.

²⁴We rely on individual hours and earnings information in the SCF data to construct wages.

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




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.

an alternative measure of employment stability. The role of age in shaping this relationship is particularly evident. Before controlling for age (Figure C.6c), there is no apparent relationship between wealth-to-income ratios and the number of employers. After controlling for age (Figure C.6d), we find a negative relationship between the number of employers and wealth-to-income ratios. Interpreting a higher number of employers as reflecting less stable employment, this implies that greater employment stability (fewer employers) is positively related to wealth-to-income ratios.

To ensure that the observed relationship between employment stability and wealth-to-income ratios is not driven by demographic characteristics, systematic differences across industries and occupations, or variation in risk attitudes, we perform additional robustness checks controlling for a rich set of observable household characteristics in the SCF. Figure C.7 shows the correlation between the SCF measures of employment stability and wealth-to-income ratios after controlling for these additional observable characteristics. In all cases, we continue to find a clear positive correlation between more stable careers and wealth accumulation. In the first column of panels, we include dummy variables for age, education, occupation, and industry. The relationship between tenure and wealth-to-income ratios (top row) remains unchanged and economically significant. The same holds for the relationship between the number of employers and wealth-to-income ratios (bottom row). In the second step, we additionally control for differences in risk attitudes, following the approach in the PSID analysis in Section B, using survey-based measures of risk attitudes from the SCF. As shown in the second column of panels, the relationship between employment stability


Figure C.6: Wealth-to-income ratios, tenure, and number of employers




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.

and wealth accumulation remains virtually unchanged when these additional controls are included. We therefore conclude that our findings based on SCF data are robust to controlling for observable characteristics and strongly corroborate the findings based on PSID data.


Figure C.7: 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

This appendix collects additional details on the model, the numerical solution approach, the calibration, and additional untargeted model predictions. Section D.1 presents the value functions for the transition phase, which are omitted from the main text for brevity. Section D.2 describes the numerical solution and the simulated method of moments calibration. Within this section, Section D.2.3 provides an intuitive discussion of how the targeted moments discipline the model parameters, and Section D.2.4 presents the estimated job-offer distribution. Finally, Sections D.3.1–D.3.5 provide additional evidence on model fit, covering earnings and consumption inequality, employment stability measures from the SCF, life-cycle earnings dynamics, the joint distribution of income and wealth, and the individual consequences of job displacement.

D.1 Value functions for the transition phase

The transition phase has stochastic duration, so agents solve a fixed-point problem. Age is no longer a state variable, and the value functions do not carry a time index. The Bellman equations for the transition phase follow directly from those of the working phase. The only difference is that, at the end of each period, the worker retires and enters the retirement phase with probability ψ . All other decisions are identical to those in the working phase.

The value function of an employed worker at the beginning of the transition phase, V_e^T , is given by the expectation over employment states resulting from 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 a non-employed worker at the production stage 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 worker at the production stage is

$$V_e^{T,P}(a, w, \lambda, h) = \max_{\{c, a' \geq 0\}} u_j(c) + \beta \left[\psi V_r(a', w, h, j_r = 1) + \right. \quad (9)$$

$$\left. (1 - \psi) (\pi_e V_e^{T,S}(a', w, \lambda, h) + (1 - \pi_e) V_e^T(a', w, \lambda, h)) \right] \quad (10)$$

$$s.t. \quad c = (1 + r)a + y(w, h, e) - a',$$

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

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

where N_w denotes the number of wage realizations in the support of the offer distribution, and N_λ denotes the number of separation-rate realizations.

