Three Essays on the Macroeconomic Consequences of Labor Market Risk

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Leanne Nam

aus Seoul, Korea

2025

Dekan:	Prof. Dr. Martin Böse
Erstreferent:	Prof. Dr. Moritz Kuhn
Zweitreferent:	Prof. Dr. Thomas Hintermaier
Tag der mündlichen Prüfung:	19. September 2025

Acknowledgements

The completion of this dissertation owes much to the support and encouragement I received from many individuals throughout this journey. I am deeply indebted to those who contributed in various ways and, while it is not possible to name everyone, I am sincerely thankful for their assistance.

First, I would like to thank my advisors Moritz Kuhn, Thomas Hintermaier, and Christian Bayer for their invaluable advice, guidance, and unwavering encouragement throughout my doctoral studies. Their insightful feedback and steady encouragement have played an important role in shaping this dissertation and my academic development. I am especially grateful for the support I received at key moments in my academic journey, which has been crucial for my progress.

I am also sincerely grateful for the generous support that I received from the Bonn Graduate School of Economics, the Institute for Macroeconomics and Econometrics, the DFG Research Training Group 2281 »Macroeconomics of Inequality«, and the Collaborative Research Center TR 224. Their support created an excellent research environment that greatly enriched my experience during my doctoral studies. I am deeply grateful to all those who contributed to making this journey possible, among others, Kim Schneider, Beate Roters, Britta Altenburg, Heide Baumung, and Silke Kinzig.

One of the most rewarding aspects of this journey was sharing it with a group of wonderful friends and colleagues, among others, Johannes, Gero, Lennard, and Max, along with many others I cannot name individually here. The long hours we spent together in the first two years of the PhD, working together on problem sets, preparing exams, and supporting each other through the early stages of research, were intense and demanding at the time, but they have since become some of my most valuable and fond memories.

Finally, and most importantly, I am deeply grateful to my family. Their unconditional support and encouragement have been a constant source of strength throughout this journey. Without their support, none of this would have been possible.

Leanne Nam

Bonn, May 2025

Contents

List of Figures

List of Tables

Introduction

This dissertation consists of three self-contained essays that examine the macroeconomic consequences of labor market risk. The first two chapters focus on heterogeneity in labor market risk and its consequences on life-cycle earnings dynamics, wealth accumulation, and the design of optimal pension systems. Chapter 1 shows that heterogeneity in employment stability and the resulting earnings dynamics are a key driver of household saving behavior and wealth accumulation. In Chapter 2, I show that accounting for heterogeneity in employment stability makes the optimal pension system more progressive relative to the current U.S. system. Chapter 3 investigates how labor market shocks affecting the working-age population spill over to older households through the housing market.

In Chapter 1, which is joint work with Moritz Kuhn and Gašper Ploj, we explore how heterogeneity in employment stability shapes earnings dynamics and savings behavior of workers. Such heterogeneity is a salient feature of labor markets where some workers hold stable, lifetime jobs, while others cycle frequently in and out of employment. Despite its prevalence, the economic consequences of this labor market feature remain underexplored. We fill this gap using both empirical analysis and quantitative modeling. We first document two new empirical facts: (1) workers with more stable employment experience faster earnings growth, and (2) they accumulate more wealth per dollar of income. Second, we develop a life-cycle model that combines a labor market search model with an incomplete markets model of saving behavior. We use the model to interpret the empirical facts, to establish a causal link from employment stability to wealth accumulation, and to quantify the drivers of the underlying mechanism. Our results support an important role of labor market heterogeneity to account for differences in wealth accumulation in financial markets. Chapter 2 explores the policy implications of heterogeneity in employment stability. Employment stability has important consequences for individual welfare as interrupted work histories and the consequent earnings losses lower pension entitlements and hinder the accumulation of life-cycle savings. This chapter studies how a progressive pension system can optimally account for such heterogeneity and quantifies the resulting welfare gains. Higher pension progressivity offers insurance against employment instability and helps to reduce lifetime earnings inequality, but comes at the cost of distorting human capital investment and retirement decisions. Using a life-cycle model with heterogeneous employment stability, endogenous human capital accumulation, and retirement decision, I find that abolishing the Social Security cap and increasing progressivity relative to the current U.S. pension system is optimal. The optimal pension system leads to a welfare gain equivalent to 0.75 % of lifetime consumption for new labor market entrants. Following the observed macroeconomic shift in the labor market towards higher employment stability since the 1990s, the optimal pension system becomes less progressive, yet still delivers substantial welfare improvements. Chapter 3 studies the consequences of labor market shocks on the housing market. Unemployment leads to large and persistent income losses for workers. Higher unemployment in the labor market therefore has spillover effects on the housing market. This paper studies such spillover effects from both empirical and theoretical perspectives. Using data from the Current Population Survey (CPS), I show that a 1 percentage point increase in the unemployment rate leads to a 1.55% decline in housing prices. Theoretically, I develop an overlapping generations model with a housing market. The calibrated model replicates the empirically observed spillover effect for the U.S. economy. Higher income uncertainty is the main driver of the spillover effect during periods of high unemployment, rather than the actual income losses. The spillover effect transmits one-third of the welfare losses of workers due to higher unemployment in the labor market to older, retired households by reducing their housing wealth. Younger workers benefit in part by buying houses at depressed prices. The magnitude of the spillover effect is shaped by the demographic structure of the population and the specific age groups affected by unemployment shocks. I find that increasing the generosity of unemployment insurance stabilizes the housing market, although it only partially mitigates the spillover effect.

Chapter 1

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

Joint with Moritz Kuhn and Gašper Ploj

1.1 Introduction

Labor markets are characterized by large heterogeneity in employment stability. Some workers hold lifetime jobs, while others cycle repeatedly in and out of unstable employment (Hall, 1982). The consequences of this salient feature of labor markets remain a largely open question today. We address this question empirically and in theory. Guided by novel empirical evidence, we develop a model with heterogeneity in employment stability and consumption-saving choice to explore the labor market, financial, and welfare consequences of such heterogeneity in employment stability. We explore the life-cycle consequences of early-career heterogeneity in employment stability on earnings and wealth.

The paper provides an empirical and theoretical contribution. First, we document the large heterogeneity in employment stability using the *Panel Study of Income Dynamics* (PSID). We provide empirical evidence for a systematic relationship between differences in employment stability, earnings growth, and accumulated

* We thank Christian Bayer, Thomas Hintermaier, Philip Jung, Per Krusell, Iourii Manovskii, Kurt Mitman, Giuseppe Moscarini, Aysegul Sahin, and Abigail Wozniak for very helpful comments and remarks. We also thank seminar participants at the Federal Reserve Bank of San Francisco, the universities in Manchester, Liverpool, Erasmus University Rotterdam, Berkeley, Zurich, and at the IIES in Stockholm, the summer school *The Macroeconomics of Labor Markets*, and the *ECONtribute* Macro Workshop 2020 for their comments and remarks to improve the paper. Financial supports from DFG (No. 433368816), the DAAD, the Bonn Graduate School of Economics, and the DFG research training group 2281 "The Macroeconomics of Inequality" are gratefully acknowledged. This paper should not be reported as representing the views of the Banka Slovenije. The views expressed are those of the author.

wealth. We document that households with more stable employment have higher earnings growth and accumulate, controlling for income, more wealth. We document strongly divergent savings patterns throughout the life cycle between workers starting careers with stable relative to unstable employment. By the end of their prime-age working lives, workers with a stable career have 20% higher labor earnings and have accumulated additional wealth equivalent to one and a half years of income.

The second contribution of the paper is theoretical. We incorporate a frictional labor market model with heterogeneity in employment stability into an otherwise standard life-cycle model of consumption-saving behavior. We demonstrate that this model is jointly consistent with life-cycle labor market, earnings, and wealth dynamics. The model accounts for the empirical relationship between employment instability, earnings growth, and wealth accumulation. In particular, it provides a quantitatively consistent explanation for the empirical differences in life-cycle income and wealth accumulation between workers with stable and unstable careers. Given the empirical success of the model, we use it to explore the economic consequences of heterogeneity in job stability. We find that differences in wage and human capital levels each account for around half of the difference in labor earnings. The difference in earnings growth accounts for 60% of the difference in the wealth-to-income ratio between workers with stable and unstable employment. While workers with unstable employment save more for precautionary motives, the realization of nonemployment leads to a dissaving of wealth to smooth consumption. This effect accounts for the remaining 40% of the difference in wealth-to-income ratio.

Our model combines a life-cycle labor search model with heterogeneity in employment stability and human capital investment with a consumption-saving model with incomplete financial markets. We interpret the heterogeneity in employment stability as arising from heterogeneity in job stability in line with empirical and structural literature (for example, Hall, 1982; Jarosch, 2015; Pinheiro and Visschers, 2015; Jung and Kuhn, 2018; Larkin, 2019). In the labor market, workers search on and off the job and jobs are heterogeneous with respect to wages and separation rates to nonemployment (Jarosch, 2015; Pinheiro and Visschers, 2015). Separation rate differences are one determinant of differences in job duration; on-the-job search with workers climbing the job ladder constitutes a second source for differences in employment duration. Human capital investment opportunities exist only for employed workers who can exert effort to accumulate human capital. Thus, unstable careers with low employment rates perpetuate low incomes by offering fewer opportunities for human capital investment. Earnings growth in the model is endogenous and we find that over the working life, employment instability leads to lower earnings growth in line with the empirical evidence. The endogenous earnings process therefore provides a tight link between life-cycle earnings growth and income risk that is absent in most macroeconomic models of life-cycle savings behavior (Kaplan and Violante, 2010). The consumption-saving part of the model is standard, with agents facing incomplete financial markets where they save in a risk-free asset subject to a no-borrowing constraint. Life-cycle variation in incomes, in combination with incomplete financial markets, provides agents with a life-cycle and precautionary savings motive. We study the model in partial equilibrium and take job offer rates and interest rates as given.

When we bring the model to the data, we estimate model parameters to jointly match life-cycle labor dynamics, earnings growth, and wealth-to-income ratios for the U.S. economy. The model also matches untargeted empirical facts on consumption inequality (Aguiar and Hurst, 2013), earnings dynamics (Topel and Ward, 1992; Blundell, Pistaferri, and Preston, 2008), earnings losses following job displacement (Jacobson, LaLonde, and Sullivan, 1993), and the distribution of earnings growth (Guvenen et al., 2019). The novel facts that our model speaks to are the differences in earnings growth and wealth accumulation dynamics from the PSID data across workers with different employment stability. We show that the model matches the novel empirical fact on the relationship between employment stability, earnings growth, and wealth accumulation, so that we can interpret this empirical correlation through the lens of the model.[1]

To explore the individual consequences of heterogeneity in employment stability, we first decompose life-cycle earnings and wealth accumulation. The life cycle is an important dimension of heterogeneity in employment stability. Young workers look for stable and high-paying jobs so that when old, the average worker has found a stable and well-paying job. In our model, the combination of low employment stability and low income when young creates an additional tension between the precautionary and life-cycle savings motive for young workers in unstable jobs.

Compared to workers with unstable employment at age 50 (top 25% of nonemployment duration), a worker with stable employment (bottom 75% of nonemployment duration) has 20% higher labor earnings at the age of 54. We decompose the earnings difference into the wage and human capital component. The differences in wage and human capital almost equally account for the difference in labor earnings for a worker at age 54. Workers with low employment stability also have lower employment rates. Lower employment rates offer fewer opportunities for human capital investment and, taken together, result in incomes that are lower and more volatile. Put differently, unstable jobs are dead-end jobs with low income today, offer few opportunities for career development, and carry a high risk of job loss. By contrast, workers who find a stable lifetime job invest in their careers, enjoy their growing incomes, and face little risk of job loss. This model prediction aligns closely with the empirical evidence in Guvenen et al. (2019), who emphasize the importance of

1. We also compare the model to data from the SCF and the joint distribution of income and wealth.

heterogeneity in nonemployment in accounting for life-cycle earnings dynamics in the U.S. data. Whereas most job offers are dead-end jobs that create high average labor market mobility for the macroeconomy, labor market mobility and dynamics along the job ladder lead to an increasing share of workers in stable employment relationships over the life cycle.

To explain the difference in wealth-to-income ratios between workers with stable and unstable employment, we construct counterfactual models. We decompose the difference in wealth-to-income ratio into three components: "precautionary saving," "income growth," and "nonemployment." Workers with higher employment risk accumulate more precautionary savings in order to smooth consumption in case of nonemployment, which is the "precautionary saving" component. The "income growth" component arises as workers with stable employment earn higher labor earnings and therefore, they accumulate more assets due to the life-cycle savings motive. Finally, the "nonemployment" component captures the difference in wealth due to workers drawing their assets to smooth consumption when they become nonemployed.

Large and persistent earnings losses after job loss are an important source of labor market risk, and heterogeneity in job stability is the crucial model ingredient to account for such earnings losses in structural labor market models (Jarosch, 2015; Jung and Kuhn, 2018). We rely on our model framework to explore the consequences for consumption-saving dynamics and how heterogeneity in previous job stability shapes the consequences of job loss. We corroborate large and persistent earnings losses from job loss for the average worker and also find, in line with the permanent income hypothesis, persistent drops in consumption. After job loss, incomes recover during a transition to their new permanent level, but consumption remains insulated from these transitional dynamics as consumption is smoothed by running down wealth. With respect to employment stability in the lost job, we find large heterogeneity in earnings losses. Losing an unstable job leads to large but transitory earnings losses, and a buffer stock of wealth insulates consumption from these transitory earnings fluctuations. By contrast, the loss of a stable job leads to very large and persistent earnings losses that translate into persistently lower consumption. This heterogeneity suggests that at the macroeconomic level, the composition of job losses from stable and unstable jobs is a key determinant of aggregate consumption dynamics and that abstracting from heterogeneity in job stability potentially severely underestimates the consumption drop from job losses. We demonstrate that consumption dynamics absent job-stability heterogeneity align closely to the dynamics after the loss of an unstable job.

The following section relates our work to the existing literature. In Section 1.2, we provide empirical evidence on job stability and wealth accumulation, employment inequality, and the sources of heterogeneity in job stability. We present the model in Section 1.3. Section 1.4 explores the consequences of heterogeneity in employment stability on income and wealth. In Section 1.5, we study the consequences of job loss. Section 1.6 concludes.

Our work relates to two large strands of literature: models of consumption-saving behavior in the presence of idiosyncratic income risk and market incompleteness (Huggett, 1993; Aiyagari, 1994; Bewley, undated) 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 wages as 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 (Krusell, Mukoyama, and Sahin, 2010; Lise, 2012; Krusell et al., 2017; Hubmer, 2018; Larkin, 2019; Cajner, Güner, and Mukoyama, 2020; Kaas, Lalé, and Siassi, 2023). We add to this literature by exploring the consequences of heterogeneity in employment stability. Our paper connects the part of the literature that focuses on macroeconomic dynamics, as in Krusell, Mukoyama, and Sahin (2010), with microeconomic behavior, as in Lise (2012).

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

Our labor market model builds on Jung and Kuhn (2018), who develop a life-cycle search model to demonstrate that heterogeneity in job stability is key to account for earnings losses following job displacement. Jarosch (2015) also highlights the importance of heterogeneity in job stability to account for observed earnings losses. While heterogeneity in job stability arises as a bargaining outcome between employers and workers in Jung and Kuhn (2018), we follow Pinheiro and Visschers (2015) and Jarosch (2015) and introduce this heterogeneity in reduced form to the job-offer distribution.[2] Guvenen et al. (2019) explore life-cycle earnings dynamics and document large heterogeneity in life-cycle nonemployment spells. They emphasize that this heterogeneity is key to account for life-cycle earnings dynamics. Additional evidence for heterogeneity in job stability comes from Morchio (2020), who documents large heterogeneity in unemployment within cohorts of U.S. workers.

Our work also relates to research on heterogeneity in earnings risk, as in Low, Meghir, and Pistaferri (2010), Karahan and Ozkan (2013), and Caplin et al. (2023). Low, Meghir, and Pistaferri (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.

1.2 Heterogeneity in employment stability and wealth accumulation in the data

In this section, we provide an empirical analysis that establishes large heterogeneity in employment stability in the United States and a tight correlation with wealth accumulation. This empirical evidence serves as the motivation and will guide our model building in the next step. The empirical analysis relies on data from the *Panel Study of Income Dynamics* (PSID) and consists of two steps. In the first step, we show empirical evidence for large heterogeneity in employment stability in the U.S. labor market. In the second step, we document the correlation between employment stability, earnings growth, and wealth accumulation. In Appendix 1.A.6, we provide corroborating evidence for both steps based on data from the *Survey of Consumer Finances* (SCF).

2. Heterogeneity in job stability across regional labor markets has recently been highlighted in Bilal (2019) as the main driver of spatial unemployment rate differences.

The PSID is the longest-running household panel for the United States beginning in 1968. The surveys were conducted annually until 1997 and biennially thereafter. Questions on household wealth have been available since 1985 at five-year intervals and biennially since 1999. Exploiting the panel dimension of the PSID, we track the labor market status and the financial situation of households over time. We restrict the sample to household heads between age 25 and 55 and drop individuals who are permanently out of labor force.

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 nonemployment 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.

1.2.1 Heterogeneity in employment stability

The seminal paper by Hall (1982) focuses on heterogeneity in employment stability stemming from heterogeneity in job stability. He documents large heterogeneity in job stability and the existence of lifetime jobs in the U.S. labor market.[3] Guvenen et al. (2019) explore life-cycle earnings dynamics in high-quality Social Security data for the United States and document large heterogeneity in nonemployment over the life cycle. They emphasize the importance of incorporating heterogeneity in nonemployment risk to account for observed life-cycle earnings dynamics. Recent work by Morchio (2020) relies on panel data from the 1979 National Longitudinal Survey of Youth (NLSY) and also documents large differences in separation rates across over the life cycle. The key advantage of the PSID data is that we can jointly study employment stability, income, and wealth over the life cycle.

For employment stability, we document in Figure 1.1 large inequality in the accumulated life-cycle nonemployment duration across workers. Figure 1.1a shows the distribution of nonemployment duration at age 50 that workers have accumulated between ages 25 and 50. Strikingly, we find that two out of ten workers experience no spell of nonemployment over their prime-age working life. Around half of all workers accumulate a nonemployment duration between 1 and 10 quarters, whereas the remaining 30% of workers have accumulated nonemployment duration of more than 10 quarters. Such a nonemployment duration of 10 quarters corresponds to a nonemployment rate of 8 percent.

3. The strand of the literature that studies large and persistent earnings losses in structural models also emphasizes heterogeneity in job stability as a salient feature of the labor market (Jarosch, 2015; Jung and Kuhn, 2018).

Figure 1.1. Histogram and Lorenz curve of nonemployment duration

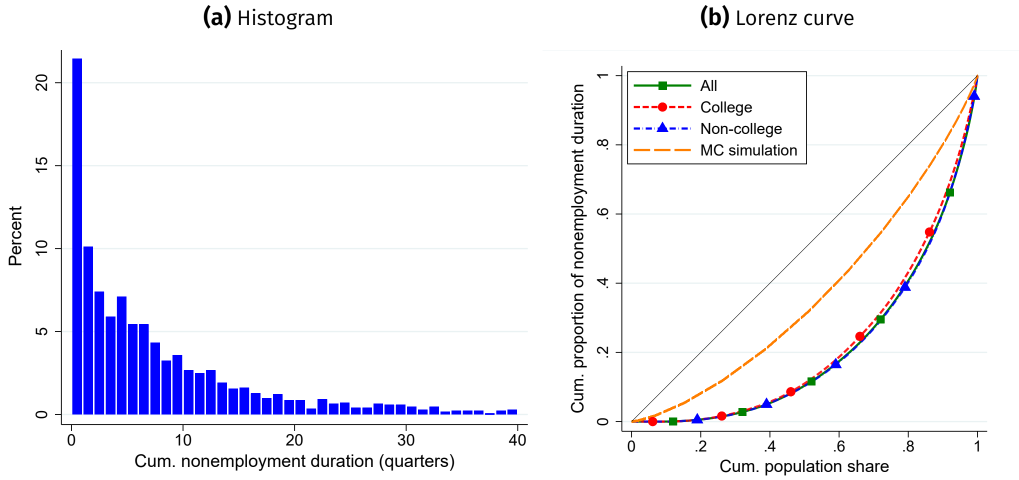

Notes: Panel (a) shows the histogram of nonemployment duration (in quarters). Panel (b) shows the Lorenz curve of nonemployment duration for all, college-educated, and non-college households, and a Monte Carlo simulation where all workers have the average age-dependent labor market transition rate (separation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

Figure 1.1b summarizes how unequally the nonemployment duration is distributed across workers by showing the Lorenz curves for nonemployment duration from Figure 1.1a. Looking first at all workers (green line with squares), we find that around 30% of workers account for more than 70% total nonemployment duration. Hence, labor market mobility is concentrated in a small group of workers. To contrast this inequality to a situation with homogeneous employment stability, we show as dashed yellow line the Lorenz curve from a Monte Carlo simulation that abstracts from heterogeneity in employment stability. The Lorenz curve without heterogeneity is much closer to the identity line. In this case, 30% of workers account for only 40% of the accumulated nonemployment duration so the distribution of nonemployment is much more equally distributed across workers. Besides heterogeneity in job stability, worker heterogeneity could be another important driver of such heterogeneity in nonemployment episodes. The arguably most important dimension of worker heterogeneity is educational attainment. Figure 1.1b shows as dotted red line with circles the Lorenz curve for nonemployment duration of college graduates, and as dashed blue line with triangles, the corresponding Lorenz curve for non-college workers. Strikingly, we find that despite very different average transition rates between these groups, the within-group inequality in nonemployment duration is virtually identical. Hence, even within two groups of workers with arguably very different average employment stability, we find the same amount of inequality in employment stability. In our structural model, we will therefore fo-

Figure 1.2. Tenure over the life cycle

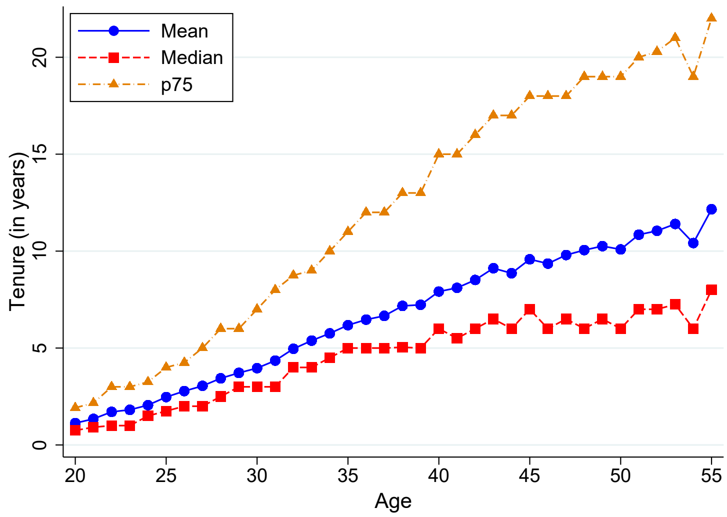

Notes: The figure shows the life-cycle evolution of the cross-sectional distribution of tenure (in years). Data are from the Panel Study of Income Dynamics.

cus on employment stability stemming from heterogeneity in job stability that is consistent with large within-worker-group heterogeneity in job stability.