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

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

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

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

D.2 Model solution and calibration

D.2.1 Solving the model

We solve the model by backward induction and use on-grid search for the joint consumption-saving and effort choice problem. We discretize the state space for assets, wages, job destruction probabilities, 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 human capital grid by $\mathcal{H}$, the full state space is given by the Cartesian product

$$\mathcal{A} \times \mathcal{W} \times \mathcal{L} \times \mathcal{H} = \{a_1, \dots, a_{N_a}\} \times \{w_1, \dots, w_{N_w}\} \times \{l_1, \dots, l_{N_l}\} \times \{h_1, \dots, h_{N_h}\}.$$

The upper bounds of the grids are chosen sufficiently large so that they do not bind in the optimal solution of the calibrated model.

Given the standardization of the offer distribution, both wages and job destruction probabilities

follow truncated exponential marginal distributions with support on $[0, 1]$.²⁵ To allow for correlation between the two marginal distributions, we construct the joint distribution over standardized wages and job destruction probabilities, $F(w^*, \lambda^*)$, using Frank's copula C_θ , where the parameter θ governs the correlation between w^* and λ^* . We then discretize this joint distribution into bins corresponding to the 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 combine the resulting policy rules with randomly generated shocks to simulate the life cycles of 200,000 agents.

““

D.2.2 Calibration of parameters

We calibrate a subset of the model parameters listed in Table 3 using a simulated method of moments. Specifically, we choose parameters to minimize the sum of squared percentage deviations between model-implied life-cycle profiles and their empirical counterparts. The targeted moments are the age profiles of separation rates, job-to-job transition rates, job-finding rates, tenure (mean, median, and 75th percentile), average log earnings, and the wealth-to-income ratio.

Let θ denote the parameter vector. The objective function is

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

where hatted variables denote empirical moments, and each summand corresponds to one targeted life-cycle profile.

D.2.3 Identification of model parameters


All model parameters are calibrated jointly, and we do not provide a formal proof of identification. Instead, we provide an intuitive discussion of how parameters map into model predictions and how the empirical moments used in the calibration discipline them.

²⁵We standardize the support of wages and job destruction probabilities to obtain a tractable numerical implementation of the joint distribution. Details are provided in Section 3.1.

We organize the discussion around three core components: job-offer dynamics, human capital accumulation, and preferences together with savings behavior. Because parameters interact throughout the model, we focus on the economic channels through which each parameter mainly shapes observable life-cycle patterns. This clarifies how the targeted moments discipline the parameter values. To make this mapping transparent, we structure the discussion around a series of comparative statics experiments in which we vary the calibrated parameters and show how they affect model moments.

We begin with the offer arrival rates for employed and non-employed workers, denoted by π_e and π_n . These parameters are disciplined by observed life-cycle patterns in job-to-job transitions and job-finding rates. Higher arrival rates increase worker mobility and accelerate matching with higher-quality jobs. The levels of the age profiles of employment-to-employment transitions and job-finding rates therefore discipline these parameters.

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




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_n (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 rates for employed and non-employed workers. When we lower the arrival rates, π_e and π_n , by 10 percentage points, both life-cycle profiles shift downward, reflecting the lower frequency with which workers receive job offers. The decline in job-to-job transitions is particularly pronounced among younger workers. Early in the career, mobility is a key mechanism for climbing the job ladder. With fewer offers arriving, opportunities to switch jobs decline, leading to an especially strong reduction in job-to-job transitions at younger ages.

Next, we explain how to pin down the joint job-offer distribution, $f(\lambda, w)$. Its parameters are disciplined by matching the joint dynamics of wages, tenure, and worker flows. In Section D.2.4, we show and discuss the calibrated joint distribution. The marginal distribution of job stability,

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




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).

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

Next, we discuss how the joint job-offer distribution, $f(\lambda, w)$, is disciplined by the data. Its parameters are identified from the joint dynamics of wages, tenure, and worker flows. Section D.2.4 presents and discusses the calibrated joint distribution. The marginal distribution of job stability, governed by ψ_λ , is disciplined by the life-cycle profile of separation rates together with the level and dispersion of tenure. The parameter ψ_λ determines the extent to which job offers are concentrated among relatively unstable versus highly stable matches. Figure D.2 illustrates the implications of shifting the marginal distribution toward fewer stable jobs. When stable matches are relatively scarce, sorting into durable employment relationships becomes more gradual. Early in their careers, workers disproportionately accept positions with higher separation risk and therefore experience repeated job loss before transitioning into more stable matches. This mechanism raises aggregate separation rates and lowers average tenure, particularly at younger ages. Over the life cycle, the slower reallocation into stable jobs also delays the accumulation of long-tenure matches, resulting

in a flatter tenure profile.