Another approach to document heterogeneity in employment stability is to look at the tenure distribution (Hall, 1982; Jung and Kuhn, 2018). Figure 1.2 shows life-cycle profiles for tenure in the PSID data.[4] 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 the finding from Figure 1.1 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 nonemployment over their entire prime-age working life. Looking at the tenure distribution at age 60, we find that more than 50 percent of workers have been with their employer for 10 years, and almost a quarter of workers at age 60 have been at the same employer for at least 25 years.

1.2.2 Earnings dynamics

What are the implications of the observed heterogeneity in employment stability on earnings and wealth accumulation of workers? To address this question, we leverage

4. Tenure is widely observed in cross sectional datasets and provides therefore an attractive variable to explore the relationship between wealth accumulation and employment stability. In Appendix 1.A.10, we document virtually the same tenure profiles in CPS and SCF data.

Figure 1.3. Log labor income

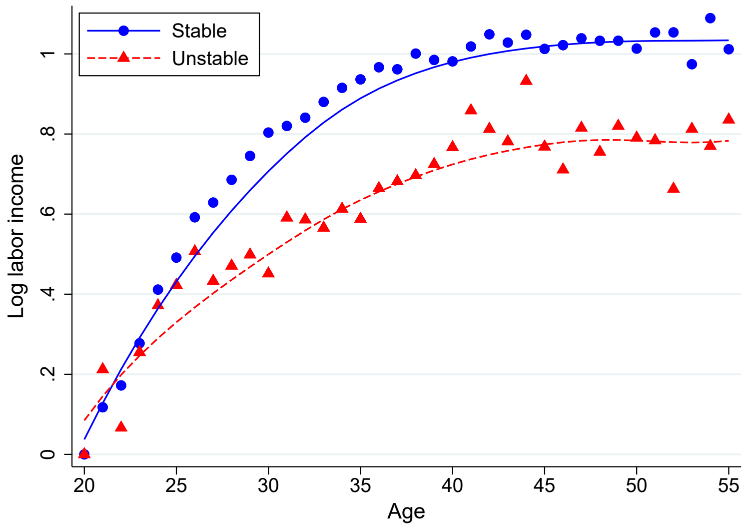

Notes: The figure shows the life-cycle profiles of log labor income for stable and unstable groups of workers in terms of employment stability. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of nonemployment duration at age 50. The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

the strength of the PSID data that it tracks workers' labor market history, earnings, and wealth accumulation jointly over time. In the first step, we examine the relationship between nonemployment and labor income for workers with stable and unstable employment histories. The unstable group consists of the 25% of workers at age 50 with the longest cumulative nonemployment duration. The remaining workers are assigned to the stable group. This sorting conditions on the ex-post realization of the employment history. We document however that results are robust to sorting workers ex ante at age 30. Consistent with the robustness of the results to the sorting age is the finding that there is high group stability. Only 13% of workers from the stable group at age 30 end up in the unstable group at age 50.

Figure 1.3 shows the life-cycle profiles of log labor income for workers with stable and unstable employment where we have removed initial differences to focus on the differences in income growth. Strikingly, we observe a large difference in earnings growth over the life cycle. At age 50, the average labor income of the stable group has grown more by 20 log points compared to the group of workers with unstable employment history. This result speaks against the common practice in large parts of the macroeconomic literature that considers the average life-cycle earnings growth and earnings risk as independent. Our results suggest that being nonemployed does not only lead to an income loss during the nonemployment spell, but that longer unemployment duration also leads to lower income growth over the life cycle. What we cannot rule out is that a repeated permanent negative shocks from job loss leads to lower earnings over the life cycle. What speaks against this interpretation is that the empirical literature finds the large and persistent shocks for long-term employed workers from stable employment relationships (Jacobson, LaLonde, and Sullivan, 1993; Davis et al., 2011).

1.2.3 Wealth accumulation

How do these differences in labor market experiences affect workers saving behavior? The PSID data allow us to trace wealth accumulation of workers with different employment trajectories over time. In Figure 1.4, we consider the relationship between the accumulated nonemployment duration from Figure 1.1b and wealth-to-income ratios and nonemployment. Looking at wealth-to-income ratios already removes the effect of lower income on wealth accumulation that we have documented in Figure 1.3. For the nonemployment duration, we look at the accumulated duration at age 50. Figure 1.4a shows a strong negative relationship with wealth-to-income ratios. Qualitatively, a negative slope implies that per dollar of income, workers in more stable employment have more wealth, or, in short, workers with more stable employment history are wealthier. Quantitatively, the observed slope is economically meaningful. The slope implies that being nonemployed for four years will, on average, lead to additional wealth corresponding to roughly one year of income.[5] One concern might be other confounding factors from fixed worker differences, most prominently, education. In Appendix 1.A.5, we repeat the analysis with additional controls for education and industry. We find that the documented relationship is both qualitatively and quantitatively robust.

Through the lens of economic theory, repeated spells of nonemployment lead workers to deplete their savings to smooth income shocks, while lower earnings growth over the life cycle (Figure 1.3) weakens the life-cycle savings motive. Hence, we should expect less wealth accumulation over the life cycle for workers with unstable employment histories. Figure 1.4b shows the life-cycle wealth-to-income ratio profiles for workers with stable and unstable employment histories as in Figure 1.3. Consistent with the prediction from theory for differences in saving behavior, we find a persistent gap in wealth-to-income ratios for workers with stable and unstable jobs over the life cycle. At age 50, workers with more stable careers are wealthier by one and a half years of income compared to workers with unstable employment. The differences in wealth accumulation are apparent over the entire life cycle starting at age 25 consistent with the idea that differences in employment stability are very persistent over the working life and that workers in stable careers typically do

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

Figure 1.4. Wealth-to-income ratios and nonemployment

(a) Binned scatter **(b)** Life cycle

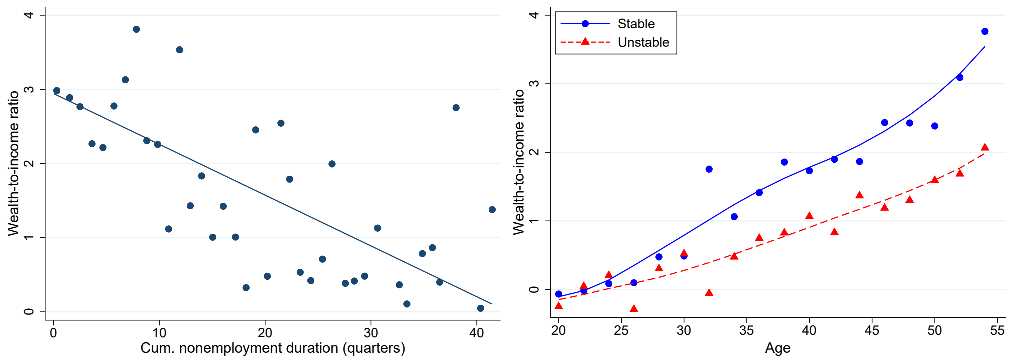

Notes: Panel (a) shows a binned scatter plot of wealth-to-income ratios against nonemployment duration of workers at age 50. Panel (b) shows the life-cycle evolution of the wealth-to-income ratios for stable and unstable groups of workers in terms of employment stability. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of nonemployment 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.

not expect losing their job. In Appendix 1.A.6, we further corroborate a positive relationship between employment stability and wealth using data from the SCF that provides information on job tenure and the number of employers that workers had over their career. All results strongly support the idea that more stable employment leads to more wealth accumulation over the life cycle.

In Table 1.1, we present the results of a regression of wealth-to-income ratio on nonemployment duration of workers. Column 1 shows that a nonemployment duration of one year (4 quarters) is associated with a 0.3 lower wealth-to-income ratio. While there is a significant negative relationship between employment stability and wealth-to-income ratio, worker heterogeneity such as education is another important determinant of differences in wealth (Bartscher, Kuhn, and Schularick, 2020). We consider education as an additional control in column 2. Higher educational attainment is associated with higher wealth-to-income ratios, but we still find a negative and significant relationship between nonemployment and wealth-to-income ratio. Even after including other worker characteristics, the coefficient of nonemployment remains negative and statistically significant.

To summarize, we document that in the data, employment stability is very heterogeneous across careers and is as large in the entire population of workers as within groups of workers with on average different labor market mobility pattern. The novel facts with respect to the consequences of this heterogeneity is that differences in employment stability are systematically related to earnings growth and wealth accumulation. The differences are economically significant. The regression

Table 1.1. Wealth-to-income ratio and nonemployment

Wealth-to-income ratio	(1)	(2)	(3)	(4)
Nonemployment (in quarters)	−0.07***	−0.06***	−0.04***	−0.03**
	(0.01)	(0.01)	(0.01)	(0.01)
Education		1.21***	0.92***	1.05***
		(0.31)	(0.01)	(0.01)
Nonemployment x education		−0.03		−0.02
		(0.03)		(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 wealth-to-income ratio. The sample consists of workers at age 50. Column (1) shows the regression results of wealth-to-income ratio on nonemployment duration of workers in quarters. Column (2) includes education and the interaction of education and nonemployment 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. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

analysis shows that, while education is an important driver of wealth differences across workers, there is also a substantial within-group heterogeneity in wealth-to-income ratio among college-educated workers and workers without college education.

The next section develops a model of household saving behavior that explicitly introduces heterogeneity in employment stability. In what follows, we focus on the within-group heterogeneity in analyzing the relationship between employment stability and wealth. We interpret the source of heterogeneity as coming from heterogeneity in employment in line with a large strand of the literature (Hall, 1982; Jarosch, 2015; Pinheiro and Visschers, 2015; Jung and Kuhn, 2018; Larkin, 2019).

1.3 Heterogeneity in employment stability and wealth accumulation in theory

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, Kuhn, and Wright, 2015). Workers start their life in the working phase that lasts for T^W periods. At the end of the working phase, workers move to the transition phase that is of stochastic length with expected duration T^T. In the end, workers leave the transition phase to the retirement phase that lasts for T^R periods. In each period before the retirement phase, a worker is either employed or nonemployed. We denote the agent's employment status by ε with $\varepsilon \in \{e, n\}$ where e stands for employed and n for nonemployed. If the worker is employed, her job is characterized by a bundle (w, λ) where w denotes the wage and λ the separation rate where the wage w captures the rental rate of human capital on the current job. We discretize wages and separation rates to grids $\{w_k\}_{k=1}^{K}$ and $\{\lambda_l\}_{l=1}^{L}$ and assume that $w_k < w_{k+1}$ for all k and $\lambda_l < \lambda_{l+1}$ for all l. To economize on notation, we denote the wage-separation rate bundle at age j only by $\{w_j, \lambda_j\}$. Each worker holds assets denoted by a and a stock of human capital denoted by h. The period budget constraint is

$$a_{j+1} + c_j = (1 + r)a_j + y(w_j, h_j, \varepsilon), \tag{1.1}$$

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

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

When the worker is in the working or transition phase, we split each period into four stages: *separation*, *investment*, *production*, and *search*. At the separation stage, employed agents separate from their job with probability λ. If the agent separates, she becomes nonemployed and moves to the production stage. Employed agents who do not separate move to the investment stage where human capital investment decisions are made. At the production stage, employed agents receive earnings, the job's wage rate times the worker's stock of human capital, and nonemployed agents receive benefits proportional to earnings on their last job. At the search stage, employed and nonemployed agents receive job offers. We allow for different job-offer arrival rates on the job and in nonemployment. We take job-offer arrival rates as exogenous and denote the arrival rate on the job by π_e and the arrival rate in nonemployment by π_n. Job offers, combinations of a wage rate w and a separation probability λ, for employed and nonemployed workers are drawn from the same joint distribution $f(w, \lambda)$ (Jarosch, 2015; Pinheiro and Visschers, 2015; Larkin, 2019). An agent who receives a job offer decides to reject or accept the job offer. If the agent accepts the job offer, she will be employed at the beginning of the next period in the new job. If the agent rejects the job offer, she remains nonemployed (employed in her current job) and there is no recall of previous job offers.

Only employed workers have the opportunity to invest in their human capital. At the investment stage, the agent decides if she wants to exert effort for human capital investment. Effort provision for human capital accumulation is a choice $t \in [0, 1]$ (training). Disutility from effort enters the utility additively separable as quadratic cost κt^2. Nonemployed agents do not have the opportunity to accumulate human capital. If agents do not exert effort, their human capital stays constant at level h until the next period.[6] One interpretation of this effort provision is as career investment with the current employer (e.g., unpaid overtime, higher work intensity, on-the-job training, or committee work). We assume that human capital levels are discrete and are members of an ordered set with largest (smallest) element h^{max} (h^{min}). We denote by h^+ the immediate successor of human capital level h and by h^- the immediate predecessor of h. Human capital investment is risky. An agent at human capital level h exerting effort t to accumulate human capital has a probability $p_H(t,j)$ of reaching human capital level h^+. We allow for age dependence of $p_H(t,j)$. The law of motion for human capital if the agent exerts effort ($t > 0$) is

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

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

The consumption-saving decision is standard. The agent chooses next period's asset level given her current state and facing a borrowing constraint that prevents negative asset holdings. 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

6. Although we do not assume human capital depreciation during nonemployment, there is on average relative depreciation of human capital because employed workers invest and accumulate human capital while nonemployed workers do not.

possible continuation states. A worker in the transition phase either remains in the transition phase or transits to the retirement phase. A worker in the working phase ages deterministically and transits at the end of prime-age working life to the transition phase. We do not allow workers from the transition phase (retirement phase) to transit back to the working (transition) phase.

Transiting from the transition phase to the retirement phase is stochastic and happens with probability ψ. Upon reaching the retirement phase, workers leave the labor market and receive social security benefits. Agents do not face any labor market risk during retirement and solve a deterministic, finite-horizon consumption-saving problem.

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

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

$$V_e(a,w,\lambda,h,j) \;=\; \lambda V_n^P(a,w,h,j) + (1-\lambda)V_e^I(a,w,\lambda,h,j), \tag{1.2}$$

where V_n^P denotes the value function of a nonemployed 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 nonemployed worker at the production stage V_n^P is identical to the value function at the separation stage V_n because for already nonemployed workers, nothing happens at the separation stage.

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

$$V_e^I(a,w,\lambda,h,j) \;=\; \max_{t\in[0,1]} -\kappa t^2 + p_H(t,j)V_e^P(a,w,\lambda,h^+,j) + (1-p_H(t,j))V_e^P(a,w,\lambda,h,j). \tag{1.3}$$

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

$$\begin{aligned} V_e^P(a,w,\lambda,h,j) \;&=\; \max_{\{c,a'\geq 0\}} u_j(c) + \beta\Big(\pi_e V_e^S(a',w,\lambda,h,j) + (1-\pi_e)V_e(a',w,\lambda,h,j+1)\Big) \\ s.t. \qquad c \;&=\; (1+r)a + y(w,h,e) - a', \end{aligned} \tag{1.4}$$

where V_e^P denotes the employed agent's value function at the production stage, V_e^S denotes the employed agent's value function at the search stage, and V_e denotes the value function of an employed worker at the beginning of the period. The time discount factor is denoted by β. The first line of equation (1.4) is composed of the flow utility for the current period and the discounted expected utility from the search stage. The probability of receiving a job offer is π_e. The distribution over job offers is $f(w, \lambda)$, so that for the value function of an employed worker at the search stage, we get

$$V_e^S(a', w, \lambda, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max\Big\{ \underbrace{V_e(a', w, \lambda, h, j+1)}_{\text{staying in current job}}, \underbrace{V_e(a', w_s, \lambda_k, h, j+1)}_{\text{accepting outside offer}} \Big\} f(w_s, \lambda_k), \tag{1.5}$$

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

The value function of a nonemployed worker at the production stage is

$$\begin{aligned} V_n^P(a, w, h, j) &= \max_{\{c, a' \geq 0\}} u_j(c) + \beta\Big(\pi_n V_n^S(a', w, h, j) + (1 - \pi_n) V_n(a', w^-, h, j+1)\Big) \\ s.t. \quad & c = (1 + r)a + y(w, h, u) - a', \end{aligned} \tag{1.6}$$

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

For the value function of a nonemployed worker at the search stage, we get

$$V_n^S(a', w, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max\Big\{ \underbrace{V_n(a', w^-, h, j+1)}_{\text{staying nonemployed}}, \underbrace{V_e(a', w_s, \lambda_k, h, j+1)}_{\text{accepting job offer}} \Big\} f(w_s, \lambda_k). \tag{1.7}$$

The value function again captures the acceptance-rejection decision over job offers.

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 1.A.7.

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

$$V_r(a, w, h, j_r) = \max_{a' \geq 0} u((1 + r)a + y(w, h, n) - a') + \beta V_r(a', w, h, j_r + 1). \tag{1.8}$$

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 1.A.8.

1.3.1 Bringing the model to the data

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

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

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

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

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

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

For the functional form of the job-offer distribution $f(w, \lambda)$, we assume that the marginal distributions of wages and job stability $(1 - \lambda)$ follow a truncated exponential distribution. We consider as support for wages $\left[\underline{w}, \overline{w}\right]$ and job stability $\left[1 - \overline{\lambda}, 1 - \underline{\lambda}\right]$. We set $N_w = 5$, $\underline{w} = 1$, and $\overline{w} = 1.85$, in line with the empirical support of mean log earnings, and use equidistant grid points in logs. For job stability $1 - \lambda$, we set $N_\lambda = 10$ and set $\overline{\lambda} = 0.35$ so that the least stable job lasts for one quarter and $\underline{\lambda} = 0.006$ to represent lifetime jobs with an expected duration of 42 years. We set the remaining grid points nonlinearly between the most and least stable jobs, with more grid points toward the least stable job.[7] To parametrize the joint

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

Table 1.2. Estimated parameters

Parameter	Value	Description
β	0.994	Quarterly discount factor
κ	0.397	Utility cost of effort
π_e	0.428	Probability of a job offer when employed
π_u	0.852	Probability of a job offer when nonemployed
ψ_w	0.536	Marginal distribution of w^*
ψ_λ	0.504	Marginal distribution of $1-\lambda^*$
θ	0.491	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.076	Skill upgrading probability
ρ	0.983	Persistence of skill upgrading probability

distribution, we map both supports to the unit interval $[0,1]$ denoting by $w^* \in [0,1]$ the standardized wage and by $1-\lambda^*$ standardized job stability. The density of w^* is $f(w^*) = (1-\exp(-\psi_w))^{-1}(\psi_w \exp(-\psi_w w^*))$ where ψ_w determines the shape of the density. The density of standardized job stability $1-\lambda^*$ follows accordingly with shape parameter ψ_λ. We parametrize the correlation between the marginal distributions by constructing the joint distribution using a copula C_θ, where the value of θ determines the correlation between w^* and $1-\lambda^*$.

We determine parameters within the model using a simulated method of moments that minimizes the difference between model moments and empirical moments. For the empirical moments, we use the life-cycle profiles of (log) earnings (mean and variance), labor market transition rates, tenure (mean, median, 75th percentile), and of the wealth-to-income ratio. For labor market transition rates, we rely on estimated life-cycle profiles from Jung and Kuhn (2018) based on CPS data.[8] In Appendix 1.A.8, we provide further details on the estimation implementation and an intuitive discussion on how the empirical profiles identify the free model parameters. We abstain from a formal proof of identification. Table 1.2 presents the model parameters together with their estimated values.

The value of the quarterly discount factor β corresponds to an annualized value of 0.97, which is well within the range of conventional values in the macroeconomic literature. The utility cost parameter κ implies average utility costs measured as life-

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

time consumption-equivalent variation between 0.35 percent during the first ten years of working life and less than one-tenth of a percent during the last ten years of working life.[9] For labor market parameters, we get that job offer probabilities in nonemployment π_u have to be roughly twice as high than in employment π_e, to match the high quarterly job-finding rates, but even during employment, workers frequently get job offers. Such a difference between contact rates is qualitatively and quantitatively consistent with the calibration in Hornstein, Krusell, and Violante (2011) for the United States. The shape parameters of the marginal distributions, ψ_w and ψ_λ, determine the relative frequency of the different wage and job stability levels in the offer distribution. They imply that one-third of job offers come with the lowest wage, and less than one out of twelve job offers come with the highest wage level. For job stability, we get that less than one out of 20 jobs are the most stable lifetime jobs, whereas almost one out of six job offers are of the least stable type, lasting in expectation for one quarter only. The copula parameter θ implies a positive correlation between wages and job stability. If wages and job stability were independent, the probability of the least-stable lowest-paying job would be 5.1 percent, but given θ, it is almost 50 percent higher with 7.4 percent. For the most-stable best-paying job, the offer probability is 0.9 percent, highlighting that stable and high-paying jobs are hard to find. Such a correlation between wages and job stability is also in line with the empirical evidence in Jung and Kuhn (2018) that high-wage jobs are more stable (lower separation rates). Figure 1.A.10a shows the estimated joint job-offer distribution over wages and separation rates and the marginal distributions of separation rates at different wage levels. The joint distribution is clearly asymmetric, with most of the probability mass concentrated at low-wage, unstable jobs. Figure 1.A.10b shows that the conditional marginal distribution of separation rates in low-wage jobs always first-order stochastically dominates the distribution of separation rates in high-wage jobs. The parameter $\bar{p}_H$ for the human capital process implies that for a labor market entrant maximum effort provision during the first year ($t = 1$) yields a 20 percent probability of career progression (human capital increase). The decay in the human capital investment technology ρ implies that after 10 years in the labor market, the same effort provision will imply a 11 percent probability of career progression. While all parameters appear economically reasonable in isolation, we will now demonstrate that they yield a close fit between the model and data along targeted and untargeted dimensions.

1.3.2 Theory meets evidence

In this section, we first demonstrate the model's ability to account for average life-cycle profiles of labor market mobility, tenure, earnings, and wealth accumulation.