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



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. This parameter is disciplined by the slope and curvature of average wage growth over the life cycle, together with the dispersion of tenure at younger ages. Economically, ψ_w determines how easily workers move up the job ladder through job-to-job transitions and how long they have to wait for better-paying job opportunities.²⁶

Figure D.3 illustrates the effect of lowering the probability of receiving high-wage offers. When such opportunities become less common, upward mobility declines. Because young workers rely heavily on job-to-job moves to generate rapid wage growth, reduced wage upgrading translates into more modest earnings growth early in the career. The figure shows that wage growth slows and wage profiles become noticeably flatter. Fewer opportunities to move into better-paying jobs, and thus slower progression up the wage ladder, also delay sorting into better matches, thereby affecting the dispersion of tenure early in working life.

The parameter θ , which governs the correlation between wages and separation risk, is disciplined by 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 moving up the wage ladder simultaneously transition into jobs with lower separation risk. This reinforces life-cycle sorting, as workers become increasingly concentrated in stable, high-wage matches. A higher correlation also reduces the trade-off between wages and job stability because high-wage jobs are

²⁶Section D.3.3 below discusses the evidence from Topel and Ward (1992) on the contribution of job switching to early-career wage growth and shows that the model matches this untargeted evidence.

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




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 0.47 to 0.74).

more likely to be stable. As a result, job-to-job transition rates and separation rates decline more strongly with age when wages and job stability are more closely aligned.

Figure D.4 further shows that a stronger correlation also affects tenure profiles. When wages and job stability are highly positively correlated, workers who sort into high-paying jobs are more likely to remain in these jobs for extended periods. Consequently, tenure rises more strongly toward the end of prime working life. The joint behavior of mobility, separations, and tenure dispersion therefore disciplines the magnitude of θ in the model.

Figure D.5: Change in parameters governing human capital $\bar{p}_H$, κ , and 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 $\bar{p}_H$ (reduced from 0.071 to 0.035), κ (increased from 0.65 to 0.96), and ρ (reduced from 0.984 to 0.982).

The parameters governing the evolution of human capital, $\bar{p}_H$, κ , and ρ , are disciplined by the life-cycle profile of average earnings. The parameter $\bar{p}_H$, which determines the baseline probability of

human capital accumulation, is disciplined by overall growth of earnings (Figure D.5(a)). The utility cost parameter κ is disciplined by the flattening of earnings late in the life cycle, when earnings growth is driven primarily by human capital accumulation rather than job mobility (Figure D.5b). A higher value of κ increases the cost of skill investment, reducing human capital accumulation and thereby lowering earnings growth at older ages. As a result, the life-cycle earnings profile becomes flatter. Finally, the persistence parameter ρ , which governs how much and how quickly human capital accumulation probabilities decline with age, is disciplined by the concavity of the life-cycle earnings profile. If human capital investment opportunities decline more strongly over the life cycle, corresponding to a lower value of ρ , the earnings profile becomes more concave.

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 0.976 to 0.971).


Finally, the time discount factor, β , is disciplined by the life-cycle profile of the wealth-to-income ratio. A lower value of β implies greater impatience, as workers place less weight on future utility. This shift in intertemporal preferences increases current consumption and reduces savings. As a result, 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

Figure D.7 shows the estimated joint distribution of wage offers and separation rates. Panel D.7a displays the full three-dimensional distribution and highlights its asymmetric shape, with most probability mass concentrated in low-wage, unstable jobs. The estimated distribution also features a negative correlation between wages and separation rates, implying that high-wage jobs tend to have low separation risk, while low-wage jobs tend to have high separation risk (Jung and Kuhn, 2018; Jarosch, 2023).