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

Figure 1.5. Transition rates and tenure

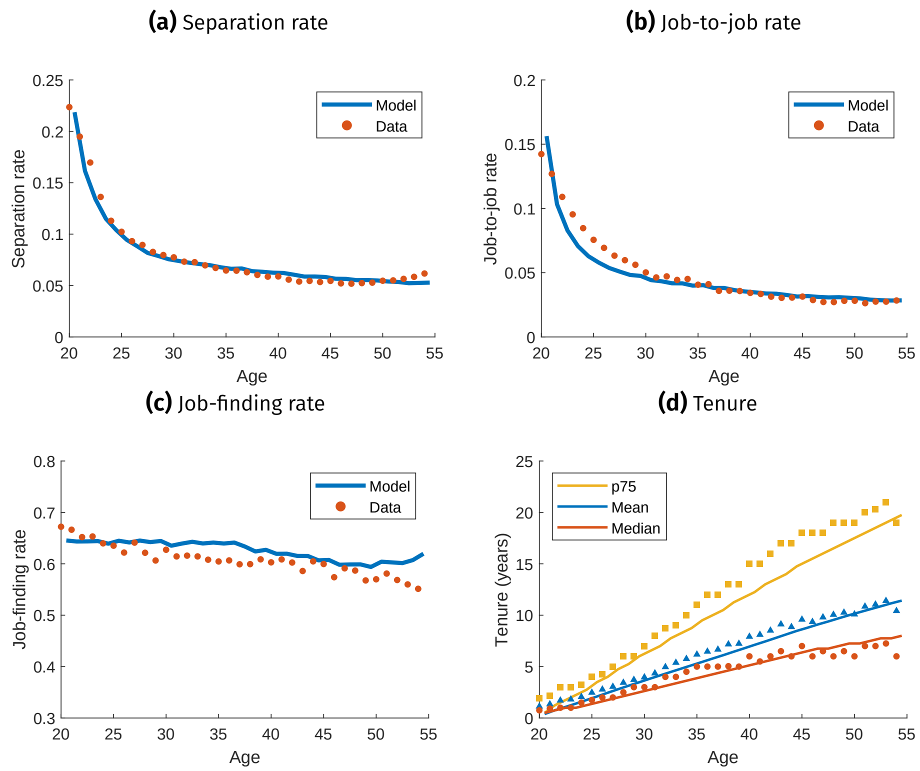

Notes: This figure shows quarterly life-cycle transition rates and tenure in years by age. The dots show the empirical profiles, while the solid lines show the corresponding model profiles. Empirical transition rates are computed using data from the CPS. The empirical tenure profiles are computed using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

Second, we demonstrate that the model is also consistent with life-cycle patterns of consumption, earnings, and employment inequality. Finally, we discuss that the model also compares favorably to the data regarding individual-level dynamics by looking at earnings and wealth mobility and the joint distribution of income and wealth.

Figure 1.5 shows the empirical life-cycle profiles for separation, job-to-job, and job-finding rates and their model counterparts. Looking at the separation rate in Figure 1.5a, we see that its evolution is matched very closely. Model and data show a strong decline up to age 30 and constantly falling separation rates between ages 30 and 50. Figure 1.5b shows that the model also matches the life cycle of job-to-job rates very well, with only a slightly steeper decline of job-to-job rates between ages 20 and 30 compared to the data. The model mechanism to match these declining life-cycle profiles consists of workers climbing the job ladder and finding more stable and better-paying jobs. Job-finding rates in Figure 1.5c are matched well in level and trend and generally show life-cycle variation. Finally, Figure 1.5d shows the life-cycle profiles of mean, median, and the 75th percentile of the tenure distribution. The model closely matches the empirical increase and heterogeneity in job stability. Importantly, this demonstrates that the model is jointly consistent with high average transition rates (Figures 1.5a and 1.5b) and high job stability for most workers (Figure 1.5d).[10] Appendix Figure 1.A.13 shows the cross-sectional distributions of employer tenure and the number of employers over a worker's career. The model compares very favorably to the data for both distributions. In particular, it accounts for a large fraction of short-term jobs but also with the substantial share of jobs with more than 10, 20, and even 30 years of tenure.[11]

Figure 1.6 turns to the life-cycle dynamics of earnings and wealth. Looking at the life-cycle profile of mean log earnings in Figure 1.6a, we find that the model matches the steep increase in earnings after labor market entry and the flattening out after age 40. It closely matches the large average increase of roughly one log points over the life cycle but shows slightly less concavity in comparison to its empirical counterpart. Figure 1.6b shows the life-cycle profile of the wealth-to-income ratio as a measure of wealth accumulation. Again, we find a close fit between model and data. Wealth-to-income ratios in model and data rise from slightly above 0 at age 20 to approximately 4 at age 55.In the Appendix 1.A.9, we also show that the model is able to match the increase in the variance of log earnings by 0.2 over the life cycle. The model also produces an increase in consumption inequality that is in line with the empirical profiles.

We have demonstrated that the model's endogenous earnings and consumption-saving dynamics match the average life-cycle earnings and wealth growth and are at the same time consistent with life-cycle inequality facts. In Appendix 1.A.13, we provide a detailed analysis on further dimensions of individual earnings dynamics. We first demonstrate that the model is consistent with standard estimates for the process of earnings using a permanent-transitory decomposition, as in Meghir and Pistaferri (2004), Blundell, Pistaferri, and Preston (2008), or Heathcote, Perri, and Violante (2010). We also corroborate the finding from Hubmer (2018) that the distribution of earnings growth in a life-cycle labor market model is consistent with the empirically observed distribution documented by Guvenen et al. (2019). We also decompose earnings growth and dispersion over the life cycle and demonstrate that our decomposition is consistent with the results in Topel and Ward (1992) on early career wage growth and resolves the tension highlighted in Hornstein, Krusell, and Violante (2011) between earnings dynamics and earnings inequality. The joint con-

10. In Appendix 1.A.10, we show that wage and tenure data from the SCF data align closely to the CPS levels.

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

Figure 1.6. Earnings and wealth

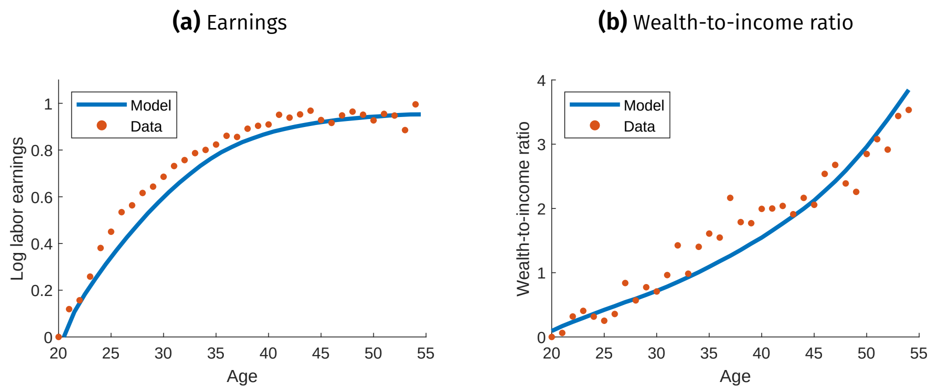

Notes: Panel (a) and (b) show the mean of log earnings and the mean wealth-to-income ratio, respectively, both normalized to 0 at age 20. Wealth-to-income ratios are calculated as the end-of-year assets divided by yearly income. In all panels, the blue lines/squares are the model profiles, while the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

sistency of the model with these facts lends support to the calibration of the human capital and wage processes as the two dimensions underlying life-cycle earnings dynamics in the model.

Finally, we discuss the mapping of the income process to wealth accumulation as a key model prediction to validate the model-implied consumption-savings dynamics. Appendix 1.A.14 demonstrates that the consumption-saving and earnings dynamics of our model result in a joint distribution of earnings and wealth that is consistent with the SCF data.

1.3.3 Employment stability and wealth accumulation

While the life-cycle profiles are targeted when bringing the model to the data, the relationship between employment stability and wealth accumulation in Figure 1.7 is not. In our empirical analysis (Section 1.2.2), we document a negative correlation between nonemployment and wealth accumulation. Figure 1.7 demonstrates that our model is consistent with this empirical fact. It shows wealth-to-income ratios by nonemployment duration in the PSID data and using model-simulated data. The model outcome implies a slope of −0.05 in Figure 1.7, which is close to the empirical estimates in Table 1.1 that range between −0.07 and −0.03.

Figure 1.7. Wealth-to-income ratio

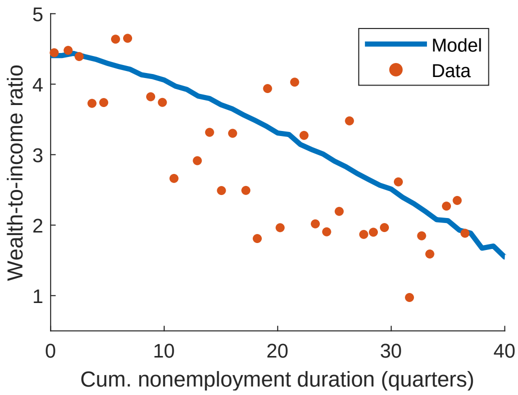

Notes: The figure shows the wealth-to-income ratios at age 50 against the accumulated nonemployment duration in the model and in the data. The blue line shows the model profile, while the red dots show the empirical wealth-to-income ratios from the PSID data.

1.4 Consequences of employment heterogeneity on income and wealth

Using our model framework, this section explores the consequences of differences in employment stability on life-cycle earnings and wealth dynamics. More specifically, we assess the channels through which employment stability shapes the differences in earnings and wealth. In a next step, we study the cost of displacement of workers.

1.4.1 Consumption-saving behavior

In Figure 1.8, we compare the life-cycle earnings and wealth-to-income ratio profiles generated in the model to the empirical counterparts for workers with stable and unstable employment history at age 50. The empirical profiles are the same profiles as in Figure 1.3 and Figure 1.4 in Section 1.2. While the average profiles of log earnings and wealth-to-income ratios were targeted when calibrating the model to the data, the profiles by stable and unstable employment groups were not explicitly targeted. Still, the model matches the empirical profiles fairly well.

Looking at labor earnings of employed workers in Figure 1.8a, we see that labor earnings are diverging quickly and already differ substantially after age 25. We find that labor earnings of workers with stable and unstable employment increases quickly in the early age and remains persistent over the life cycle. At age 50, workers with stable employment have 20% higher labor earnings. The income difference would be even larger if we include the income of nonemployed workers as they only receive benefits due to nonemployment.

Figure 1.8. Earnings and wealth-to-income ratio

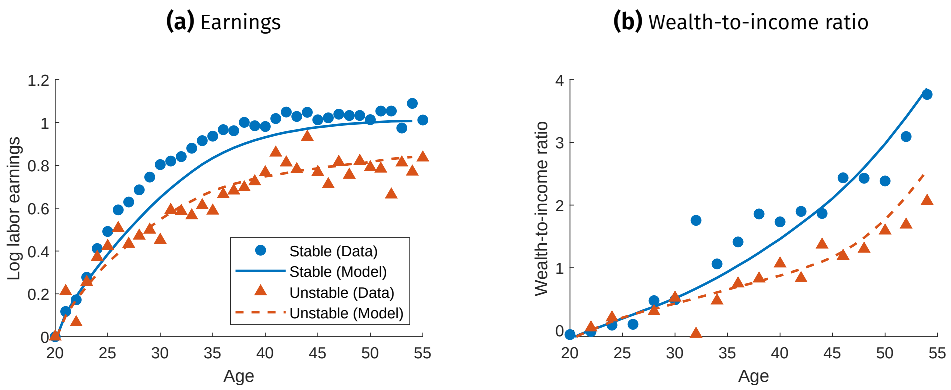

Notes: Panel (a) shows the log of labor earnings of employed workers in the data and in the model, normalized to 0 at age 20. Panel (b) shows the wealth-to-income ratio in the data and in the model, normalized to 0 at age 24. The solid lines are the model profile, while the dots show the estimated empirical 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 nonemployment duration at age 50. Data are from the Panel Study of Income Dynamics.

In the following, we decompose the earnings difference between workers with stable and unstable employment into a wage and human capital component in Figure 1.9a. Both wages and human capital strongly differ across the two groups of workers and increase in age. The earnings difference in early life is mainly explained by the gap in the wage level. Over time, the gap in human capital gains in importance and at age 55, around half of the total earnings difference is explained by the difference in human capital. The reason why the gap in human capital starts to grow after age 25 is because in the beginning of the life cycle, employment rates do not differ strongly across the two groups of workers. We show the employment rate difference between workers with stable and unstable employment in Figure 1.9b. The employment rate difference jumps at age 25 as we group workers into stable and unstable employment using the employment history between age 25 and 50. The increasing gap in the employment rate leads to an increasing gap in human capital. Workers in a less stable employment path spend less time in employment so that she has fewer opportunities to invest in human capital, especially when young and when human capital investment is most productive.[12]

12. It is important to note that productivity differences stemming from unobserved differences in human capital are empirically part of frictional wage dispersion.

Figure 1.9. Earnings difference decomposition

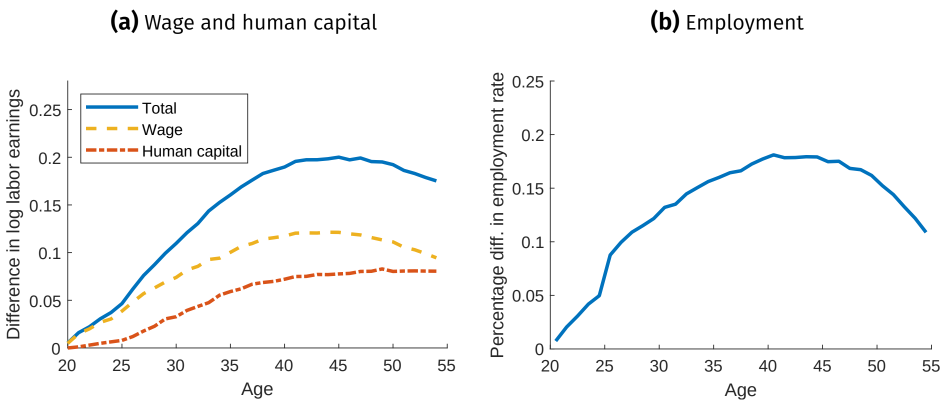

Notes: The figure shows the decomposition of the difference in log of labor earnings between the stable and unstable groups of workers in the model. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of nonemployment duration at age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

1.4.2 Decomposing differences in wealth accumulation dynamics

Employment stability affects the wealth accumulation of workers in three dimensions. First, workers with higher employment risk accumulate more precautionary savings in order to smooth consumption in case of nonemployment. We label this effect as "precautionary saving." Second, workers with a stable employment history experience a higher income growth as they have a better opportunity to climb the job ladder and accumulate more human capital, which is the "income growth" effect of employment stability on wealth accumulation. Lastly, nonemployment duration has a direct effect on wealth because workers who become nonemployed do not receive labor earnings and have to decumulate their savings. We label this effect as the "nonemployment" effect.

To trace out these different channels how employment stability affects workers' wealth accumulation, we construct counterfactual models. For the "precautionary saving" component, we assign the same age-dependent levels of human capital and wage levels to the workers with stable and unstable employment. While workers face different separation rates, we assume that there is no realization of job separation. In this case, differences in employment risk generate different savings patterns for the two groups of workers, while we shut down other channels through which employment stability affects wealth. In a second step, we remove heterogeneity in employment stability from the model by assigning the same age-dependent job-separation rate to all workers. At the same time, we fix the labor earnings for the two groups of workers at the baseline level by assuming that workers face the

Figure 1.10. Decomposition of difference in wealth-to-income ratio

(a) Income and consumption smoothing　　**(b)** Unemployment and buffer stock

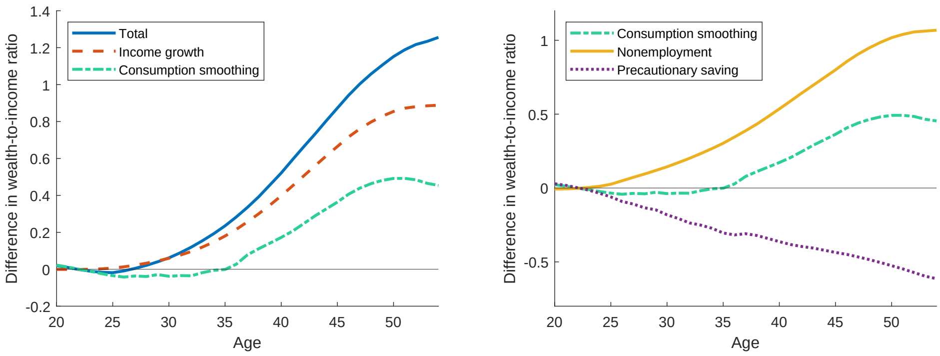

Notes: This figure shows the decomposition of the difference in wealth-to-income ratio between the stable and the unstable employment groups of workers. The stable and unstable groups consist of workers in the bottom 75% and the top 25%, respectively, of the distribution of nonemployment duration at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

same human capital and wage levels from the baseline model. Because workers have different life-cycle profiles of labor earnings while there is no separation risk, the difference in saving pattern between the two groups of workers is generated by the difference in labor earnings only.

The "nonemployment" effect is computed as follows. We assume that workers have the human capital and wage levels from the baseline model and face heterogeneity in employment stability, but there is no realization of nonemployment. We derive the wealth-to-income ratios of stable and unstable workers in the counterfactual model and subtract from the wealth-to-income ratios of these workers in the baseline model. The difference in wealth-to-income ratios measures the "nonemployment" effect.

Figure 1.10 displays the decomposition of the difference in wealth-to-income ratio between the stable and unstable employment groups of workers over the life cycle. The difference in wealth-to-income ratio is largely driven by differences in income growth. Figure 1.10a shows that the "income growth" effect grows over the life cycle and accounts for more than 60% of the total wealth difference at the end of the prime-age working life. The remaining difference in wealth-to-income ratio is explained by a consumption smoothing effect. As workers with low employment stability have the desire to save for precautionary motives, the “consumption smoothing” component is negative at the beginning of higher separation risk. The consumption smoothing effect consists of the sum of precautionary savings and the nonemployment effect. We show the precautionary savings and the wealth difference generated by nonemployment spells separately in Figure 1.10b. As workers with unstable employment always accumulate more precautionary savings, the wealth difference due to precautionary savings is always negative and almost linearly decreasing over the life cycle. In contrast, the wealth difference due to asset decumulation when nonemployed spells is always positive and increasing over the life cycle. Workers who become nonemployed are subject to a drop in their income and use their savings to smooth consumption. As workers with stable employment rarely become nonemployed by construction, they almost never have to dissave to smooth negative income shocks. Hence, the difference in wealth due to the realization of nonemployment is always positive and increasing over the life cycle.

1.5 Employment heterogeneity and the consequences of job loss

To provide further intuition for the consequences of employment stability on life-cycle dynamics, this section explores the consequences of job loss and their relationship to heterogeneity in employment stability. To do so, we adapt the approach from the empirical literature on job displacement (Jacobson, LaLonde, and Sullivan, 1993) and compare identical workers at age 40 where one worker is losing the job while the other worker remains employed and only faces the probability of future job loss. Specifically, we compare a cross section of 40-year-old employed workers to the same group of workers who have been sent (exogenously) into nonemployment at age 40.[13] Empirical studies document that such job displacements lead to large and persistent earnings losses for workers (Jacobson, LaLonde, and Sullivan, 1993; Couch and Placzek, 2010; Davis et al., 2011), and heterogeneity in employment stability has been identified as a key ingredient in accounting for large and persistent earnings losses in structural models (Jarosch, 2015; Jung and Kuhn, 2018). Models without heterogeneity in employment stability struggle to account for the persistence in earnings losses (Low, Meghir, and Pistaferri, 2010), a fact we also highlight in Appendix 1.A.12 where we show earnings losses for a model without employment stability heterogeneity and how the resulting earnings losses are only transitory.

Figure 1.11a shows that our baseline model with heterogeneity in employment stability implies large and persistent earnings losses for the average 40-year-old worker. The initial earnings drop in displaced workers amounts to around 15%. Over the subsequent five years, displaced workers are able to cut initial earnings losses by half, but there is little further catch-up. The figure also shows the evolution of the components of earnings (wage and human capital) to uncover the underlying

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

Figure 1.11. Cost of displacement

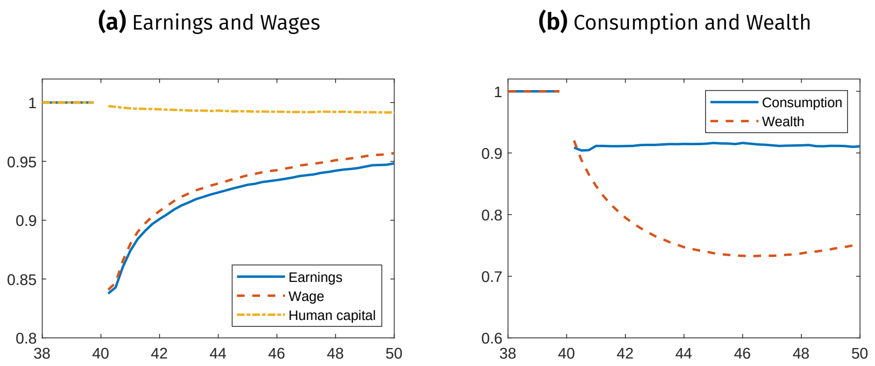

Notes: This figure shows the evolution of earnings, consumption, and wealth of workers who become unemployed at age 40 relative to a control group of workers who remain employed. Prior to displacement, both groups are identical.

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

Figure 1.11b turns to the consequences of job loss for consumption and wealth. Looking at consumption, we see a sharp (roughly 10%) drop in consumption directly on impact. After the one-time persistent shock, consumption dynamics show only a very slight upward trend. These consumption dynamics can again be rationalized by the permanent income hypothesis. Directly upon job loss, agents anticipate that they enter an employment trajectory with lower and more volatile income. Income after job loss will be persistently lower due to lower earnings and lower employment rates, and income will be more volatile because of lower employment stability. As a consequence of lower permanent income and higher volatility, agents permanently reduce consumption and increase their precautionary savings to smooth consumption in the future. On impact, consumption drops less than in-

Figure 1.12. Effects of displacement by employment stability

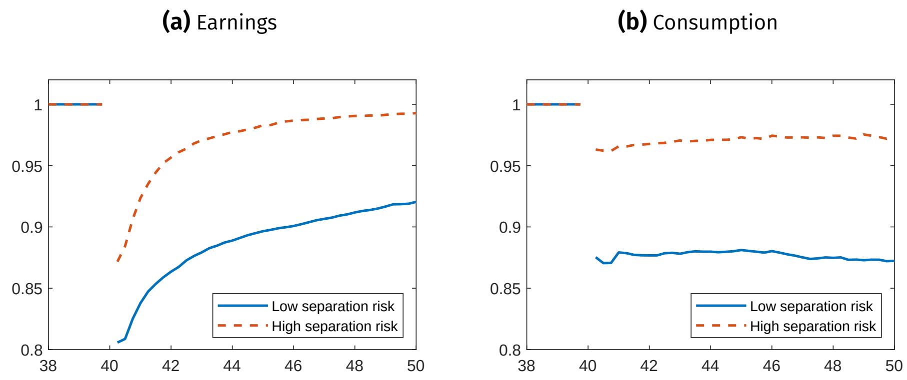

Notes: This figure shows the evolution of earnings and consumption of workers who become unemployed at age 40 relative to the control group. Workers with low-separation risk are employed in jobs belonging to the bottom quartile of jobs by separation rate at the time of displacement. Workers with high-separation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

come, and this difference shows up directly in wealth dynamics as wealth is used to smooth the transition to the new, lower permanent income level (Kuhn, 2013). Four years after job loss, wealth levels stabilize 25% below the level of nondisplaced workers and remain persistently lower in line with recent empirical results in Barnette (2020). Three reasons account for this lower average wealth after job loss. First, income is lower so that the wealth level adjusts, too. Second, the job loss has flattened the life-cycle income profile. Current income is now lower relative to income during the rest of the life cycle, which reduces the need for life-cycle savings that aim at reshuffling age-varying income over time. Third, the lower employment stability after the job loss sets agents on the kind of Sisyphus saving cycle: While cycling through unstable jobs, workers' ability to accumulate wealth is mitigated by the fact that consumption smoothing over repeated spells of nonemployment reduces any accumulated savings. These consumption and labor market dynamics also provide intuition for the positive relationship between wealth and nonemployment duration implied by the model and observed in the data (Figure 1.7). Workers who lose their jobs experience significant decreases in income. During the transition, wealth declines, earnings recover so that wealth-to-income ratios fall, and we get the negative relationship between nonemployment duration and wealth-to-income ratios, as observed in Figure 1.7.