Panel D.7b presents slices of the joint distribution by showing the conditional distribution of separation rates at different wage levels. These slices show that the distribution of separation rates

Figure D.7: Job-offer distribution



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.

for low-wage jobs first-order stochastically dominates the corresponding distribution for high-wage jobs. Hence, low-wage jobs are more likely to be less stable. Jung and Kuhn (2018) provide direct empirical evidence for this relationship based on SIPP data. Taken together, the calibrated distribution implies a positive correlation between job stability and wages, with stable and high-paying jobs being relatively scarce, while most offered jobs are unstable and low paying. In short, good jobs are hard to find.

D.3 Additional results on model fit


This appendix assesses the model's additional untargeted predictions by comparing them to their empirical counterparts. We first show in subsection D.3.1 that the model is consistent with empirical evidence on earnings and consumption inequality over the life cycle. Subsection D.3.2 then documents that the model also matches measures of employment stability observed in the SCF data, specifically the cross-sectional distribution of employer tenure and the number of employers. Subsection D.3.3 shows that the model is not only consistent with earnings growth (Section 3.2) and inequality (subsection D.3.1), but also matches key statistical properties of the earnings process. Subsection D.3.4 discusses the model's endogenous predictions for the joint distribution of earnings and wealth and their conditional dynamics. Finally, subsection D.3.5 studies the individual worker consequences of job displacement that underlie the macroeconomic consequences analyzed in Section 5.2.

D.3.1 Earnings and consumption inequality

As a first step, we explore the model's ability to account for the life-cycle pattern of earnings variance. The calibration targets average earnings growth and heterogeneity in employment stability, but not the variance of earnings. Figure D.8a shows that the model nevertheless closely matches the life-cycle increase in log earnings variance observed in the data. In both the model and the data, earnings variance increases strongly with age.

This pattern is difficult to generate in standard labor market search models. As discussed in Lise (2012), Jung and Kuhn (2018), and Hubmer (2018), search models often struggle to produce the observed rise in earnings dispersion over the life cycle. Augmenting job search models with differences in human capital accumulation provides one way to account for this increase (Jung and Kuhn, 2018; Hubmer, 2018). We build on this approach by allowing employment histories to affect human capital accumulation over the life cycle. This channel helps translate differences in search outcomes and employment stability into persistent earnings differences, as discussed in Section 4.1. The model's ability to match the untargeted life-cycle increase in earnings variance therefore supports the broader mechanism linking employment stability to long-run earnings dynamics.

Figure D.8: Earnings and consumption inequality



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.


In a second step, we examine the model's predictions for consumption inequality. Consumption inequality reflects the interaction between the endogenous earnings process and the consumption-saving decision, an interaction that we further analyze in Section D.3.4 when studying the joint distribution of earnings and wealth and their dynamics. Figure D.8b shows that the model also

aligns with empirical estimates of the life-cycle increase in the variance of log consumption. These empirical estimates differ across studies (Deaton and Paxson, 1994; Primiceri and Van Rens, 2009; Aguiar and Hurst, 2013), and the model lies in the middle of the range of existing estimates. Comparing Figures D.8a and D.8b, we also note that the increase in the variance of consumption is about half as large as the increase in the variance of earnings. Hence, consumption is through self-insurance partly insulated from earnings dynamics.

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

In Appendix C.3, we discuss evidence on employment stability and wealth accumulation based on SCF data, using tenure and the number of employers as alternative measures of employment stability. Here, we show that the model also matches the distributions of these measures. Figure D.9 compares the cross-sectional distributions of tenure and the number of employers in the model and in the data. The number of employers refers to the number of employers for which a worker has worked for at least one year over her working life. We pool all workers and report the corresponding distributions using histograms. When pooling model data, 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



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.

We find a close alignment between model and data, supporting both the consistency of the survey evidence between SCF and PSID and the model's labor market dynamics. This fit is notable because the SCF data were not used in the model calibration. The result therefore suggests that cross-sectional survey data, such as the SCF, can provide useful information for disciplining model parameters when long panel histories are unavailable. This is important because it points to a

substantially less data-intensive route for calibrating models of employment stability and wealth accumulation.

D.3.3 Life-cycle earnings dynamics

The model matches the life-cycle profiles of earnings growth (Figure 7) and earnings dispersion (Figure D.8a). In this section, we provide a more demanding assessment of the model by examining whether it also matches the underlying dynamics of individual earnings. This is important because consumption-saving models often take earnings processes as exogenous inputs, typically parameterized using reduced-form statistical representations. Our search-and-saving model, by contrast, generates earnings dynamics endogenously through labor market mobility and human capital accumulation. Comparing the model-implied earnings process to these empirical representations therefore provides a direct way to evaluate whether the structural labor market mechanism generates realistic earnings dynamics.

We proceed in three steps. First, we estimate a standard decomposition of earnings dynamics into permanent and transitory components and compare the variances of the corresponding innovations in model and data (Meghir and Pistaferri, 2004; Blundell et al., 2008; Heathcote et al., 2010). Second, following Guvenen et al. (2021a), we examine higher moments of the earnings growth distribution. Third, we decompose earnings growth into the contributions of human capital accumulation and job switching and show that the model aligns with the evidence on early-career wage growth in Topel and Ward (1992). We close by discussing the model’s implications for frictional wage dispersion.


In the first exercise, 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 of Blundell et al. (2008) to the simulated data. The estimated variance of the permanent component is 0.027, which lies well within the range of empirical estimates. Blundell et al. (2008) estimate time-varying variances of the permanent component between 0.01 and 0.03 for the period from 1980 to 1990.²⁷ Empirical estimates of the variance of transitory shocks are harder to compare because they also include the contribution of measurement error, which is likely substantial in the data. It is therefore not surprising that empirical estimates of the variance of transitory shocks, ranging from 0.03 to 0.05, are substantially larger than the model-implied estimate of 0.015. We interpret this difference as reflecting the contribution of measurement error.

In a second step, we consider the findings of Guvenen et al. (2021a), who emphasize that earnings growth rates are not normally distributed but instead exhibit substantial negative skewness and high

²⁷Heathcote et al. (2010) provide a detailed discussion of different estimation approaches. We use this estimation as a reduced-form description of earnings dynamics without requiring the process to be the true underlying process or the estimates to be unbiased. See Daly et al. (2016) for further discussion.

excess kurtosis. As shown by Hubmer (2018), these patterns can be well explained by a life-cycle job ladder model with a human capital process. Our model generates similar implications. Figure D.10 shows the distribution of one-year earnings growth rates in the model, together with a normal distribution with the same standard deviation. Earnings changes in the model are left-skewed and strongly clustered around zero, consistent with the empirical evidence.

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.

Finally, we use the model to decompose life-cycle earnings growth and the increase in earnings variance into a component due to search for higher wages and a component due to human capital accumulation. Figure D.11 presents this decomposition for both the mean and the variance of earnings and highlights human capital accumulation as the main driver of life-cycle earnings dynamics.

This decomposition aligns well with empirical evidence. Topel and Ward (1992) estimate the contribution of employer switching to wage growth after labor market entry and find that search for better-paying employers accounts for about one-third of wage growth during the first ten years in the labor market. Looking at the decomposition of mean earnings in Figure D.11a, we find a similar pattern in the model. Between ages 20 and 30, climbing the job ladder is an important driver of earnings growth and accounts for roughly one-third of the increase, close to the estimate in Topel and Ward (1992). After age 30, wage growth from job search flattens out, in line with the slowdown in employer switching shown in Figure 6. Human capital accumulation then accounts for almost the entire increase in earnings once most workers have found stable jobs.