Figure 1.12 explores as the next step how differences in employment stability shape the consequences of job loss. It shows the results of the previous displacement experiment but compares workers with initially high-separation risk and low-separation risk. Except for employment stability, the workers are otherwise identical. On impact, losing the job with low-separation risk or high-separation risk leads to earnings losses of 12% and 18%, respectively (Figure 1.12a). While these initial earnings losses are similar, the recovery from the initial shock is strikingly different between the stable and unstable employment. To understand the reasons behind these differences, it is important to keep in mind that the counterfactual earnings dynamics provided by the control group of workers in stable and unstable jobs differ.

For the unstable employment, we see a recovery that is very quick and shows almost full mean reversion within five years. The reason for the fast recovery is that the group of workers in unstable employment who did not lose their job initially are very likely to lose their job moving forward, so that differences between job losers and initial job stayers quickly vanish. Put differently, unstable employment exhibits a lot of mean reversion. This strong mean reversion also explains why earnings losses in a model matching average separation rates but abstracting from heterogeneity in employment stability are only transitory (Appendix 1.A.12). By contrast, the consequences of job loss are strikingly different for workers who are initially in stable employment. Now the same logic applies but with different consequences. If workers in initially stable employment had not lost their job, the high employment stability would imply that they would have been unlikely to lose their job in the future. Hence, high employment stability implies little mean reversion and high persistence of the earnings process. This implies that labor market search models that aim at generating persistent earnings dynamics need at least some jobs that are highly stable in order to reduce mean reversion in labor market outcomes.

We have already seen that after a job loss, workers adjust their consumption immediately to their expectations about the level and volatility of their future earnings path. We have also seen that with heterogeneity in employment stability, earnings paths after job loss differ substantially, so that workers who have lost an unstable job expect the shock to their earnings to be much smaller and less persistent, and precautionary savings allow these workers to smooth consumption after the job loss (Figure 1.12b). Workers who lose their stable employment experience a much larger and more persistent drop in earnings, and their wealth allows them to smooth only the transitory part of the income loss but not the permanent shock to income, so that their consumption path moves persistently down by 13%. The additional drop in consumption in excess of the persistent earnings drop results from the differences in employment rates that lead to a larger drop in income compared to earnings. The employment effect is substantially larger for workers who had a low separation risk as employment rates starting from a stable employment are much higher than employment rates for workers after a job loss.

Our analysis highlights the large heterogeneity in the consumption responses after a job loss. Such heterogeneity provides a potentially important link between individual consumption behavior and macroeconomic dynamics. If, for example, all jobs that are lost at the macroeconomic level are low-stability jobs, the consumption drop would be 4% on average. By contrast, if all job losses were in stable jobs, then the consumption drop would be 13% —more than three times as large.

1.6 Conclusions

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

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

1.A Appendix

1.A.1 Heterogeneity in job destruction rates

This section provides additional evidence for the differences in job destruction rates across firms of different ages. Figure 1.A.1 shows that the heterogeneity in job destruction rates persists even after controlling for year and Metropolitan Statistical Area (MSA) fixed effects.

Figure 1.A.1. Heterogeneity in job destruction rates by firm age

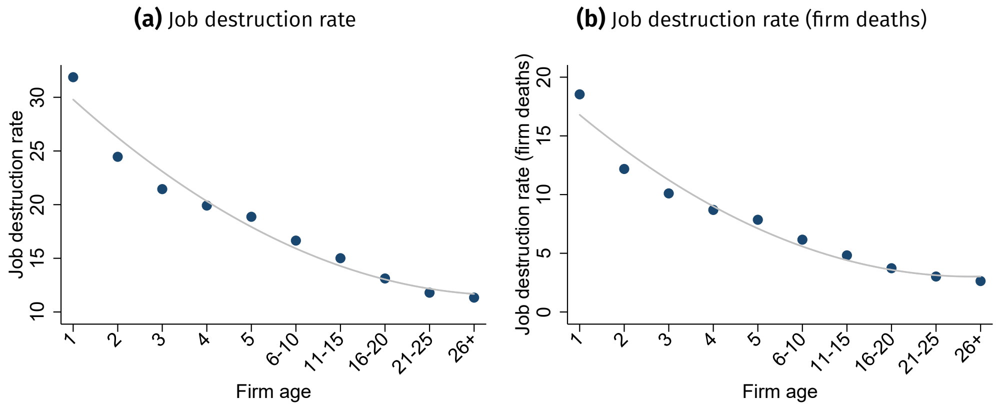

Notes: Panel (a) shows the relationship between job destruction rate and firm age from the Business Dynamics Statistics. Panel (b) shows the relationship between job destruction rate due to firm deaths and firm age. Job destruction rates are computed as the number of jobs destroyed over the last 12 months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

1.A.2 Tenure distribution with heterogeneity in worker types

A possible explanation for the fanning out of the tenure distribution over the life cycle could lie in the existence of worker type heterogeneity. It is plausible to imagine a situation in which some workers, because of their intrinsic characteristics, change jobs frequently, whereas others keep the same jobs for long periods of time. As shown in this section, introducing worker types can indeed lead to an increasing dispersion of tenure over the life cycle; however, the resulting tenure profiles fail to fully represent the empirical patterns. Furthermore, as already highlighted in the main part of the paper, the resulting profiles of average transition rates are inconsistent with the empirical profiles.

To illustrate this point, Figure 1.A.2 presents results from a simulation exercise where workers ex ante differ in their labor market mobility. There are two types of workers: a stayer type and a mover type. The population of workers consists of equal

Figure 1.A.2. Worker types and tenure distribution

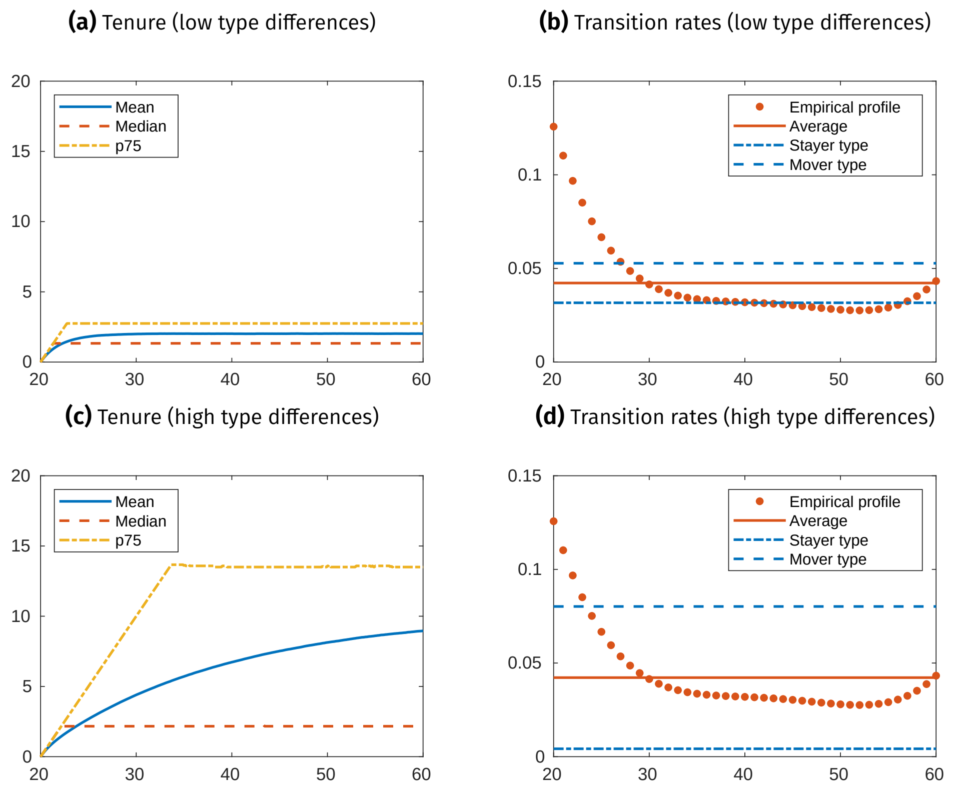

Notes: This figure shows the consequences of heterogeneity in worker types on life-cycle tenure and transition rate profiles. The left panels show life-cycle tenure dynamics from a simulation where workers have different age-invariant labor market transition rates (separation and job-to-job transitions). The three lines show mean tenure, median tenure, and the 75th percentile of the tenure distribution. The right panels show the monthly transition rates used in the simulation and the empirical life-cycle profile. Panels (a) and (b) show results from a simulation with low type differences, where workers of the stayer type have a transition rate that is 25% lower than the average transition rate and workers of the mover type have a 25% higher transition rate. Panels (c) and (d) show results from a simulation with high type differences, where differences for both types of workers relative to the average transition rates are increased to 90%.

shares of both types. On average, workers have an age-invariant transition rate that corresponds to the average empirically observed monthly transition rate resulting from separations into nonemployment and job-to-job transitions. We present two cases of worker type heterogeneity with different degrees of worker heterogeneity that preserve the same average transition rate. In panels (a) and (b) of Figure 1.A.2, the stayer type has a transition rate that is 25% lower than the average transition rate, whereas the mover type has a transition rate that is 25% higher than the average. In panels (c) and (d), we consider an alternative case in which type heterogeneity is more substantial: the stayer type now has a transition rate that is 90% lower, and the mover type has a transition rate that is 90% higher than the average transition rate. The left panels show the resulting tenure distribution, and the right panels show the transition rates.

Compared to the empirical profiles shown in Figure 1.5, it is clear that none of the considered cases matches empirical tenure profiles. Although the increase in tenure dispersion is fairly substantial with high type differences, the profiles of median tenure and the 75th percentile flatten out relatively early in the working life. Even more important, worker types cannot provide a good explanation for the decreasing convex profile of transition rates.

1.A.3 Employment inequality

For our structural model, we have to take a stand on the extent and the source of heterogeneity in job stability. The natural alternative view to job-related differences in stability is that heterogeneity is worker related, with some workers being of a "mover" type with less stable employment and others of a "stayer" type with more stable employment. The following analysis will do two things. First, we provide corroborating evidence for heterogeneity in job stability and quantify the extent of employment inequality using a simple summary statistic. Second, we provide empirical evidence from the BDS and Monte Carlo simulations to argue that job heterogeneity must be the important driver of heterogeneity in job stability.

The increasing life-cycle dispersion of job stability mirrors the widely studied increase in wage inequality with age (Heathcote, Perri, and Violante, 2010), and we will refer to this dispersion in job stability respectively as employment inequality. To quantify the extent of employment inequality in the data, we propose a simple summary statistic: the ratio of expected tenure of a representative worker to observed average tenure. Both components are observed in the data. Without heterogeneity in worker flow rates (i.e., in the representative-worker case), mean tenure is simply the inverse of the average job outflow rate. If there is heterogeneity in job stability that constitutes a mean-preserving spread of transition rates, then, according to Jensen's inequality, mean tenure increases and the ratio of the two tenure statistics provides a measure of underlying heterogeneity. Specifically, denote mean tenure in a labor market with homogeneous outflow rates from jobs by $\mathbb{E}[T] = (\bar{\lambda} + \bar{\pi}_{ee})^{-1}$ where $\bar{\lambda}$ denotes the average transition rate to nonemployment and $\bar{\pi}_{ee}$ the average job-to-job transition rate. If we denote mean tenure as observed in the data by $\bar{T}$, then we summarize the extent of heterogeneity in job stability, or employment inequality, σ_E, by the ratio of $\bar{T}$ to $\mathbb{E}[T]$:

$$\sigma_E = \frac{\bar{T}}{\mathbb{E}[T]} = \frac{\overbrace{\bar{T}}^{\substack{\text{average} \\ \text{observed tenure}}}}{\underbrace{(\bar{\lambda} + \bar{\pi}_{ee})^{-1}}_{\substack{\text{expected tenure} \\ \text{w/o heterogeneity}}}} = \bar{T} \times (\bar{\lambda} + \bar{\pi}_{ee})$$

We derive in Appendix 1.A.4 the approximate equivalence between σ_E and the coefficient of variation of outflow rates, justifying the intuition that σ_E is a measure of employment inequality closely related to the variance of (log) wages as a typical measure of life-cycle wage inequality. Regarding the interpretation, remember that if there is no heterogeneity in job stability (i.e., in the representative-worker case), the ratio will be one as average tenure $\bar{T}$ equals expected tenure $\mathbb{E}[T]$. By contrast, the ratio will exceed one whenever there is heterogeneity in job stability. The level of σ_E also has a very intuitive interpretation. Consider, for example, the case $\sigma_E = 3$. In this case, average tenure is three times larger than expected based on the average observed transition rates.

Figure 1.A.3. Employment inequality

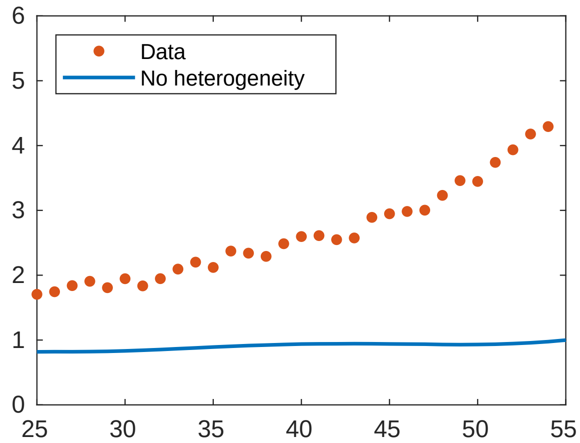

Notes: Estimated life-cycle profile of employment inequality σ_E. Employment inequality is computed as the ratio of observed to expected tenure given average job outflow rates. Underlying data on outflow rates and tenure come from CPS and are taken from Jung and Kuhn (2018). The *no-heterogeneity* case is constructed from a Monte Carlo simulation where all workers have the average age-dependent labor market transition rate (separation and job-to-job transitions) but no cross-sectional heterogeneity conditional on age.

Figure 1.A.3 shows the empirical life-cycle profile of employment inequality based on CPS data from Jung and Kuhn (2018) together with a counterfactual Monte Carlo simulation with no heterogeneity in job stability. Most importantly, we see immediately that employment inequality is always above one in the data, indicating that there is heterogeneity in job stability. Over the life cycle, we find, similar to wage inequality, an almost linear increase with age. At age 25, employment inequality starts at slightly below 2 and increases to above 4 at age 55. During the middle of working life, the level of employment inequality is around 3. Hence, a job lasts three times longer than the average transition rates suggest. The simulated no-heterogeneity case shows no life-cycle increase. It is initially slightly below one as a result of the transitional dynamics after starting all workers from zero tenure.

We interpret this result as suggesting that during a worker's prime-age working life, heterogeneity in job stability in the U.S. labor market is large and economically significant.

In a final step, we explore the potential sources of heterogeneity in job stability. We start in Figure 1.A.4 with evidence from the BDS on heterogeneity in job loss probabilities across employers of different age. We consider two definitions of job loss: total job destruction rate (Figure 1.A.4a) and job loss due to firm closure (Figure 1.A.4b). We remove year and industry fixed effects in both cases.[14] In Figure 1.A.4a, we observe large heterogeneity in job loss across employers, with the least stable employer having job loss rates that are twice as large as the most stable employers. Such differences in observed job loss could be the result of differences in job-to-job transitions or worker quits. Figure 1.A.4b therefore considers the more restrictive definition of job loss where we consider only job loss due to firm closure. For this case, we find the differences in job stability to be even larger, with the least and the most stable employers differing by a factor of four. Such large differences in the probability of job loss across employers are also supported by existing research. Larkin (2019) documents large heterogeneity in separation rates into unemployment in U.S. CPS data, and Jarosch (2015) documents such heterogeneity across German employers. Next, we extend a theoretical argument supporting this conclusion from Jung and Kuhn (2018) using Monte Carlo simulations.

14. Appendix 1.A.1 shows that controlling for year and MSA fixed effects yields very similar results. The BDS data do not provide publicly available data where industry and geographical breakdown is available in the same file.

Figure 1.A.4. Heterogeneity in job destruction rates by firm age

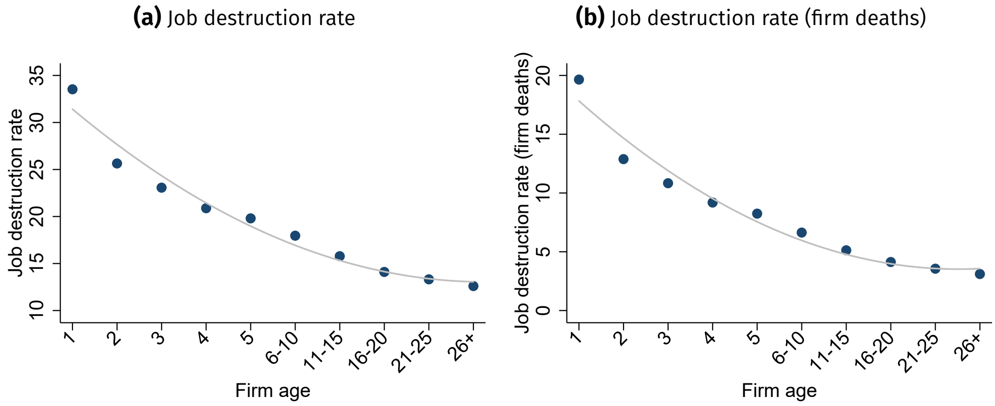

Notes: Panel (a) shows the relationship between job destruction rate and firm age from the BDS. Panel (b) shows the relationship between the job destruction rate due to firm deaths and firm age. Job destruction rates are computed as the number of jobs destroyed over the last 12 months divided by average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

We provide two Monte Carlo simulations to revisit the extent of heterogeneity in job stability and its sources. Figure 1.A.5a revisits the extent of heterogeneity and shows a simulation of the tenure distribution for a representative-worker case where we only feed in age heterogeneity in job stability but rule out any cross-sectional heterogeneity by age (i.e., we use the average age profiles of transitions to nonemployment and job-to-job rates). The age pattern of the tenure distribution differs starkly from its empirical counterpart in Figure 1.A.8a. The moments of the tenure distribution increase much less, there is less dispersion at each age, and even the 75th percentile of the tenure distribution remains bounded at about four year while it increases to 25 years in the data. Consistently, Figure 1.A.3 shows no employment inequality in such a simulated economy. This simulation therefore supports the conclusion that there is large heterogeneity in job stability in the U.S. labor market.

Figure 1.A.5b revisits the question on the sources of heterogeneity based on a second Monte Carlo experiment. It considers a stylized case with workers of a mover type and a stayer type. Workers of the mover type have low job stability, whereas workers of the stayer type have high job stability. For simplicity, we assume that both groups are of equal size. The life-cycle profiles of transitions into nonemployment of the two types are shown as a dashed blue line. By construction, the transition rates are a fixed worker characteristic and do not change over the life cycle, thereby resulting in flat age profiles. The average profile (solid red line) corresponds to the unconditional average of the transition rates, and the red dots show the empirical profile. The average profile and its empirical counterpart show a life-cycle pattern

that appears strongly inconsistent. We find a strongly declining empirical profile in the first 10 years of working life, whereas the simulated profile is flat and does not show any life-cycle variation. In Appendix 1.A.2, we present additional simulation results for this model and show that the life-cycle pattern of the tenure distribution is also at odds with the data. We consider this Monte Carlo evidence as strongly supportive of the conclusion that job differences are a key driver of heterogeneity in job stability in the data.[15]

Figure 1.A.5. Life-cycle profiles and heterogeneity

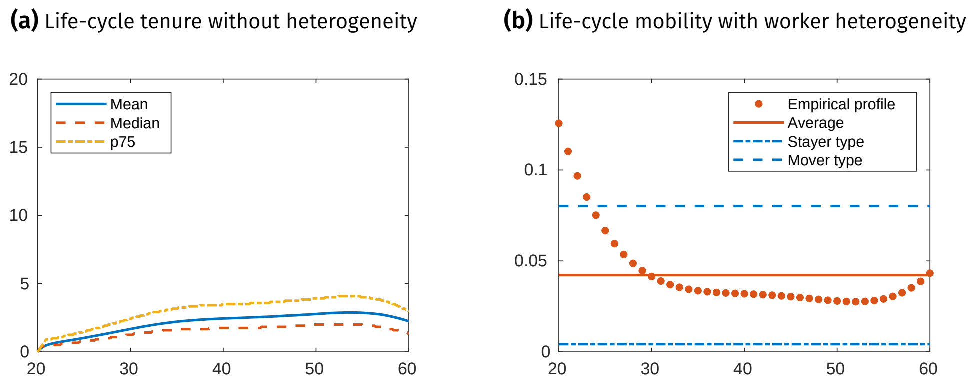

Notes: The figure shows the consequences of heterogeneity in job stability on life-cycle tenure and separation rate profiles. The left panel shows life-cycle tenure dynamics from a simulation where all workers have the average age-dependent labor market transition rate (separation and job-to-job transitions). The three lines show mean tenure, median tenure, and the 75th percentile of the tenure distribution. The right panel shows stylized life-cycle profiles for separation rates with fixed worker heterogeneity in separation rates (dashed blue line) and heterogeneity in job stability (solid red line).

1.A.4 A measure of employment inequality

Suppose there are N different jobs with outflow rates $\{\pi_i\}_{i=1}^N$. Job outflow rates capture all outflow events from jobs to unemployment, out of the labor force, and other employers. To make things simple, assume that the average outflow rate is

$$\bar{\pi} = \frac{1}{N}\sum_{i=1}^{N}\pi_i$$

This assumes that workers are uniformly distributed across jobs. Average tenure in this economy is

$$T_H = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\pi_i}$$

15. The result must not be interpreted as the absence of any fixed worker differences. We will provide a model extension where we allow for differences in worker types, for example, due to differences in educational attainment.

where subscript H denotes explicitly that we consider average tenure in an economy with heterogeneous job stability. Average tenure assuming a representative agent (i.e., one agent with separation rate $\bar{\pi}$) is

$$T_R = \frac{1}{\bar{\pi}}.$$

A measure of employment inequality is

$$\sigma_E = \frac{T_H}{T_R} = T_H \times \bar{\pi}.$$

To see this, consider first

$$\begin{aligned} T_H - T_R &= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\pi_i} - N\left(\sum_{i=1}^{N}\pi_i\right)^{-1} = \left(\sum_{i=1}^{N}\pi_i\right)^{-1}\left(\left(\frac{1}{N}\sum_{i=1}^{N}\frac{\sum_{j=1}^{N}\pi_j}{\pi_i}\right) - N\right) \\ &= \left(\sum_{i=1}^{N}\pi_i\right)^{-1}\left(\left(\sum_{i=1}^{N}\frac{\frac{1}{N}\sum_{j=1}^{N}\pi_j}{\pi_i}\right) - N\right) = \left(\sum_{i=1}^{N}\pi_i\right)^{-1}\left(\left(\sum_{i=1}^{N}\frac{\bar{\pi}}{\pi_i}\right) - N\right) \end{aligned} \tag{1.A.1}$$

Using a second-order approximation of $f(\pi_i) = \frac{\bar{\pi}}{\pi_i}$ around $\bar{\pi}$ and plugging it into equation (1.A.1) yields

$$\begin{aligned} T_H - T_R &= (\bar{\pi})^{-1}\left(\frac{1}{N}\sum_{i=1}^{N}\left(1 - \frac{\pi_i - \bar{\pi}}{\bar{\pi}} - \left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}}\right)^2\right)^2 - 1\right) \\ &= T_R\frac{1}{N}\sum_{i=1}^{N}\left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}}\right)^2 \\ \frac{T_H - T_R}{T_R} &= \frac{1}{N}\sum_{i=1}^{N}\left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}}\right)^2 \\ \sigma_E &= 1 + \frac{1}{N}\sum_{i=1}^{N}\left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}}\right)^2. \end{aligned} \tag{1.A.2}$$

Hence, σ_E corresponds (up to first order) to the coefficient of variation of employment stability π_i. The key advantage of $\frac{T_H}{T_R}$ is that while $\{\pi_i\}_{i=1}^{N}$ remains unobserved, mean tenure and the average separation rate to estimate T_R can be estimated from the data.