For the increase in earnings variance, we find a similar decomposition. At age 20, all workers start with the same level of human capital, so differences in entry wages account for all earnings dispersion. Over time, workers climb the job ladder by leaving lower-paid jobs and accepting

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



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.

better-paid jobs. This process compresses wage differences and therefore contributes negatively to the increase in life-cycle earnings inequality. The reason is that workers who already hold well-paid jobs receive fewer attractive outside offers, as many job offers pay lower wages, and these workers are therefore more likely to remain with their current employer. As a result, initial wage differences decline over the life cycle.

This mechanism highlights the general challenge of accounting for rising life-cycle earnings dispersion through employer differences alone. As our decomposition shows, differences in human capital accumulation are the main 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% of the search component. The small contribution of search frictions, or frictional wage dispersion, is consistent with Hornstein et al. (2011) and Hagedorn and Manovskii (2010), who point to 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 is that the model jointly matches evidence on earnings dynamics and earnings inequality. Hornstein et al. (2011) show that models with on-the-job search and a homogeneous separation rate, calibrated to match the observed average separation rate, can generate large wage inequality among otherwise identical workers. In such models, however, wage differences are highly transitory because jobs are short-lived on average. Large wage inequality arises from a long and stretched wage ladder: workers start at the bottom, and the gradual process of moving up the ladder spreads workers across wage levels.

Our model relies on a different mechanism. Heterogeneity in employment stability increases the persistence of search outcomes because some workers sort into stable jobs while others repeatedly fall off the job ladder. This persistence allows the model to be jointly consistent with cross-sectional earnings inequality, persistent employment relationships, and observed earnings dynamics over the life cycle.

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 search-and-saving model combines this consumption-saving block with a labor market block that generates earnings dynamics endogenously. This joint modeling of earnings dynamics and wealth accumulation makes the joint distribution of income and wealth, as well as the wealth transition dynamics implied by the model, natural objects for evaluating the model’s performance. This comparison is particularly useful because it links the model’s endogenous earnings dynamics directly to wealth accumulation, the central outcome of interest in this paper. In addition, unlike panel data on consumption and income dynamics, the joint distribution of income and wealth is readily observed in datasets such as the SCF (Brendler et al., 2024).


For the comparison of the joint distribution in the model and the data, we split households into income and wealth quintiles and study how households are distributed across wealth quintiles conditional on their income quintile. To abstract from differences in age composition, we focus in model and data on households aged 40 to 50. Panels (a) through (c) in Figure D.12 show selected conditional wealth distributions by income quintile. Table D.1 reports the corresponding population shares across all income-wealth cells for the model and the SCF data.

We then study wealth dynamics by comparing how households move across wealth quintiles over time. Panels (d) through (f) in Figure D.12 show conditional distributions over wealth quintiles five periods ahead, starting from a given wealth quintile at age 40. Equivalently, these panels show selected rows of a five-step wealth transition matrix. Table D.2 reports all corresponding conditional probabilities. To study these wealth dynamics, we follow previous research (Díaz-Giménez et al., 2011; Kuhn et al., forthcoming) and use PSID data from 1984 to 1999, as the repeated cross sections of the SCF do not allow us to trace households over time.²⁸

Looking first at the joint distribution in Panels D.12a, D.12b, and D.12c, we find that the model provides a very good fit to this untargeted dimension of the data. This comparison is particularly informative about the model’s wealth accumulation mechanism. Income is a flow and can change

²⁸Pfeffer et al. (2016) compare SCF and PSID data and conclude that, except for the very top of the wealth distribution, the two surveys provide consistent wealth distributions for the vast majority of households.

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



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

relatively quickly, whereas wealth is a stock that reflects the accumulated history of earnings, saving decisions, and past shocks. The close fit of the joint income-wealth distribution therefore suggests that the model not only generates realistic earnings dynamics, but also maps these dynamics into wealth accumulation patterns that are consistent with those observed in the data.

Looking directly at wealth dynamics in Panels D.12d, D.12e, and D.12f, we find that the model closely matches the main transition patterns in the PSID data. As the PSID regularly surveyed wealth only every five years, we focus on five-year transition probabilities. Both in the model and in the data, households' positions in the wealth distribution are highly persistent, with most of the probability mass concentrated on the diagonal and the adjacent wealth quintiles. This pattern is especially clear at the bottom of the distribution: 94% of households in the model and 84% in the PSID data remain within the bottom two wealth quintiles over a five-year horizon. The model also closely matches the middle of the distribution. For households starting in the middle quintile, both model and data imply that close to 90% remain within the three middle quintiles. At the top of the distribution, persistence is somewhat stronger in the model: 82% of households in the model and

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	0.05	0.03	0.01	0.00
I20-40	0.05	0.07	0.05	0.03	0.00
I40-60	0.03	0.04	0.06	0.06	0.02
I60-80	0.01	0.02	0.05	0.08	0.05
I80-100	0.00	0.01	0.02	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

65% in the PSID data remain in the top wealth quintile. Hence, the model generates somewhat more persistence than observed in the PSID data, especially at the top, but it captures well the overall pattern of limited wealth mobility and predominantly local movements across the wealth distribution. It is important to note that measurement error and imputation of wealth information in the PSID may generate spurious transitions across wealth quintiles. Such measurement issues may therefore overstate measured wealth mobility in the data and account for part of the lower empirical persistence.


Table D.2: Wealth transition matrix, age 40

	W0-20	W20-40	W40-60	W60-80	W80-100
Model					
W0-20	0.70	0.24	0.05	0.01	0.00
W20-40	0.23	0.50	0.24	0.04	0.00
W40-60	0.06	0.20	0.50	0.22	0.02
W60-80	0.01	0.04	0.22	0.55	0.17
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

Finally, we provide additional results on the consequences of job displacement that complement the analysis of 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 displaced workers (Jacobson et al., 1993). We compare otherwise identical workers at age 40, where one worker loses her job while the other 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 after they have been exogenously displaced into non-employment at age 40.²⁹ Empirical studies document that job displacement leads 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 these losses in structural models (Jung and Kuhn, 2018; Jarosch, 2023). Models without heterogeneity in employment stability struggle to account for the persistence of earnings losses (Low et al., 2010).

Figure D.13: Cost of displacement



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.

Figure D.13a shows that the baseline model with heterogeneity in employment stability implies large and persistent earnings losses for the average 40-year-old worker. The initial earnings drop after displacement amounts to more than 15%. Over the subsequent five years, displaced workers recover about half of the initial earnings loss, but there is little further catch-up afterward. The figure also decomposes earnings into wages and human capital, which helps uncover the mechanism

²⁹This 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, which are the key concern in the empirical implementation (see Jacobson et al. (1993)).

behind the persistent earnings loss.

The initial earnings loss is driven by wages. Upon job loss, workers are unlikely to immediately find a well-paying job through off-the-job search, as most job offers come with low wages and high separation rates (Appendix D.2.4). On-the-job search then allows workers to partially recover by climbing the job ladder toward better-paying jobs. However, the speed of convergence declines substantially after the first five years.

The evolution of human capital reveals an additional source of persistence. Job loss has a negative effect on human capital accumulation that builds up dynamically over time. Two forces explain this divergence. First, workers cannot accumulate human capital while non-employed immediately after job loss. Second, new jobs are on average less stable when workers restart the job ladder, so displaced workers spend more time in non-employment in the future, which further limits opportunities for human capital accumulation. Still, in line with Stevens (1997) and Jung and Kuhn (2018), lower wages account for the largest part of the long-run earnings losses of the average worker.