1.A.5 Job stability and wealth accumulation: PSID data

Figure 1.A.6 shows a binned scatter plot of wealth-to-income ratios against nonemployment of workers at age 50 after controlling for education and industry. Even after adding the additional controls, there is a significant negative relationship between wealth-to-income ratios and the nonemployment duration.

Figure 1.A.6. Wealth-to-income ratios and nonemployment

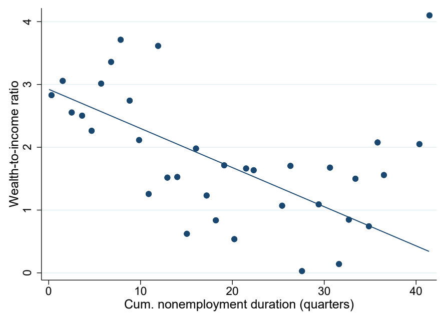

Notes: This figure shows binned scatter plot of wealth-to-income ratios against nonemployment of workers at age 50 after controlling for education and industry.

1.A.6 Job stability and wealth accumulation: SCF data

Using the Survey of Consumer Finances (SCF), we also document the large heterogeneity in job stability in the cross-sectional data of U.S. households. The SCF is a triennial household survey that has become the key resource on the distribution of income and wealth for the United States (see, for example, Kuhn and Ríos-Rull, 2016; Bricker et al., 2017; Kuhn, Schularick, and Steins, 2020). Besides the detailed information on household income and wealth, the SCF also offers information on household members' labor market situation. We pool data across survey waves form 1992 to 2016 and restrict the sample to households with employed household heads of ages 20 to 60.

In addition to the results from the PSID data, we explore the relationship between the labor market situation and wealth accumulation in the SCF. For our analysis, we pool data across survey waves from 1992 to 2016 and restrict the sample to households with employed household heads ages 20 to 60. As our model will abstract from self-employment, we drop households with self-employed household heads and households with extreme wage observations, defined as wages lower than 75% of the minimum wage.[16] Additionally, we exclude the top 1% of households by wealth and earnings as we do not provide a theory of the very right tail

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

of these distributions, and, similarly, we exclude households in the bottom 1% by earnings and households with negative wealth.

Figure 1.A.7. Wealth-to-income ratios, tenure, and number of employers

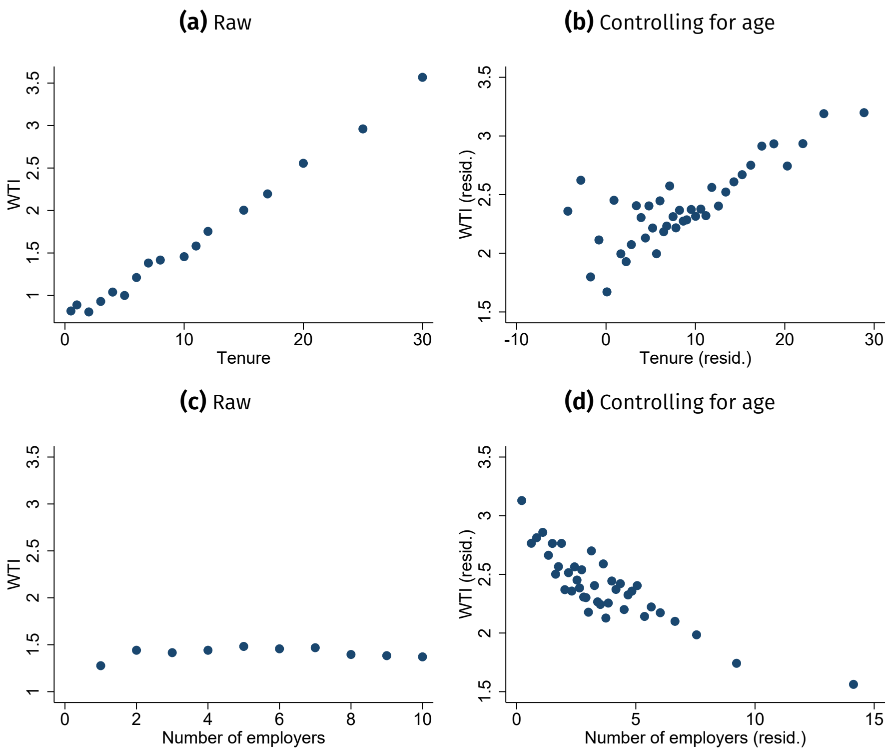

Notes: This figure shows binned scatter plots of wealth-to-income ratios against tenure or number of employers for which a person has worked full-time jobs lasting one year or more. In panels (a) and (c), each dot represents a median wealth-to-income ratio for a given bin. Panels (b) and (d) show binned 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.

Regarding the construction of variables, we consider household wealth as the difference between household assets and debt. Household income is gross income from all sources including transfers, and earnings is income from wages and salaries.[17] We control for income differences nonparametrically by always considering wealth-to-income ratios. Job stability itself is unobserved in the data, and we only observe retrospectively whether an employer-employee relationship has been stable by looking at a worker's employer tenure or the number of employers during a worker's

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

career.[18] Using these statistics as measures of job stability will lead to measurement error regarding the true level of underlying job stability for two reasons. First, realized tenure can be high despite low job stability due to luck. Second, realized tenure can be low despite the worker having a stable job because the worker might have received a better job opportunity and therefore changed employers. We will therefore interpret the observed correlations from this section through the lens of our structural model in Section 1.3. In the structural model, we will also consider realized job tenure when mapping the model to the data. We will also have job-to-job mobility as observed in the data so that we can impose consistency in the measurement of job stability between the model and data. The structural model will, in addition, offer us the opportunity to consider meaningful counterfactuals to study the causal effect of job stability on wealth accumulation in isolation.

Figure 1.A.8 shows life-cycle profiles for tenure and the number of employers in the SCF data. We find that both profiles are positively correlated with age. Looking at the mean, the median, and the 75th percentile of the tenure distribution in Figure 1.A.8a, we observe a spreading out of the distribution as workers age. As pointed out in Hall (1982), the typical U.S. worker has a stable employment history. At age 60, more than 50% of workers have been with their employer for 10 years, and almost a quarter of workers at age 60 have been at the same employer for at least 25 years. The life-cycle profiles for the number of employers in Figure 1.A.8b provide a similar picture. We find that the mean number of employers increases linearly up to age 40 when the growth starts slowing down in the second part of working life. On average, an American worker has worked for four employers at the end of his/her working life.

Figure 1.A.7 provides the estimates of the empirical measures of job stability and their relationship to wealth-to-income ratios. In Figure 1.A.7a, we observe an almost linear relationship between tenure and wealth-to-income ratios. This raw correlation could be the result of tenure and wealth-to-income ratios both increasing with age. In Figure 1.A.7b, we therefore show the correlation between wealth-to-income ratios and tenure, controlling nonparametrically for age. In this case, we still find a positive relationship between job stability and wealth accumulation but with a smaller slope than before. Qualitatively, a positive slope implies that per dollar of income, workers in more stable jobs have more wealth, or, in short, workers with more stable jobs are wealthier.

Figure 1.A.7b offers a second interesting observation. While the relationship in the raw data appears almost perfectly linear (Figure 1.A.7a), the relationship turns into a U-shape for low tenures after controlling for age. This U-shape relationship

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

Figure 1.A.8. Tenure and number of employers over the life cycle

(a) Employer tenure **(b)** Number of employers

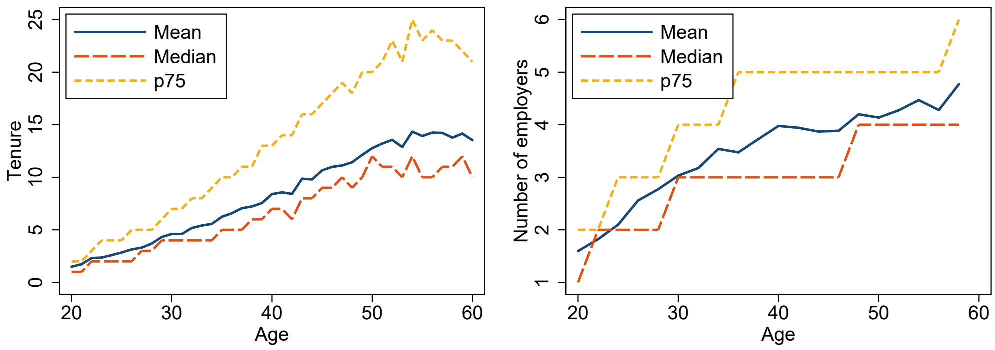

Notes: Panel (a) shows the life-cycle evolution of the cross-sectional distribution of tenure (in years). Panel (b) shows the life-cycle evolution of the cross-sectional distribution of number of employers for which a person has worked full-time jobs lasting one year or more. Two-year age bins are used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

means that workers who have low tenure relative to their age group tend to have higher wealth-to-income ratios. As we will see below, this is a characteristic property of the model when job losers accumulate precautionary wealth before layoffs and then get a negative shock to income, so that wealth-to-income ratios increase. Consumption smoothing of job losers will lower wealth-to-income ratios over time and allow them to converge back to their target wealth-to-income ratio if no further job loss will occur.

Figures 1.A.7c and 1.A.7d corroborate the previous findings using the total number of employers as an alternative measure of job stability. The effect of age on the relationship can now be seen even more clearly. Before controlling for age (Figure 1.A.7c), there is no apparent relationship between wealth-to-income ratios and the number of employers. After age effects are taken out, we find a declining relationship between the number of employers and wealth-to-income ratios. Following the interpretation that more employers are a consequence of less stable jobs, we find again that job stability (fewer employers) is positively related to wealth-to-income ratios. These results on tenure and number of employers point toward a positive relationship between job stability and wealth accumulation, with workers in more stable jobs being wealthier.

To make sure that the observed relationship between labor market experience and wealth is not driven by demographic characteristics of workers, systematic differences in jobs across industries and occupations, or differences in risk attitudes among workers, we perform further robustness checks where we control for additional observable characteristics of households in the SCF. Figure 1.A.9 shows that our findings are not affected by the inclusion of additional controls. In the first column, we nonparametrically control for age, education, occupation, and industry. The relationship between tenure and wealth-to-income ratios (the top row) remains unaffected and significant. The same holds for the relationship between the number of employers and wealth-to-income ratios (the bottom row). In the second step, we additionally control for differences in risk attitudes of workers by nonparametrically controlling for different levels of risk attitudes, as elicited in the SCF survey. As shown in the second column, the relationship between labor market experience and wealth is not affected by the inclusion of all these additional variables.

1.A.7 Value functions for the transition phase

In the transition phase, agents solve a fixed point problem. As a result, value functions do not have any time index. The value functions for the transition phase follow directly the value functions of the working phase. The only difference is that they comprise a probability ψ that at the end of the period, the worker retires and enters the retirement phase. All decisions are otherwise identical to the working phase.

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

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

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

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

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

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

$$\begin{aligned} V_e^{T,P}(a,w,\lambda,h) \;&=\; \max_{\{c,a'\geq 0\}} u_j(c) + \beta\Big[\psi V_r(a',w,h,j_r \;=\; 1) + \\ &\qquad (1-\psi)\left(\pi_e V_e^{T,S}(a',w,\lambda,h) + (1-\pi_e)V_e^T(a',w,\lambda,h)\right)\Big] \\ s.t. \qquad & c \;=\; (1+r)a + y(w,h,e) - a', \end{aligned} \tag{1.A.3}$$

where V_r denotes the agent's value function in the retirement phase, $V_e^{T,P}$ denotes the employed agent's value function at the production stage, $V_e^{T,S}$ denotes the employed agent's value function at the search stage, and V_e^T denotes the value function

Figure 1.A.9. Wealth-to-income ratios, tenure, and number of employers (with additional controls)

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

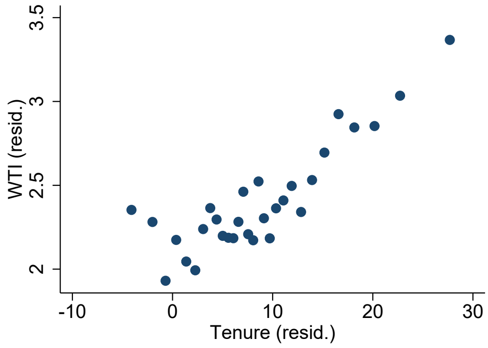

(b) Controls: age, education, occupation, industry, risk attitude

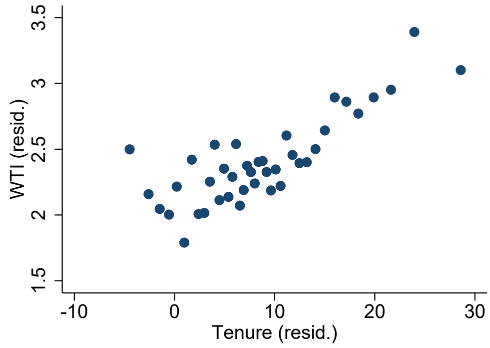

(c) Controls: age, education, occupation, industry

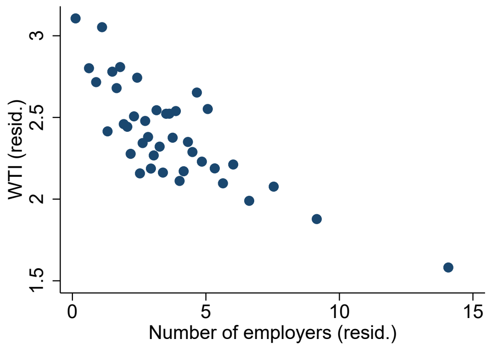

(d) Controls: age, education, occupation, industry, risk attitude

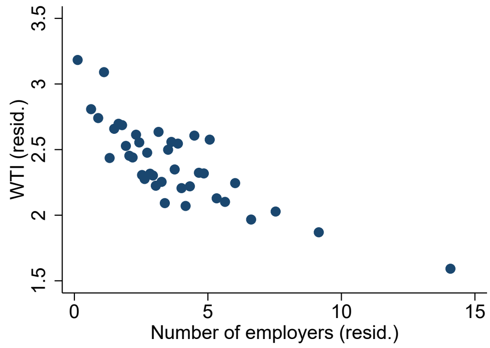

Notes: This figure shows binned scatter plots of wealth-to-income ratios against tenure or number of employers for which a person has worked full-time jobs lasting one year or more. Panels (a) and (c) show binned scatter plots of wealth-to-income ratios against tenure or number of employers after nonparametrically 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.

of an employed worker at the beginning of the transition phase. The value function of an employed worker at the search stage of the transition phase is

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

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

The value function of a nonemployed worker at the production stage is

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

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

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

1.A.8 Model solution and estimation

1.A.8.1 Solving the model

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

We assume that both wages and job destruction probabilities have after standardization a truncated exponential marginal distribution with support of $[0,1]$.[19] To allow for a possible correlation between both marginal distributions, we construct a joint distribution over standardized wages and job destruction probability $F(w^*, \lambda^*)$ using Frank's copula C_θ, where the value of θ determines the correlation between w^* and λ^*. Finally, we discretize this distribution into bins that correspond to grids for w and λ.

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

19. We standardize the support of wages and job destruction probabilities to allow for an easier numerical implementation of the joint distribution. We discuss the details in Section 1.3.1.

1.A.8.2 Parameter estimation

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

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

where the empirical profiles are denoted with a hat.

1.A.8.3 Discussion of identification of model parameters

All parameters of the model are jointly determined, and we refrain from providing a formal identification proof. Here we provide an intuitive discussion on how model parameters are related to the model predictions, which we match to the data to determine the parameter values.

The job offer probabilities when employed or nonemployed, π_e and π_u, are informed by the average job-to-job and job-finding rate over the life cycle. The shape of the joint distribution of job offers $f(\lambda, w)$ is informed by the life-cycle profiles of earnings, tenure, and transition rates. The parameter of the marginal distribution of job destruction probabilities ψ_λ is informed by the life-cycle profiles of the separation rate and tenure. The relative proportion of stable jobs in the job-offer distribution influences how quickly workers sort into stable jobs and as a result accumulate higher tenure due to lower incidence of nonemployment. Consequently, if stable jobs are frequently sampled, the separation rate will quickly decline after labor market entry and tenure dispersion will increase substantially. Similarly, the parameter of the marginal distribution of wages ψ_w is informed by the shape of the life-cycle profile of the average wage. If high-wage offers arrive frequently, the life-cycle growth of average wage will be faster compared to a situation in which high-wage offers arrive very infrequently. The parameter θ, which governs the correlation between wages and separation rates in the job-offer distribution, is informed by the joint life-cycle evolution of the job-to-job rate and the separation rate. If job stability and wages are strongly positively correlated, workers quickly find the best jobs, and the job-to-job rate and separation rate synchronously decline. On the other hand, if the correlation is weak, workers take longer to find a stable and well-paying job, and the job-to-job rate declines more slowly over the life cycle.

Parameters governing human capital dynamics, ρ, p_H and p_H^*, are informed by the life-cycle profile of the variance of earnings. The higher the probability of human capital upgrading, the higher the life-cycle increase in the variance of earnings. On the other hand, the profile of mean earnings in the second half of the working life helps to identify the utility cost of effort κ. At this stage in the working life, earnings growth comes almost exclusively from human capital accumulation, and the utility cost of effort controls when human accumulation starts to slow down. Finally, the wealth-to-income profile informs the discount factor β.

1.A.8.4 The estimated job-offer distribution

Panel 1.A.10a of Figure 1.A.10 shows the estimated job-offer distribution for wages and separation rates, which is asymmetric with most of the probability mass concentrated at low-wage, unstable jobs. Additionally, we also find that wages and separation rates are negatively correlated, implying that high-wage jobs have a low separation rate and low-wage jobs have high separation rates. Panel 1.A.10b additionally shows the conditional distribution of separation rates for different levels of wages. The distribution of separation rates in low-wage jobs first-order stochastically dominates the distribution of separation rates in high-wage jobs.

1.A.9 Earnings and consumption inequality

In the next step, we explore the model's ability to account for the life-cycle pattern of earnings variance. Figure 1.A.11a shows the life-cycle increase in log earnings variance and the close match between model and data. We see the typical almost linear increase in the variance by age in both the model and the data. Labor market search models oftentimes struggle to account for this increase, as discussed in Lise (2012), Jung and Kuhn (2018), or Hubmer (2018). Augmenting models of job search with differences in human capital accumulation provides one way to account for the observed increase (Hubmer, 2018; Jung and Kuhn, 2018). While we build on this approach, we further refine it by endogenizing the human capital ac-

Figure 1.A.10. Job-offer distribution

(a) Joint distribution **(b)** Conditional distributions

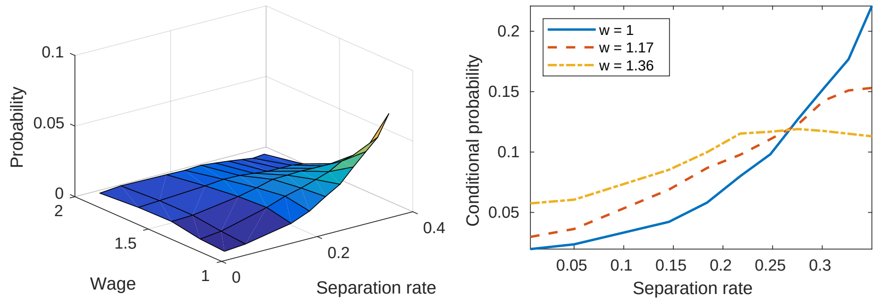

Notes: Panel (a) shows the estimated job-offer distribution over wages and separation rates used in the numerical implementation. Panel (b) shows the conditional distribution of separation rates for different levels of wage.

Figure 1.A.11. Earnings and consumption inequality

(a) Earnings variance **(b)** Consumption variance

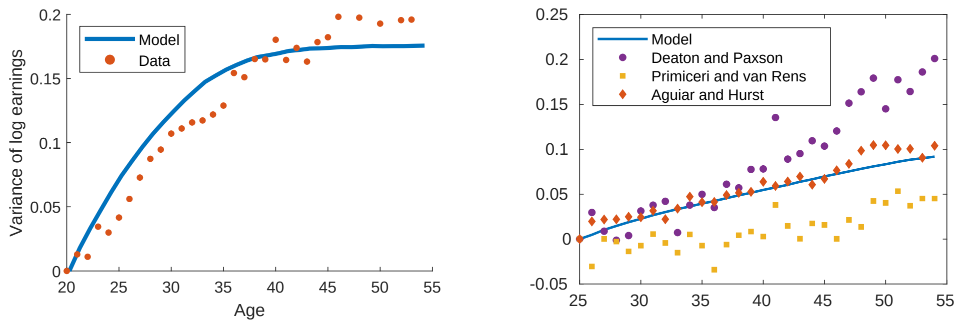

Notes: Panel (a) shows the variance of log earnings in the PSID data, normalized to 0 at age 20. The blue line is the model profile, while the red dots show the estimated empirical profile from the PSID data. Panel (b) shows the variance of log consumption from Deaton and Paxson (1994), Primiceri and Van Rens (2009), and Aguiar and Hurst (2013), normalized to 0 at age 25.

cumulation decision. As we discuss in Section 1.4.1, the endogenous human capital accumulation is key for the question of this paper because it provides a mechanism to transform transitory differences in search outcomes into persistent earnings differences. Figure 1.A.11b demonstrates that the model also aligns with the empirical estimates of the life cycle increase in the variance of log consumption. Empirical estimates of consumption variance differ across studies (Deaton and Paxson, 1994; Primiceri and Van Rens, 2009; Aguiar and Hurst, 2013), and the model falls in the middle of the range of existing estimates. Comparing Figures 1.A.11 and 1.A.11b, we also note that the increase in the variance of consumption is roughly one-third lower than that for earnings. Hence, consumption is partly insulated from earnings dynamics.

1.A.10 Comparison of life-cycle profiles in SCF and CPS

Figure 1.A.12 compares the life-cycle profiles for earnings and tenure in the Current Population Survey and the Survey of Consumer Finances. To be consistent with the construction of the CPS tenure profiles, we use labor market information on household heads and spouses. We find that evidence from both data sources is consistent and shows similar life-cycle patterns for earnings and tenure.

Figure 1.A.12. Earnings and tenure in SCF and CPS

(a) Earnings **(b)** Tenure

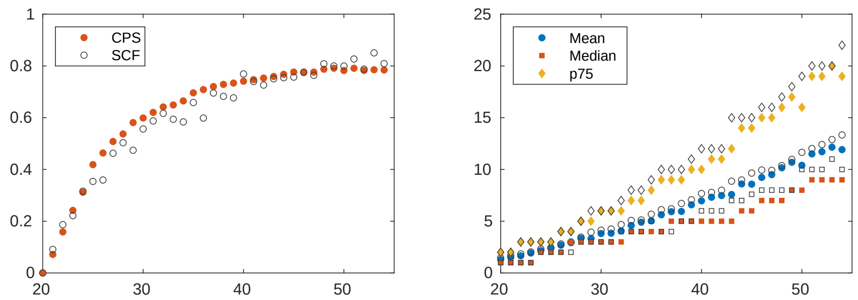

Notes: This figure compares life-cycle profiles of earnings and tenure in the SCF and CPS data. Panel (a) shows the life-cycle profile of mean log earnings, normalized to 0 at age 20. Panel (b) shows the mean, median, and 75th percentile of tenure. Filled dots show the SCF profiles; unfilled dots are the CPS profiles.

1.A.11 Cross-sectional distributions of tenure and the number of employers

Figure 1.A.13 shows the cross-sectional distribution of tenure and the number of employers for which a worker has worked for at least one year during her working life. We combine all workers and show the corresponding distribution using histograms. When pooling data from the model, we assume that each age group has the same share in the pooled sample.

1.A.12 A model without heterogeneity in job stability

This section presents results for an economy in which jobs do not differ in terms of job stability. The structure of the model is the same as in the main part of the paper, with the only difference that now job offers that workers randomly draw from the job-offer distribution differ only across the wage dimension. In contrast to the baseline model, the job separation rate is exogenous, and as a result, all workers of

Figure 1.A.13. Cross-sectional distribution of tenure and number of employers

(a) Tenure distribution **(b)** Number of employers

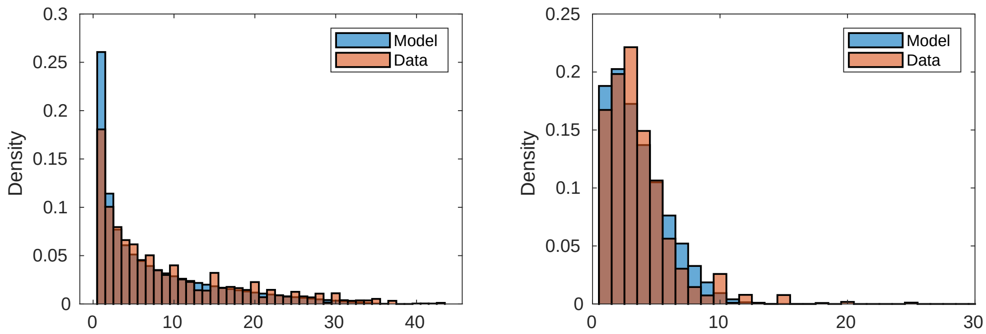

Notes: This figure shows the distribution of tenure and the number of employers from the SCF and the model when all ages are pooled together. Red bars are the SCF data; blue bars are the model equivalent. In line with the SCF design, only employment spells with a duration of at least a year are used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

the same age have an equal probability of becoming nonemployed, as in Michelacci and Ruffo (2015). The separation rate that workers of a given age face is the same as the average separation rate in the baseline model. Consequently, the life-cycle profiles for the separation rate are identical in both models.

We find that this alternative model significantly underperforms the baseline model when it comes to matching several documented empirical facts. The model in which all workers face the same job loss probability produces a tenure distribution that does not match the documented empirical distribution. As shown in Figure 1.A.14a, the distribution of tenure is much more compressed, and the life-cycle increase in tenure is substantially lower compared to the baseline model and the empirical evidence. Additionally, we also find that this model performs poorly in capturing the empirical relationship between wealth accumulation and job stability. In Figure 1.A.14b, we show binned scatter plots of wealth-to-income ratios and tenure after controlling for age effects. As is clearly visible, the dispersion in tenure is lower than the empirically observed one, and the slope of the model-based relationship deviates from the empirical one.

Furthermore, without any cross-sectional heterogeneity in the separation rate, the model also cannot replicate large and persistent earnings losses following displacement. In Figure 1.A.15, we show the cost of displacement for the model without cross-sectional heterogeneity in the separation rate. Contrary to the results from the baseline model in Figure 1.11, we find that earnings losses following displacement largely disappear ten years after displacement when jobs do not differ in terms of their job stability. Consistent with less persistent negative effects of nonemployment on earnings, we also find that consumption declines substantially less com-

Figure 1.A.14. Tenure and wealth

(a) Tenure **(b)** Wealth-to-income ratios and tenure

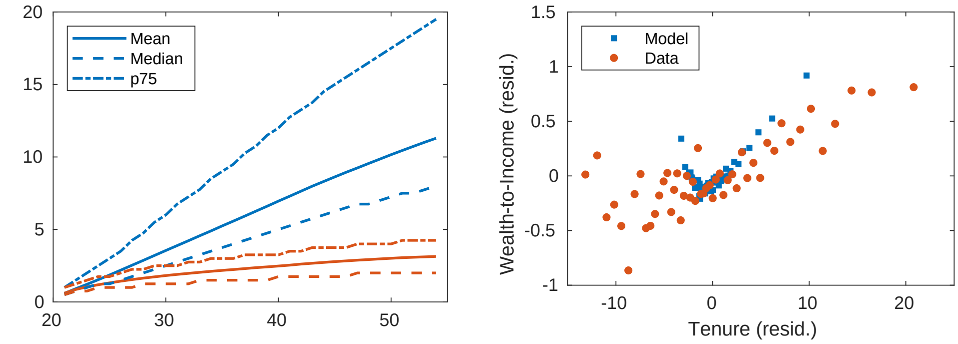

Notes: Panel (a) compares the life-cycle evolution of the distribution of tenure for the model without heterogeneity in separation rates and for the baseline model. The red profiles are for the model without heterogeneity in job stability, and the blue profiles correspond to the baseline model. Panel (b) shows the relationship between wealth-to-income ratios and tenure after nonparametrically controlling for age.

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

Figure 1.A.15. Cost of displacement without heterogeneity in job stability

(a) Earnings **(b)** Consumption and wealth

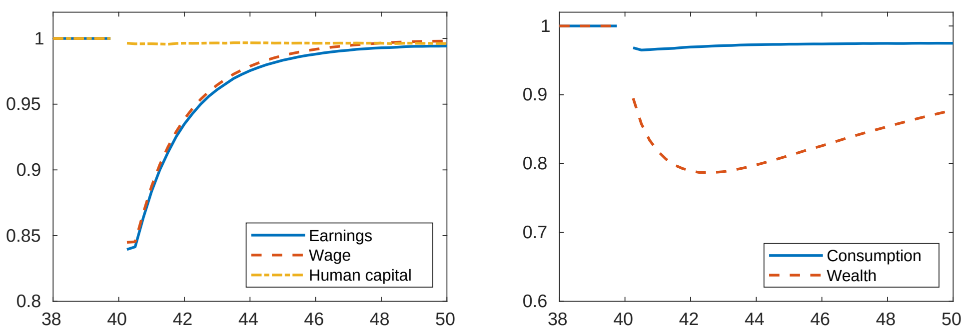

Notes: This figure shows the evolution of earnings, consumption, and wealth of workers who become unemployed at age 40 relative to the control group. Prior to displacement, both groups are identical.

1.A.13 Life-cycle earnings dynamics

Results presented in Section 1.3.2 demonstrate that the model matches the life-cycle profiles for means and variances. Here we provide additional evidence that the model also provides a good fit along other dimensions. To explore the fit for earnings dynamics, we compare how the model-implied earnings dynamics align with statistical representations of earnings processes as typically estimated in applied work and used to parametrize exogenous earnings dynamics in consumption-saving models. Such a description of earnings dynamics by a reduced-form statistical representation allows for a straightforward comparison of earnings dynamics between model and data. First, we perform a standard decomposition of earnings dynamics into a permanent and transitory component and estimate the variances of the innovation terms (Meghir and Pistaferri, 2004; Blundell, Pistaferri, and Preston, 2008; Heathcote, Perri, and Violante, 2010). Second, as emphasized in Guvenen et al. (2019), we look at the higher moments of the distribution of earnings growth. Third, we decompose earnings growth into contributions from human capital accumulation and job switching and demonstrate that the model aligns with the evidence on early career wage growth by Topel and Ward (1992). We close with a discussion of the model predictions for frictional wage dispersion.

In the first case, we estimate the variance of the permanent component of earnings dynamics using simulated earnings series from the model aggregated to an annual frequency. We apply the identification approach, as in Blundell, Pistaferri, and Preston (2008), to the simulated data. We estimate a variance of the permanent component of 0.025 that falls well within the range of empirical estimates. Blundell, Pistaferri, and Preston (2008) estimate time-varying variances of the permanent component ranging from 0.01 to 0.03 for the period from 1980 to 1990.[20] Empirical estimates for the variance of transitory shocks are harder to compare as they also comprise the contribution from measurement error that is likely substantial in the data, so it is not surprising that our finding is that the empirical estimates for the variance of the transitory shocks (0.03-0.05) are substantially larger than the model-implied estimate (0.016). We interpret the difference as the contribution from measurement error but also unmodeled earnings components such as bonuses and overtime pay.

In a second step, we consider the findings by Guvenen et al. (2019), who emphasize that earnings growth rates are not normally distributed but exhibit large negative skewness and high excess kurtosis. As has been demonstrated by Hubmer (2018), these patterns can be well explained by a life-cycle version of the job ladder model with a human capital process. The same also applies to our model. Figure 1.A.16 shows the distribution of one-year earnings growth rates in the model with

20. Heathcote, Perri, and Violante (2010) provide a detailed discussion of different estimation approaches. We use the estimation as a reduced-form description of earnings dynamics without requiring the process to be the true underlying process or estimates to be unbiased. See Daly, Hryshko, and Manovskii (2016) for further discussion.

a superimposed normal distribution that has the same standard deviation. Earnings changes in our model are left-skewed and strongly clustered around 0.

Figure 1.A.16. One-year earnings changes

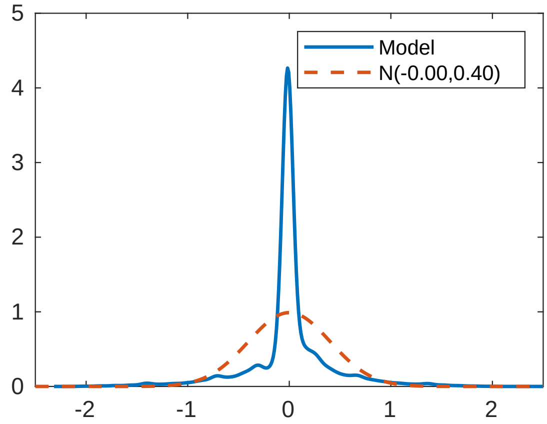

Notes: Figure shows the estimated kernel density of the model-based one-year earnings changes superimposed on Gaussian densities with the same standard deviation. Earnings growth net of average age effect shown. Kernel density estimation with a bandwidth of 0.05 used.

Last, we use the model to decompose life-cycle earnings growth and the increase in the variance into a component from search for higher wages and a component from human capital accumulation. Figure 1.A.17 shows the decomposition for mean and variance and highlights human capital accumulation as the key driver of life-cycle earnings dynamics. This decomposition aligns well with empirical evidence. Topel and Ward (1992) provide estimates for the contribution of employer switching to wage growth after labor market entry. They find that search for better-paying employers accounts for about one-third of wage growth within the first ten years in the labor market. Looking at the decomposition of mean earnings in Figure 1.A.17a, we find, in line with their results, that climbing the wage ladder is an important driver of early career wage growth. Between ages 20 and 30, it accounts for roughly one-third of earnings growth, close to the Topel and Ward (1992) estimate. After age 30, wage growth from job search flattens out, in line with a slowdown in employer switching (Figure 1.5). Human capital investment accounts for almost the entire increase in earnings once most workers have found stable jobs.

For the increase in the variance, we find a similar decomposition. At age 20, all workers start from the same level of human capital, so differences in entry wages account for all of the dispersion in earnings. Over time, workers climb the wage ladder leaving less well-paid jobs and accept better-paid jobs, which leads to wage compression and contributes negatively to the increase in life-cycle earnings inequality. Workers in well-paid jobs receive fewer opportunities to climb the wage ladder as many jobs offer lower wages, and therefore, these workers are more likely to

Figure 1.A.17. Decomposing earnings dynamics over the life cycle

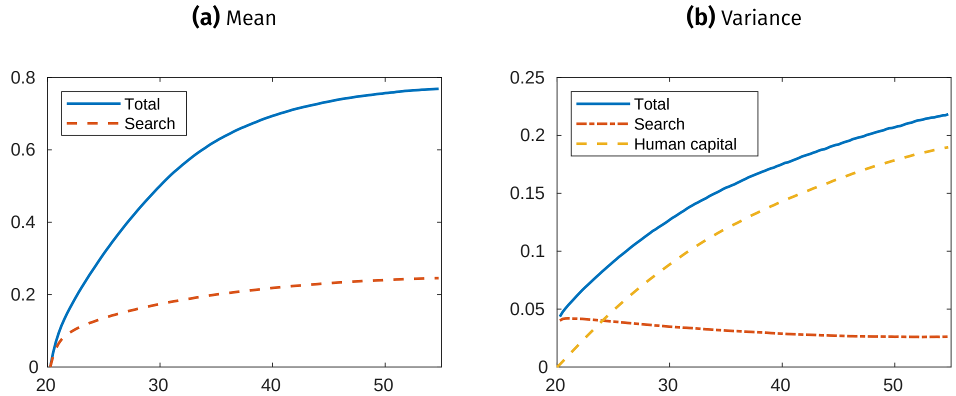

Notes: Panel (a) shows the profile of mean of log earnings and the contribution of the wage component to the growth of earnings over the life cycle. Panel (b) shows the contribution of human capital dispersion and wage dispersion to the overall earnings dispersion over the life cycle.

stay with their current employer. As a result, initial wage differences decrease. This mechanism also highlights the general challenge when trying to account for wage dispersion relying on employer differences alone. As our decomposition shows, differences in human capital accumulation are the driver of rising earnings inequality over the life cycle. The covariance between human capital and wages (not shown) is small but positive and contributes little to earnings dispersion in the model. At age 40, the contribution of the covariance accounts for about 10% of the search component. This decomposition with a small contribution from search frictions (frictional wage dispersion) is consistent with results in Hornstein, Krusell, and Violante (2011) and Hagedorn and Manovskii (2010) that point toward low levels of frictional wage dispersion in the data. Bayer and Kuhn (2018) decompose the increase in life-cycle wage dispersion using German administrative data and also find a negligible contribution of employer differences to the life-cycle increase in wage dispersion.

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

1.A.14 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 (Huggett, 1993; Aiyagari, 1994). At their core, these models provide a mapping from earnings dynamics to wealth accumulation. Our model combines this consumption-saving block with a labor market block endogenizing earnings dynamics. This combination of building blocks suggests that a comparison of the joint distribution of income and wealth and the implied cross-sectional wealth dynamics induced by the model's earnings dynamics are particularly well suited to compare how our model performs in accounting for the data. Unlike panel data on consumption and income dynamics, the approach to compare the joint distribution of income and wealth has the additional advantage that these data are easily observed in datasets such as the SCF.

For this comparison, we split households into income and wealth quintiles. We consider the joint distribution by comparing how households are distributed across wealth quintiles conditional on their income quintile. For wealth dynamics, we compare how households move across wealth quintiles over time. Panels (a) through (c) in Figure 1.A.18 show selected conditional distribution functions for wealth by income quintile from model and SCF data.[21] Panels (d) through (f) show the conditional distribution function for starting from a given wealth quintile over all wealth quintiles five periods in the future (i.e., we show rows of a five-step wealth transition matrix).[22] For data on the joint distribution, we rely on the SCF data. We follow previous research (Díaz-Giménez, Glover, and Ríos-Rull, 2011; Kuhn, Schularick, and Steins, 2020) and rely on PSID data from 1984 to 1999 to trace out individual-level wealth dynamics. The repeated cross sections of the SCF data prevent such an analysis.[23]

Looking at the joint distribution in Figures 1.A.18a to 1.A.18c, we find that the model aligns closely with the SCF data for households ages 40 to 50. We focus on a single age group to alleviate concerns regarding the age structure in the model relative to the age structure in the data. The fit of this untargeted dimension is very good. The most notable difference between the model and data is that too many households from the fifth income quintile are at the bottom of the wealth

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

22. We report all conditional probabilities in Table 1.A.2.

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

Figure 1.A.18. Joint distribution of income and wealth and wealth dynamics

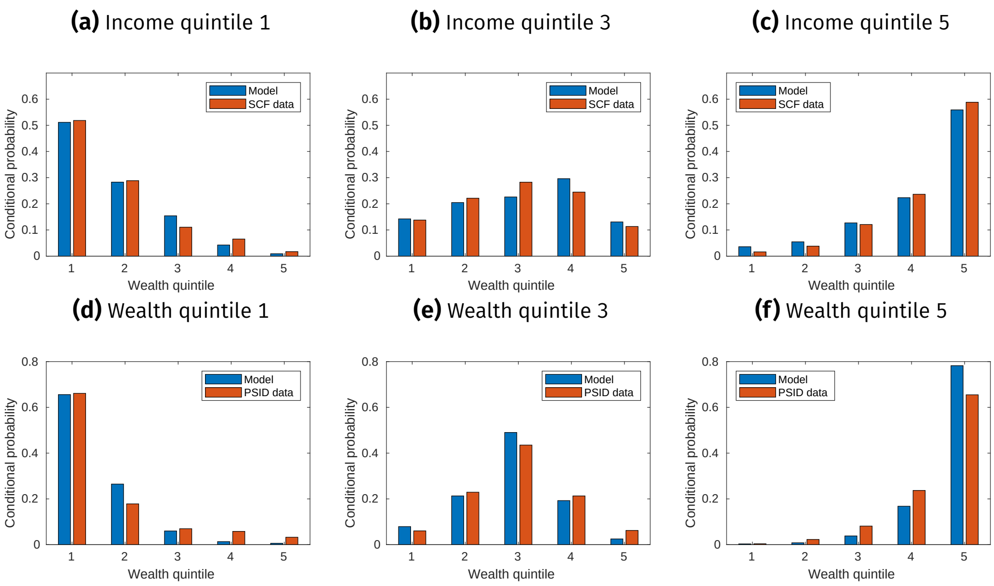

Notes: Panels (a) through (c) show the joint distribution of income and wealth of households ages 40 to 50. We split households into quintiles along the income and wealth dimension and show the conditional cumulative distribution functions for wealth by income quintile from model and SCF data. Panels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

distribution. The model hence generates too many income-rich but wealth-poor households, but overall the distributions from the model and data align very closely. While income, being a flow, might change quickly, wealth as a stock moves much more slowly. The close fit of the joint distribution therefore suggests that the model-implied wealth dynamics compare favorably to their data counterpart.

Looking at wealth dynamics in Figures 1.A.18d to 1.A.18f, we find that the model matches closely the observed wealth dynamics in PSID data. The PSID surveys wealth only every five years during this time period, so we focus on five-year transition probabilities. We observe a high persistence of households' position along the wealth distribution. More than 80% of households from the first wealth quintile remain within the two bottom wealth quintiles over a five-year horizon. Similarly at the top, only about 30% of households from the top quintile end up in a lower quintile five years later. Comparing the model and data, we find that the model produces slightly too little wealth mobility at the bottom and the top. The middle of the distribution is closely matched. It is important to note that imputation of wealth information in the PSID likely leads to overstating the estimates of wealth mobility in the PSID data.