Figure D.13b turns to the consequences of job loss for consumption and wealth. Consumption falls sharply on impact, by roughly 10% and remains permanently lower. These dynamics are consistent with the permanent-income logic underlying the model. Upon job loss, workers anticipate that they have entered an employment trajectory with lower and more volatile income. Income is persistently lower because displaced workers move down the job ladder and face lower employment rates, and income is more volatile because subsequent jobs are less stable. As a result, workers reduce consumption permanently and increase precautionary saving to smooth consumption against future income risk.


On impact, consumption falls by less than income. The difference is absorbed by wealth, as workers use accumulated assets to smooth the transition to a lower permanent-income path (Kuhn, 2013). Four years after job loss, wealth stabilizes around 25% below the level of non-displaced workers and remains persistently lower, in line with the empirical findings in Barnette (2020). Three forces account for the lower average wealth after displacement. First, lower income mechanically reduces the level of wealth that households accumulate. Second, job loss flattens the remaining life-cycle income profile, reducing the need for life-cycle saving aimed at smoothing age-varying income over time. Third, lower employment stability after displacement moves workers into the Sisyphus cycle of wealth accumulation discussed in Section 4.2: workers rebuild precautionary savings while employed but repeatedly draw them down during subsequent non-employment spells.

E Model extensions

E.1 A model without heterogeneity in employment stability

This section provides additional details on the model version without heterogeneity in employment stability discussed in the main text. The general structure of the model remains unchanged relative to the baseline specification in Section 3, except that all workers of a given age j now face the same exogenous, age-dependent 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 average separation rates is identical in the models with and without heterogeneity in employment stability. Job offers drawn from the job-offer distribution continue to differ in wages, but no longer differ in separation risk.

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.

This alternative specification performs substantially worse than the baseline model in matching the dispersion of employer tenure over the life cycle. When all workers of a given age 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. The reason is that representative-worker employment dynamics translate high average mobility into low average employment stability for all workers. At age 40, the average job in the model lasts about three years, less than half of its empirical counterpart. This shortfall reflects the absence of lifetime jobs and other highly stable employment relationships in the model.

More broadly, the failure of this specification highlights that heterogeneity in employment stability is essential for jointly generating persistent employment relationships, high average labor market mobility, and the dispersion in labor market dynamics that underlies our model’s earnings and wealth implications. This point is particularly important for understanding job displacement. A model without heterogeneity in employment stability fails to replicate the large and persistent earnings losses following displacement documented in the seminal work by Jacobson et al. (1993) and the large subsequent literature. Relatedly, the discussion of frictional wage dispersion in Section D.3.3 shows why models with on-the-job search and homogeneous separation rates, as in Hornstein et al. (2011), can generate wage dispersion but struggle to jointly match persistent employment relationships and realistic earnings dynamics.

E.2 A model with dynamic worker types

This section provides additional details on the model with dynamic worker types. In this extension of the baseline model, the separation rate depends not only on the job, as in the baseline specification, but also on the worker’s human capital. Specifically, we assume

$$\lambda(h) = \lambda_l \cdot \max\{(2.5 - 1.8h + 0.3h^2), 0.5\},$$

where λ_l is drawn from the baseline distribution $\{\lambda_l\}_{l=1}^L$ and captures the job component of separation risk, while h denotes individual human capital and captures the worker component. We follow Jung and Kuhn (2018) in calibrating the relative importance of the worker and job components based on life-cycle profiles of separation rates. Here, we report the calibrated functional form for $\lambda(h)$. The calibrated parameters imply that, over the relevant range, higher human capital reduces separation risk and thereby reinforces employment stability over the life cycle.

This modification affects only the employed worker’s value function in equation (2), which becomes

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

so that the probability of transitioning into non-employment now depends on both the job component, λ_l , and the worker’s human capital. All other value functions in Section 3 remain unchanged.