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

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

Table 1.A.2. Wealth transition matrix, age 40

	W0-20	W20-40	W40-60	W60-80	W80-100
			Model		
W0-20	0.66	0.26	0.06	0.01	0.01
W20-40	0.25	0.45	0.24	0.05	0.01
W40-60	0.08	0.21	0.49	0.19	0.03
W60-80	0.02	0.06	0.16	0.57	0.19
W80-100	0.00	0.01	0.04	0.17	0.78
			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

References

Acabbi, Edoardo Maria, Andrea Alati, and Luca Mazzone. 2024. "Human Capital Ladders, Cyclical Sorting, and Hysteresis." *Available at SSRN 4068858.* [7]

Aguiar, Mark, and Erik Hurst. 2013. "Deconstructing life cycle expenditure." *Journal of Political Economy* 121 (3): 437–92. [5, 52]

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *Quarterly Journal of Economics* 109 (3): 659–84. [7, 59]

Barnette, Justin. 2020. "Wealth After Job Displacement." [32]

Bartscher, Alina K, Moritz Kuhn, and Moritz Schularick. 2020. "The college wealth divide: Education and inequality in America, 1956-2016." *Available at SSRN 3587685.* [14]

Bayer, Christian, and Moritz Kuhn. 2018. "Which Ladder to Climb? Wages of workers by job, plant, and education." [58]

Bewley, Truman. undated. "Interest bearing money and the equilibrium stock of capital." Working paper. mimeo. [7]

Bilal, Adrien. 2019. *The Geography of Unemployment.* [8]

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5): 1887–921. [5, 24, 56]

Bricker, Jesse, Lisa J Dettling, Alice Henriques, Joanne W Hsu, Lindsay Jacobs, Kevin B Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A Windle. 2017. "Changes in US family finances from 2013 to 2016: Evidence from the Survey of Consumer Finances." *Fed. Res. Bull.* 103: 1. [43, 44]

Burdett, Kenneth, and Dale T Mortensen. 1998. "Wage differentials, employer size, and unemployment." *International Economic Review,* 257–73. [7]

Cajner, Tomaz, Ilhan Güner, and Toshihiko Mukoyama. 2020. "Gross Worker Flows over the Life Cycle." [7]

Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud. 2023. "Subjective Earnings Risk." Working Paper, Working Paper Series 31019. National Bureau of Economic Research. [8]

Couch, Kenneth A, and Dana W Placzek. 2010. "Earnings losses of displaced workers revisited." *American Economic Review* 100 (1): 572–89. [30]

Daly, Moira, Dmytro Hryshko, and Iourii Manovskii. 2016. "Improving the measurement of earnings dynamics." Working paper. National Bureau of Economic Research. [56]

Davis, Steven J., Till von Wachter, Robert E. Hall, and Richard Rogerson. 2011. "Recessions and the Costs of Job Loss [with Comments and Discussion]." *Brookings Papers on Economic Activity,* 1–72. [13, 30]

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal choice and inequality." *Journal of political economy* 102 (3): 437–67. [52]

Díaz-Giménez, Javier, Andy Glover, and José-Víctor 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 (1): 2–31. [59]

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2019. "What do data on millions of US workers reveal about life-cycle earnings dynamics?" *working paper,* no. 710. [5, 8, 9, 24, 56]

Hagedorn, Marcus, and Iourii Manovskii. 2010. "Search frictions and wage dispersion." *Manuscript, University of Zurich and University of Pennsylvania.* [58]

Hall, Robert E. 1982. "The Importance of Lifetime Jobs in the US Economy." *American Economic Review* 72 (4): 716–24. [3, 4, 9, 11, 15, 45]

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13 (1): 15–51. Special issue: Cross-Sectional Facts for Macroeconomists. [24, 37, 56]

Hornstein, Andreas, Per Krusell, and Giovanni L Violante. 2011. "Frictional Wage Dispersion in Search Models: A Quantitative Assessment." *American Economic Review* 101 (7): 2873–98. [22, 24, 58]

Hubmer, Joachim. 2018. "The job ladder and its implications for earnings risk." *Review of Economic Dynamics* 29: 172–94. [7, 24, 51, 56]

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of economic Dynamics and Control* 17 (5-6): 953–69. [7, 59]

Iacono, Roberto, and Marco Ranaldi. 2020. "The wage curve across the wealth distribution." *Economics Letters* 196: 109580. [13]

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709. [5, 13, 30]

Jarosch, Gregor. 2015. "Searching for Job Security and the Consequences of Job Loss." [4, 6–9, 15, 17, 30, 39]

Jung, Philip, and Moritz Kuhn. 2018. "Earnings Losses and Labor Mobility Over the Life Cycle." *Journal of the European Economic Association,* jvy014–jvy014. 10.1093/jeea/jvy014. [4, 6–9, 11, 15, 17, 21, 22, 30, 31, 38, 39, 51]

Kaas, Leo, Etienne Lalé, and Nawid Siassi. 2023. "Job Ladder and Wealth Dynamics in General Equilibrium." Working paper. [7]

Kaplan, Greg, and Giovanni L Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2 (4): 53–87. [5]

Karahan, Fatih, and Serdar Ozkan. 2013. "On the persistence of income shocks over the life cycle: Evidence, theory, and implications." *Review of Economic Dynamics* 16 (3): 452–76. [8]

Krebs, Tom, Moritz Kuhn, and Mark LJ Wright. 2015. "Human capital risk, contract enforcement, and the macroeconomy." *American Economic Review* 105 (11): 3223–72. [16]

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2017. "Gross Worker Flows over the Business Cycle." *American Economic Review* 107 (11): 3447–76. [7]

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Sahin. 2010. "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations." *Review of Economic Studies* 77 (4): 1477. [7]

Kudlyak, Marianna, and Fabian Lange. 2014. "Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories." *IZA DP,* no. 8663. [21]

Kuhn, Moritz. 2013. "Recursive Equilibria in an Aiyagari-Style Economy with Permanent Income Shocks." *International Economic Review* 54 (3): 807–35. [32]

Kuhn, Moritz, and José-Víctor 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 (1): 2–73. [43, 44]

Kuhn, Moritz, Moritz Schularick, and Ulrike Steins. 2020. "Income and Wealth Inequality in America, 1949-2016." *Journal of Political Economy.* [43, 59]

Larkin, Kieran. 2019. "Job risk, separation shocks and household asset allocation." In *2019 Meeting Papers.* 1058. Society for Economic Dynamics. [4, 7, 15, 17, 39]

Lise, Jeremy. 2012. "On-the-Job Search and Precautionary Savings." *Review of Economic Studies* 80 (3): 1086–113. [7, 51]

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review* 100 (4): 1432–67. [8, 30]

Meghir, Costas, and Luigi Pistaferri. 2004. "Income variance dynamics and heterogeneity." *Econometrica* 72 (1): 1–32. [24, 56]

Michelacci, Claudio, and Hernán Ruffo. 2015. "Optimal Life Cycle Unemployment Insurance." *American Economic Review* 105 (2): 816–59. [7, 54]

Morchio, Iacopo. 2020. "Work Histories and Lifetime Unemployment." *International Economic Review* 61 (1): 321–50. [8, 9]

Mortensen, Dale T, and Christopher A Pissarides. 1994. "Job creation and job destruction in the theory of unemployment." *review of economic studies* 61 (3): 397–415. [7]

OECD. 2015. *Pensions at a Glance 2015.* 376. [20]

Pfeffer, Fabian T, Robert F Schoeni, Arthur Kennickell, and Patricia Andreski. 2016. "Measuring wealth and wealth inequality: Comparing two US surveys." *Journal of economic and social measurement* 41 (2): 103–20. [59]

Pinheiro, Roberto, and Ludo Visschers. 2015. "Unemployment risk and wage differentials." *Journal of Economic Theory* 157: 397–424. [4, 8, 15, 17]

Primiceri, Giorgio E, and Thijs Van Rens. 2009. "Heterogeneous life-cycle profiles, income risk and consumption inequality." *Journal of monetary Economics* 56 (1): 20–39. [52]

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49. [20]

Stevens, Ann Huff. 1997. "Persistent effects of job displacement: The importance of multiple job losses." *Journal of Labor Economics* 15 (1, Part 1): 165–88. [31]

Topel, Robert H, and Michael P Ward. 1992. "Job Mobility and the Careers of Young Men." *Quarterly Journal of Economics* 107 (2): 439–79. 10.2307/2118478. [5, 24, 56, 57]

Chapter 2

Optimal Progressive Pension Systems and Heterogeneous Labor Market Risk*

2.1 Introduction

Heterogeneity in employment stability is a salient feature of the labor market. While some workers have lifetime jobs, others have unstable work histories and frequently experience unemployment. Employment stability is key for individual welfare, as interrupted work histories and earnings losses following job displacements reduce pension entitlements and hinder the accumulation of life-cycle savings. Early work by Hall (1982) and a growing literature (Jung and Kuhn, 2019; Molloy, Smith, and Wozniak, 2020; Gregory, Menzio, and Wiczer, 2021; Ahn, Hobijn, and Şahin, 2023; Jarosch, 2023; Kuhn, Nam, and Ploj, 2025) have explored this persistent heterogeneity. However, we still know little about its implications for optimal policy, in particular, the design of the pension scheme. This paper fills this gap.

In this paper, I develop a life-cycle model with heterogeneity in employment stability, endogenous human capital accumulation, and consumption-saving and retirement decisions. I calibrate the model to the U.S. economy to derive the optimal degree of pension progressivity. I also examine how macroeconomic shifts toward higher employment stability, as observed in the United States since the 1990s affect optimal pension design. I find that abolishing the Social Security cap and an increase in the degree of pension progressivity is optimal. The optimal pension system achieves a welfare gain of 0.75% of lifetime consumption for labor market entrants

* I thank Moritz Kuhn for invaluable advice, guidance, and encouragement throughout this project. I also thank Mariacristina De Nardi, Jonathan Heathcote, Marco Bassetto, Joseph Mullins, Camille Landais, Thomas Hintermaier, Eugenia Gonzalez-Aguado, Johan Sæverud, and Hans Fehr as well as participants of numerous conferences and seminars for their helpful feedback and comments. I thank the DFG for financial support (DFG No. 433368816). I also gratefully acknowledge financial support from the Bonn Graduate School of Economics and the DFG research training group 2281 "The Macroeconomics of Inequality." The usual disclaimer applies.

relative to the current U.S. pension system. Following the observed macroeconomic shift in the employment stability distribution, the optimal pension system becomes less progressive. However, the welfare gain from implementing the optimal pension system remains sizeable in the economy with higher employment stability.

The life-cycle model features a frictional labor market where workers search both on and off the job, make consumption-saving decisions in incomplete financial markets, and invest in risky human capital. When workers become eligible for pension benefits, they choose their retirement age. Workers are heterogeneous with respect to their job-separation rates at each point in time of their life cycle, which builds on previous literature on heterogeneity in employment stability (Pinheiro and Visschers, 2015; Jarosch, 2023; Kuhn, Nam, and Ploj, 2025). The novel feature of this model is the incorporation of heterogeneity in employment stability while treating human capital accumulation and retirement as endogenous. Endogenous human capital accumulation is an important determinant of earnings (Becker, 1964; Ben-Porath, 1967) and is key to explain increasing life-cycle earnings inequality (Hubmer, 2018; Jung and Kuhn, 2019). The interaction between human capital and employment stability amplifies earnings losses upon job loss into long-run consequences (Jarosch, 2023). Workers get stuck in bad, unstable employment as they fall down the job ladder and fail to accumulate human capital and life-cycle savings. These long-run consequences on earnings put forward the importance of heterogeneity in employment stability in analyzing optimal pension systems. In the absence of progressivity, the pension system transmits interrupted work histories into the old age as pension entitlements depend on the pre-retirement history of earnings. I calibrate the model to the U.S. economy and the Social Security system, specifically, the empirical profiles of labor market transition rates, earnings profile (mean and variance), tenure distribution, and wealth-to-income ratios, and show that the model matches the data remarkably well.

In analyzing the optimal pension system, I focus on the classical trade-off between providing insurance and preserving incentives. A progressive pension system, which offers higher replacement rates for low-income workers, insures against unstable employment and low lifetime earnings. However, this comes at the cost of reducing the return of human capital investment, discouraging workers to invest in human capital. The progressivity level also affects retirement incentives of workers as the amount of pension wealth is one of the most important components in forming retirement decisions. Higher pension progressivity reduces labor supply at the extensive margin as low-income workers retire earlier. Additionally, financing pensions through payroll taxes introduces distortions, especially for young, low-income individuals facing borrowing constraints. While higher progressivity increases future pension wealth, higher payroll taxes can depress consumption for those unable to smooth income over time. Thus, increasing the provision of pension coverage through higher payroll taxes may lower welfare for liquidity-constrained young workers.

Using the calibrated life-cycle model, I analyze the optimal design of the U.S. pension system. The model shows that heterogeneity in employment stability translates into large inequality in labor market outcomes and is a key driver of inequality in earnings and consumption. I show that increasing the progressivity of pension benefits and abolishing the U.S. Social Security cap is optimal. The optimal benefit schedule is close to a benefit floor so that workers at the 25th percentile of the lifetime income distribution face an increase from 74% to 88% in their replacement rate. For workers at the 75th percentile, the replacement rate decreases by 6 percentage points. The abolishment of the Social Security cap increases the tax revenue from high-income workers so that the payroll tax rate decreases by 0.3 percentage points. The optimal pension system induces a welfare gain of 0.75% in terms of lifetime consumption for labor market entrants. I show that in a model without heterogeneity in employment stability conditional on age, the optimal progressivity level of the pension system is much lower compared to the model with heterogeneity in employment stability.

The optimal pension system with a higher progressivity level reduces consumption inequality over the life cycle by nearly one-third relative to the current U.S. Social Security system. It provides insurance to workers who suffer from frequent career interruptions and low lifetime earnings. While increased progressivity introduces distortions—reducing the average human capital stock by 1.5% and lowering the average retirement age by 0.4% of total years worked—these effects are modest and do not outweigh the insurance benefits. This is for two reasons. First, young workers do not reduce their human capital investment as they strive to accumulate human capital to increase their prospective labor earnings. As the return to human capital is largely realized in early and mid-career stages, a reduction in pension-related returns in the distant future has limited impact. This intuition aligns with Michelacci and Ruffo (2015), who show that an optimal age-dependent unemployment insurance system should provide higher insurance to young workers as this creates limited distortions due to the high marginal returns from human capital of the young. Second, changes in retirement decisions depend on the accumulated stock of human capital. Less productive workers decide to retire earlier, whereas productive workers postpone their retirement in face of the increase in pension progressivity. The delay of retirement of productive workers partly offsets the distortion of retirement incentives on low-productivity workers such that the aggregate distortionary effect remains small.

I also analyze the transition from the current economy to the economy with the optimal pension system, assuming that the progressivity level is gradually increased over a 40-year horizon. While all labor market entrants benefit from the optimal progressive pension system as they are ex-ante identical, the welfare effects vary among older workers based on their employment histories. At age 50, workers with low employment stability (bottom 25% of employment-stability distribution) experience a welfare gain of 0.6%, while those with high employment stability (top 25%

of employment-stability distribution) see a welfare gain of 0.3%. Importantly, the optimal policy leads to positive welfare gains for the vast majority of workers except for those at the very top income. By abolishing the Social Security cap, the tax revenue from high-income workers increases, enabling a reduction in the payroll tax rate and leading to welfare improvements.

In the last part of this paper, I analyze the consequences of a shift in the employment-stability distribution in the U.S. labor market on the optimal design of progressive pension systems. This study is motivated by the wide range of empirical findings in the literature that employment stability in the U.S. labor market has been shifting in the last few decades (Hyatt and Spletzer, 2013; Pries and Rogerson, 2019; Molloy, Smith, and Wozniak, 2020). In particular, empirical studies have consistently found an increase in employment stability for the United States since the late 1990s. While holding the job-separation rate of the most stable job constant, I reduce the separation rate of the most unstable job for the economy with higher employment stability. In the model, a shift towards higher employment stability reduces average job separations in the economy. However, the earnings inequality remains almost unchanged. As jobs get more stable, the earnings dispersion generated by the job ladder and risky human capital process becomes stronger so that on average, the earnings inequality does not decrease with higher employment stability. This model implication is in line with the literature on the recent evolution of income inequality in the United States (Braxton et al., 2021; Guvenen et al., 2021; Heathcote et al., 2023) and reconciles the fact that we observe lower labor market dynamism and at the same time stable or increasing earnings inequality. As higher employment stability does not lead to significant changes in earnings inequality, the optimal pension system remains largely unchanged. I only find that the implied level of progressivity by the optimal pension system slightly decreases. The replacement rate of workers at the 25th percentile of the lifetime income distribution decreases from 88% to 84% compared to the economy before increasing employment stability, while the replacement rate slightly increases for workers at the 75th percentile.

Importantly, the welfare gain from implementing the optimal pension system becomes even higher as employment stability increases. In the economy with more stable jobs, workers still value the insurance provided by the progressive pension system as heterogeneity in employment stability and income inequality are still large, but the cost of introducing the optimal pension system decreases. As average earnings increase with higher employment stability, the payroll tax rate can be lowered compared to the baseline economy. The cost of implementing a more progressive pension system is mitigated in the economy with higher employment stability and hence, the welfare gain from implementing the optimal pension system becomes higher. I show that the importance of the optimal pension system does not diminish in the economy with higher employment stability as the degree of heterogeneity in employment stability is still large in today's economy.

The remainder of this paper is structured as follows. Section 2.2 relates this paper to the existing literature. Section 2.3 presents the life-cycle model, followed by the baseline calibration in Section 2.4. Section 2.5 analyzes the effects of heterogeneity in employment stability on inequality in labor market outcomes and its life-cycle consequences. I analyze the optimal pension system for the baseline economy in Section 2.6. Section 2.7 explores the consequences of a shift in employment-stability distribution in the U.S. labor market on the optimal design of pension systems. Section 2.8 concludes.

2.2 Related literature

This paper contributes to three important strands of macroeconomic literature. First, by analyzing the policy implications of heterogeneity in employment stability, it contributes to the growing strand of the literature that provides evidence for heterogeneity in employment stability and labor market dynamics (Molloy, Smith, and Wozniak, 2020; Gregory, Menzio, and Wiczer, 2021; Ahn, Hobijn, and Şahin, 2023). Among others, Jung and Kuhn (2019) find considerable heterogeneity in employment stability for the United States and show that it is key to explain persistent earnings losses after job displacement in a life-cycle model with search. Using German Social Security data and a model with heterogeneity in employment stability, Jarosch (2023) points out that the interaction between human capital and employment stability is crucial to explain the observed earnings losses. Morchio (2020) shows that unemployment is concentrated among certain groups of workers. Workers significantly differ in their job-separation rates and such differences are large from the start of the career. Kuhn, Nam, and Ploj (2025) also provide evidence for employment-stability heterogeneity in the U.S. labor market. Based on a life-cycle model, they show that employment stability at labor market entry significantly affects income and consumption level over the whole life. I follow the strand of the literature that assumes that heterogeneity in employment stability is generated due to differences across jobs (Jung and Kuhn, 2019; Larkin, 2019; Bilal, 2023; Jarosch, 2023). Another strand of the literature considers heterogeneity in employment stability as worker types rather than job types (Hall and Kudlyak, 2019; Gregory, Menzio, and Wiczer, 2021; Ahn, Hobijn, and Şahin, 2023). The model implications in this paper are, however, also consistent with this strand of the literature. This model generates very stable and unstable career paths across workers and in fact, a worker with a very stable job is observationally indistinguishable from a very stable worker type within the model. Hence, differences in employment stability in this paper can be interpreted as worker types as in the above-mentioned empirical approaches.

Second, this paper relates to the literature on unemployment risk and inequality in lifetime earnings. Using the life-cycle model, I discuss the interaction between heterogeneity in employment stability and how unstable career paths affect lifetime earnings of workers, which consequently affects their pension entitlements. These results complement the large strand of empirical literature on the relationship between inequality in lifetime earnings and labor market risks. For example, Bonhomme and Robin (2009) explore inequality in employment and lifetime earnings for France and find that unemployment is a key driver of lifetime earnings inequality. Similarly, Bowlus and Robin (2012) show that unemployment mobility has important consequences on lifetime earnings inequality. Bönke, Corneo, and Lüthen (2015) find that inequality in lifetime earnings has been increasing in Germany over time, with a larger increase for individuals at the bottom of the earnings distribution. Longer unemployment durations are one of the major sources of the increase in lifetime earnings inequality. Haan, Kemptner, and Prowse (2019) find that heterogeneity in labor market outcome is a major source of inequality in lifetime earnings for workers with same abilities.

Lastly, this study contributes to the literature on optimal design of pension systems by applying a life-cycle framework which explicitly takes heterogeneity in employment stability into account. The application of a life-cycle model is a frequently used approach in the previous macroeconomic literature on optimal pension systems. Beginning with the work by Feldstein (1974), a large body of literature investigates the effects of pension systems in a life-cycle framework. For example, Hubbard and Judd (1987) consider a model with capital-market imperfections to show that the introduction of Social Security does not increase an economy's welfare as much as in an economy without borrowing constraints. A more recent work is the study by Fehr and Habermann (2008) who analyze the welfare effects of introducing a more redistributive pension system to Germany. They find that the positive liquidity and income insurance effects are large enough such that an increase in progressivity enhances welfare despite its distortive effects on labor supply. Fehr, Kallweit, and Kindermann (2013) analyze the impact of higher progressivity of the German pension system in a model with disability risk and labor supply decisions at the extensive margin. Golosov et al. (2013) take the current form of the U.S. retirement benefit function and determine its optimal structure, assuming that workers have heterogeneous productivities. They find that a more redistributive system increases the economy's welfare by a considerable amount. O'Dea (2018) finds that introducing means-tested income floors for the elderly is beneficial. Some other prominent studies on pension systems are, for example, Krueger and Pischke (1992); De Nardi, Imrohoroglu, and Sargent (2001); Coile and Gruber (2001); Cremer, Lozachmeur, and Pestieau (2004); Gruber and Wise (2004); Bloom et al. (2007); Haan and Prowse (2014); and Borella, De Nardi, and Yang (2022). This paper differs from the earlier work on pension reforms in that I study heterogeneity in employment stability as a source of inequality in lifetime earnings. This also allows me to analyze how changes in the employment-stability distribution affect the optimal pension system.

2.3 Life-cycle model

This section outlines the life-cycle model with heterogeneity in employment stability which combines consumption-saving and labor market behavior. In this model, risk-averse agents maximize expected lifetime utility. The utility function is additively separable in utility from consumption and disutility from providing effort to accumulate human capital. The labor supply of an employed worker amounts to one unit of time. A worker's age is denoted by j.

The life cycle of a worker consists of a working phase and a retirement phase. Let T^W denote the maximal number of periods that agents can remain in the working phase, and let T^R denote the minimum number of periods that agents spend in the retirement phase. The total length of life is fixed to $T = T^W + T^R$. Starting life in the working phase, agents make retirement decisions in the last T^C periods of the working phase. When agents do not decide to retire until the last period of the working phase T^W, agents are shifted to the retirement phase in the next period. Once an agent enters the retirement phase at age $t_R \in [T^W - T^C, T^W]$, the agent remains retired for the remaining life of $T^R + (T^W - t_R)$ periods.

In the working phase, a worker is characterized by her employment status each period. The employment status is summarized by $\epsilon \in \{e, n\}$, where e and n represent employment and unemployment, respectively. The job of an employed worker is described by a combination (w, λ) where w denotes the wage level and λ denotes the separation rate of the underlying job. Employed agents separate from their jobs with probability λ. Wages and separation rates are discretized to $\{w_k\}_{k=1}^{K}$ and $\{\lambda_l\}_{l=1}^{L}$. In each period before retirement, agents receive job offers from a distribution $f(w, \lambda)$. The arrival rates are exogenously given by the parameter π_e when agents are employed and by the parameter π_n for agents in unemployment.

The period budget constraint of an agent is given by

$$a_{j+1} + c_j \;=\; (1 + r)\; a_j + y^n(w_j, h_j, \epsilon)$$

where a and h denote the risk-free asset and the stock of human capital of the agent, respectively. Moreover, r denotes the risk-free rate in the economy, and $y^n(w_j, h_j, \epsilon) = y(w_j, h_j, \epsilon) - \tau(y(w_j, h_j, \epsilon))$ where y^n denotes the labor income of the current period y including transfers minus the pension contribution $\tau(y(w_j, h_j, \epsilon))$ as a function of the labor income. The assets are restricted to be non-negative ($a_j \geq 0$) implying a borrowing limit of zero.

The income of an employed agent in the current period is given by the product of the wage level of current period's job and the human capital stock of the agent, which yields $y(w_j, h_j, e) = w_j h_j$. Although the model does not feature an explicit decision of intensive labor supply, the human capital component acts as a proxy for the effort and skill accumulation that are the result of intensive labor supply and investment in skills. When unemployed, a worker receives a transfer of $y(w_j, h_j, n) = b w_j h_j$, that is, an agent gets a benefit proportional to the labor earnings from the last job. The replacement rate of the unemployment insurance system is denoted by b. The model captures declining benefits in the spell of unemployment by assuming that the wage of the last job drops from w_k to $\max\{w_{k-1}, w_1\}$ if agents stay unemployed and therefore continue receiving unemployment benefits.

In the retirement phase, agents receive retirement benefits $y_r(w_j, h_j, n) = \omega(\bar{y})$ where $\bar{y}$ is the approximation to a worker's average lifetime earnings using the labor earnings in the last period before retirement. The function ω determines the level of benefits assigned to an agent with an average lifetime earnings $\bar{y}$. The next subsection provides a detailed explanation for the approximation of average lifetime earnings and the shape of the function ω. Retirement benefits remain constant during the retirement phase and therefore, agents face no income risk once they enter the retirement phase. There is no bequest motive in the model and agents die at the end of the retirement phase.

Each period of the working phase consists of a separation, an investment, a production, and a search stage. If agents are employed, they lose their jobs with separation probability λ at the separation stage. This separation probability is heterogeneous across jobs and therefore, the probability to lose one's job at the separation stage differs across workers. Agents who do not separate from their job are moved to the investment stage at which they make their human capital investment decisions. In case of a job loss, agents immediately move from the separation stage to the production stage. Employed agents obtain their labor earnings and unemployed agents get unemployment benefits at the production stage. Finally, at the search stage, all agents get job offers with exogenous job-offer arrival rates. These arrival rates differ for employed and unemployed agents. For employed workers, the arrival rate is denoted by π_e and for unemployed agents by π_n. The job offers which consist of a combination of the wage rate w and the separation probability λ are drawn from a joint distribution $f(w, \lambda)$. Upon receiving a job offer, agents can decide whether to accept or to reject the job offer. In case of accepting the offered job, agents are employed in the new job in the next period. A rejection of the job offer does not change the current employment status of the agent: employed agents stay in their current job in the next period and unemployed agents remain unemployed. It is not possible to recall a job offer which was previously rejected.

The investment decision with regard to human capital is a choice of an effort level $t \in [0, 1]$ which entails a quadratic disutility κt^2.[1] For a given level of effort t, the realization of human capital investment is stochastic. More specifically, assum-

1. I also consider an alternative model where human capital accumulation corresponds to the Ben-Porath (1967) model. While all model results of the alternative model remain largely unchanged and the data moments are also matched well, the earnings variance over the life cycle is slightly larger in the alternative model and does not exhibit concavity as observed in the data. See Appendix 2.A.3 for details.

ing that human capital levels are discrete and that h^+ denotes the immediate successor and h^- the immediate predecessor of h, the law of motion for human capital is

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

where $p_H(t,j)$ denotes the age-dependent probability of achieving the next higher level of human capital h^+ for a given effort level t. Without exerting effort, a worker's human capital stock remains constant over time. Because only employed workers have the opportunity for human capital investment, this implies that human capital levels do not change for unemployed workers.

2.3.1 Recursive formulation of the decision problem

In each period, the expected outcome of the separation stage gives the value function for an employed worker V_e as

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

Here, V_n^P denotes the value function of an unemployed agent at the production stage, and V_e^I represents the value function of an employed agent at the investment stage. Because unemployed agents do not face a risk of job loss and cannot invest in human capital, the value function at the separation stage V_n is equal to the value function at the production stage. An employed agent who does not separate from the job makes a human capital investment decision at the investment stage. Since the realization of the human capital investment is stochastic, the value function at the investment stage is given by

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

where V_e^P denotes the value function of an employed agent at the production stage. At the production stage, agents make consumption-saving decisions where the Bellman equation of an employed agent is as follows:

$$\begin{aligned} V_e^P(a,w,\lambda,h,j) &= \max_{\{c,a'\geq 0\}} u(c) + \beta\left(\pi_e V_e^S(a',w,\lambda,h,j) + (1-\pi_e)V_e(a',w,\lambda,h,j+1)\right) \\ \text{s.t.} \qquad c &= (1+r)a + y^n(w,h,e) - a' \end{aligned}$$

In the above equation, V_e^S denotes the value function of an employed agent at the search stage. Moreover, $u(c)$ denotes the period-utility function and β denotes the time preference parameter. Future utility is given by the expected value function as an outcome of job search at the search stage where π_e denotes the job-offer arrival rate. The value function of an employed worker at the search stage depends on the job-offer distribution $f(w, \lambda)$ such that

$$V_e^S(a', w, \lambda, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max\left\{V_e(a', w, \lambda, h, j+1), V_e(a', w_s, \lambda_k, h, j+1)\right\} f(w_s, \lambda_k).$$

In the above expression, N_w and N_λ denote the number of possible realizations in the support of the offer distribution for wages and separation rates, respectively. This value function comprises the decision of acceptance and rejection of expected arrival of outside job offers. Turning to unemployed workers, the value function at the production stage is given by

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

where π_n is a probability of a decline in unemployment benefits to a lower wage level w^-. Note that an unemployed worker receives unemployment benefits $y(w, h, n)$ which is reduced in the next period if the worker remains unemployed. At the search stage, an unemployed worker faces the value function

$$V_n^S(a', w, h, j) = \sum_{s=1}^{N_w} \sum_{k=1}^{N_\lambda} \max\left\{V_n(a', w^-, h, j+1), V_e(a', w_s, \lambda_k, h, j+1)\right\} f(w_s, \lambda_k)$$

which, as for an employed worker, comprises the decision of acceptance and rejection of expected arrival of outside job offers.

Between period $T^W - T^C$ and T^W, workers have the option to leave the labor force and enter the retirement phase. At the beginning of each of these periods, workers observe a shock ε, drawn from a logistic distribution with a location parameter μ and a scale parameter σ. After observing this shock, they decide whether to remain in the labor force and continue working or to move to the retirement phase. This shock captures health shocks and a preference shock related to the value of leisure, representing utility derived from leaving the labor market and entering retirement. Given the realization of the shock, if lifetime utility from retirement is larger than the expected utility of remaining in the labor force, agents decide to retire. Hence, agents face the following discrete choice problem

$$V_{\max}(a, w, \lambda, h, j) \;=\; \max\{V(a, w, \lambda, h, j), V^r(a, w, h, j) + \varepsilon\}, \quad \epsilon \sim \text{Logistic}(\mu, s)$$

where V denotes the value function from remaining in the labor force and V^r denotes the value function of the retirement phase. The expected utility of agents in these periods is given by

$$\begin{aligned} \mathbb{E}\left[V_{max}(a, w, \lambda, h, j)\right] =& pV(a, w, \lambda, h, j) + (1-p)V^r(a, w, h, j_r) \\ & - \sigma\left((1-p)\log(1-p) + p\log(p)\right) + \mu(1-p), \end{aligned} \tag{2.1}$$

where $p = \left(1+\exp\{-\sigma^{-1}(V(a,w,\lambda,h,j)-V^r(a,w,h,j_r)-\mu)\}\right)^{-1}$. Appendix 2.A.1 derives this equation.

Agents who enter the retirement phase receive constant retirement benefits and hence, do not face any earnings uncertainty. All agents die at the end of the retirement phase and there is no bequest motive. The Bellman equation after retirement is

$$V^r(a,w,h,j) = \max_{a'\geq 0} u\left((1+r)a + y_r(w,h,n) - a'\right) + \beta V^r(a',w,h,j+1).$$

The retirement benefits $y_r(w,h,n)$ depend on the labor market history of agents. The next section describes in detail how the retirement benefits are determined.

The current model abstracts from mortality risk, bequest motives, and the early claiming penalties in the Social Security system. This simplification is intended to isolate and clearly highlight the core mechanisms relevant for analyzing the optimal design of the pension system. Appendix 2.A.7 presents several extensions of the model that incorporate these additional features. Importantly, the main results presented in the following sections remain robust under these extended specifications.

2.3.2 Payroll tax finance of the pension system

The pension system is financed by payroll taxes. That is, the expected present value of retirement benefits obtained by all workers in the economy has to be compensated by the expected revenues from payroll taxes levied on employed workers. Hence, the condition

$$\mathbb{E}\left[\sum_{t=t^R}^{T} \frac{\omega(\bar{h})}{(1+r)^{t-1}}\right] = \mathbb{E}\left[\sum_{t=1}^{t^R-1} \frac{\tilde{y}_t(w,h,\epsilon)}{(1+r)^{t-1}} \cdot \tau\right] \tag{2.2}$$

must be satisfied where

$$\tilde{y}_t(w,h,\epsilon) = \begin{cases} y_t(w,h,\epsilon) & \text{if } y_t(w,h,\epsilon) < \text{cap} \\ \text{cap} & \text{if } y_t(w,h,\epsilon) \geq \text{cap} \end{cases}$$

The cap denotes the maximum taxable earnings. The parameter τ in Equation (2.2) denotes the payroll tax rate. Taxes are only levied on labor earnings of employed agents in the working phase. Unemployed agents and agents in the retirement phase do not pay taxes. The tax rate is assumed to be constant over all periods and across all agents.

2.4 Calibration

This section describes the data, parametric assumptions, the functional forms, and the estimated procedure applied to bring the model to the data.

2.4.1 Data

The Survey of Consumer Finances (SCF) is a triennial household survey that offers comprehensive data on income and wealth across a representative sample of U.S. households. In addition to detailed information on financial situation of households, the SCF includes data on the labor market status of household members. To calibrate the model to the data, I pool the SCF waves from 1992 to 2016 and limit the sample to households with employed heads. As the model abstracts from self-employment, I drop households with self-employed household heads, as well as those with extreme wage observations: specifically, observations with wages below 75% of the minimum wage are dropped. As the framework is not designed to explain the very right tail of distributions, I also exclude the top 1% of households by wealth and earnings. Similarly, I exclude the bottom 1% by earnings and households with negative wealth.

2.4.2 Functional forms

In the model, one period is set to match one quarter of a year. Agents derive logarithmic utility from consumption so that $u(c) = \log(c)$. Human capital is discretized on a grid $h_{i,t} \in \{h_1, ..., h_{N_h}, h^*\}$ with $h_1 = 1$ and $h_{N_h} = 6.5$. Human capital levels between these two grid points are equidistant in log space. To capture the right tail of the earnings distribution, the last grid point h^*, which represents the highest level of human capital, is set to $h^* = 25$. The probability to attain the next higher level of human capital when exerting effort t is

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

where the probability p_H starts from a value of $\bar{p}_H$ and decreases in age in a geometric fashion. Upon reaching the human capital level h_{N_h}, the probability to reach the highest level of human capital h^* is an exogenous parameter p_H^* which is independent from age.

Agents enter the labor market at age 20 with the lowest level of human capital $h_1 = 1$ and with an asset level of $a_0 = 0$. The replacement rate for unemployment benefits is set to 0.4 as in Shimer (2005). The total span of the life cycle $T^W + T^R$ is set to 70 years such that workers live up to age 90. Following the U.S. Social Security legislation, workers can start receiving retirement benefits at age 62.[2] Between ages 62 and 67, workers make retirement decisions. At age 67, all workers who decided to remain in the labor force up to this age are shifted to the retirement phase.

Wages and job-separation probabilities are discretized on grids with $[\underline{w}, \overline{w}]$ and $[\underline{\lambda}, \overline{\lambda}]$, respectively. The lowest quarterly separation probability is $\underline{\lambda} = 0.006$ which

2. Information about retirement ages in the United States is available at: https://www.ssa.gov/benefits/retirement/planner/agereduction.html.

corresponds to lifetime jobs with an expected job duration of 42 years. The highest separation probability is set to $\overline{\lambda} = 0.35$. The job-offer distribution consists of a joint distribution of job-separation probability and wage. See Appendix 2.A.4.1 for a detailed description of the job-offer distribution.

The labor earnings in the model are calibrated so as to match the net earnings in the data and therefore, the labor earnings are subject to progressive earnings taxation. The analysis of the optimal pension system in this paper takes the redistribution of earnings through annual earnings taxation as given.

2.4.3 U.S. pension system and retirement elasticity

The pension system of the economy is restricted to the parametric class given by

$$\omega(\bar{y}) = \begin{cases} \phi \cdot (\bar{y})^{1-\gamma} & \text{if } \bar{y} < \text{cap} \\ \overline{\omega} & \text{if } \bar{y} \geq \text{cap} \end{cases} \tag{2.3}$$

where cap denotes maximum taxable labor earnings and $\overline{\omega} = \phi \cdot (\text{cap})^{1-\gamma}$ if average lifetime earnings exceed the cap. This parametric class is often used in the public finance literature to represent tax systems, see, for example, Bénabou (2000, 2002) and Heathcote, Storesletten, and Violante (2017).

The progressivity level of the pension system is determined by the parameter γ and the Social Security cap. The parameter γ determines the shape of the benefit function. For example, $\gamma = 0$ represents a system with flat replacement rate (benefits linearly increase in lifetime earnings) and $\gamma > 0$ corresponds to a strictly concave benefit function in lifetime earnings up to the cap. If the benefit function is strictly concave, then the replacement rate (pension benefit divided by lifetime earnings) decreases in lifetime earnings. Hence, workers with low lifetime earnings face a higher replacement rate compared to workers with high lifetime earnings if $\gamma > 0$. This formulation of the pension system implies that, all else equal, an increase in the parameter γ makes the pension system more progressive.

The cap enters both the contribution and the benefit scheme of the pension system. Therefore, it is not straightforward how the cap affects the progressivity level of the pension system. In the benefit function, the lower the cap, the more progressive becomes the pension system, everything else equal. In contrast, on the contribution side, the lower the cap the less progressive becomes the pension system. As high-income workers with labor earnings above the cap make less contribution to the pension system, it is not clear whether the replacement rate of high-income workers increases or decreases in response to a change in the cap.

In order to calibrate the pension system in the model to the U.S. pension system, I consider the U.S. Social Security legislation for 2019. Based on this, I set the Social Security cap to \$132,900 and the first and second bendpoints to \$926 and \$5,583,

Figure 2.1. Pension System

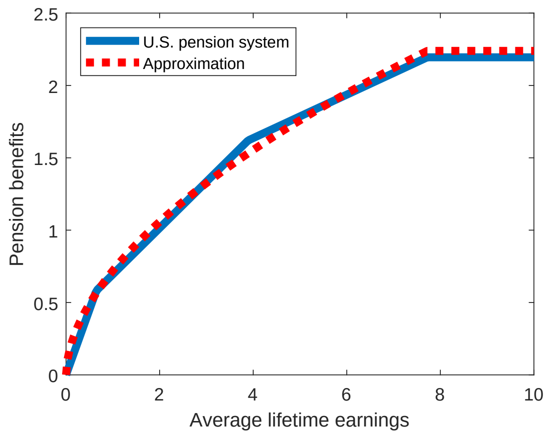

Notes: Comparison of the U.S. Social Security system (solid line) and the approximating function (dotted line).

respectively. Then, I apply the retirement benefits formula given as follows:[3]

$$\Omega(\bar{y}) = \begin{cases} 0.9\bar{y} & \text{if } \bar{y} < bp_1, \\ 0.9bp_1 + 0.32(\bar{y} - bp_1) & \text{if } bp_1 \leq \bar{y} < bp_2, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(\bar{y} - bp_2) & \text{if } bp_2 \leq \bar{y} < cap, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(cap - bp_2) & \text{if } \bar{y} > cap \end{cases} \tag{2.4}$$

where $\bar{y}$ denotes the average lifetime earnings, $\Omega(\bar{y})$ denotes the assigned benefit level, and bp_1 and bp_2 denote the two bendpoints. To match the pension system in the model to the U.S. Social Security, I target the net pension replacement rate for an average worker in terms of lifetime earnings in the model to the observed net pension replacement rate of 49.4% for an average worker in the United States reported by OECD (2019). The estimated parameters are $\phi = 0.72$ and $\gamma = 0.44$. Figure 2.1 shows the U.S. Social Security benefit function and the approximating function in the model.

I use the final level of human capital and wages of individual workers attained at the end of the working phase to infer the lifetime earnings of workers. Human capital is a direct measure of lifetime earnings as workers accumulate human capital throughout their entire career and labor earnings are a function of human capital. The final level of human capital also contains information on workers' earnings histories because human capital accumulation is only possible in employment. Thus, it

3. The retirement benefit formula follows the documentation http://www.ssa.gov/OACT/COLA/piaformula.html.

Figure 2.2. Labor market transitions and tenure

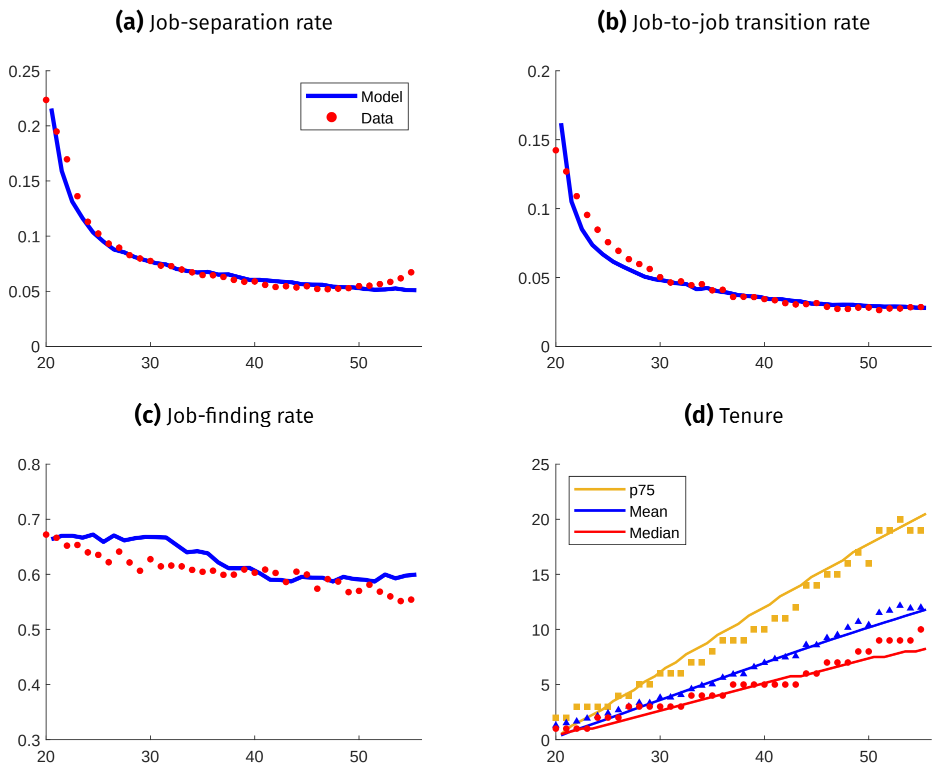

Notes: This figure shows quarterly life-cycle transition rates and tenure in years by age. Panel 2.2a shows the average separation rate, Panel 2.2b the job-to-job transition rate, and Panel 2.2c the job-finding rate. Panel 2.2d shows the tenure distribution for the mean, median, and the 75th percentile of the distribution. Tenure in the model is computed as the time spent on the same job without separation or job-to-job transitions. The dots display the empirical profiles and the solid lines show the model profiles. The empirical profiles are derived from the Current Population Survey.

is informative about workers' average labor earnings and their duration of employment during the working phase. I regress the average lifetime earnings of workers on cubic polynomials of the final level of labor earnings of each worker which consist of human capital and wages. In fact, the final level of human capital and wages attained by workers until retirement explains more than 90% of variation in lifetime earnings. Appendix 2.A.4.3 discusses the accuracy of this approximation of average lifetime earnings in detail.

The parameters of the logistic distribution, which govern the retirement decision of the agents, are chosen so as to match the retirement age distribution in the U.S. data and the retirement elasticity typically found in the macroeconomic literature. In the model, recall that an agent chooses to retire if

$$V^r(a,w,h,j_r) + \varepsilon \;\geq\; V(a,w,\lambda,h,j), \quad \varepsilon \sim \text{Logistic}(\mu,\sigma)$$

where μ denotes the location parameter and σ denotes the scale parameter of the logistic distribution. I set on average $\mu = 13.68$ and $\sigma = 2.9$. The shock ε drawn

Figure 2.3. Earnings and wealth

(a) Earnings **(b)** Earnings variance **(c)** Wealth-to-income ratio

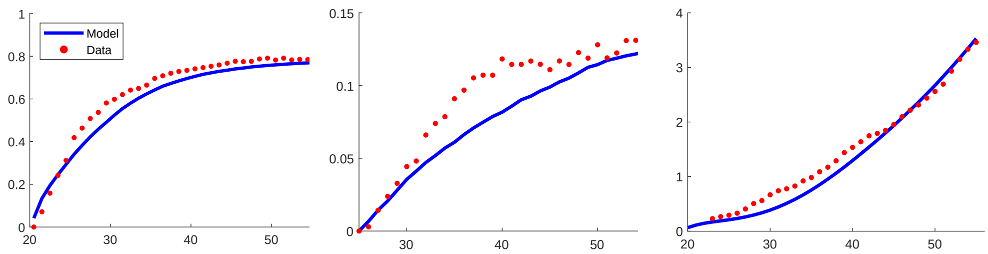

Notes: This figure shows quarterly life-cycle profiles of earnings and wealth-to-income ratio by age. Panel 2.3a shows the average labor earnings, Panel 2.3b the variance of labor earnings, and Panel 2.3c the wealth-to-income ratios. The dots display the empirical profiles and the solid lines show the model profiles. The empirical profiles are derived from the Current Population Survey and the Survey of Consumer Finances.

from the logistic distribution can be interpreted as allowing for heterogeneity in preferences for leisure. It can be also interpreted as deriving an unexpected positive utility from retirement such as health problems, which are an important cause for retirement (see, for example, (Dwyer and Mitchell, 1999; Disney, Emmerson, and Wakefield, 2006; Jones, Rice, and Roberts, 2010; Galama et al., 2013; Trevisan and Zantomio, 2016).

The remaining parameters are estimated by applying a simulated method of moments. In particular, I match the model profiles to the following empirical moments: the life-cycle profiles of labor market transition rates, mean and variance of earnings, tenure distribution of mean, median, and the 75th percentile, and the wealth-to-income ratio. Appendix 2.A.4 provides further details on the estimation procedure and Table 2.A.1 summarizes the estimated parameters.

Figure 2.2 compares the life-cycle profiles of job-separation, job-to-job, job-finding rates, and the tenure distribution in the model and in the data. Both job-separation and job-to-job rates strongly decline between age 20 and 30 and decline until age 50. The model profiles fit the empirical profiles well with a steep decrease of job-separation rate and job-to-job transition rate at the beginning of the life cycle, followed by a flattening of the profiles after age 30. The empirical profile of job-finding rates also exhibit a decreasing trend over the life cycle which is closely matched by the model. In Figure 2.2d, I compare the empirical profiles of mean, median, and the 75th percentile of the tenure distribution with their model counterpart. The increasing gap over the life cycle between the 75th percentile, the mean, and the median of the tenure distribution indicates that the tenure distribution gets more dispersed over time. Heterogeneity in employment stability generates differences in tenure across workers, and these differences grow over the life cycle both in the data and in the model. Figure 2.3 displays the profiles of mean log earnings, earnings variance, and wealth-to-income ratio. The model is able to match the strong increase in earnings at the beginning of the life cycle and the decreasing income growth over time when workers get closer to the end of prime age. The model also matches the steep increase of the earnings variance as well as the life-cycle profile of wealth-to-income ratio.

2.4.4 Retirement age distribution and retirement elasticity

Figure 2.4 plots the distribution of retirement age in the data for the United States in 2019 and in the baseline calibration of the model. The retirement pattern in the United States exhibits two huge peaks at age 62 and at age 66; around 32.6% of workers entitled for retirement benefits in 2019 were at age 62, which is the early eligibility age of the current U.S. Social Security rule; around 25.3% were at age 67. The proportion of retired workers between ages 63 and 65 were around 27.5%. Around 14.5% were at age 67 or older. Figure 2.4 shows that the specification of the model matches the U.S. retirement pattern well. In particular, the model is able to fit the huge peaks at ages 62 and 67.

I define retirement elasticity as the percentage change in the retirement hazard relative to a change in the generosity of retirement benefits and apply the estimated retirement elasticity by Coile and Gruber (2007). Using the Health and Retirement Survey, Coile and Gruber (2007) analyze the effect of pension incentives on retirement behavior. They find an elasticity of retirement with respect to retirement benefits of 0.16 which I target for the retirement elasticity in the model. The policy change takes place for workers at the age of 55 holding the degree of progressivity of the pension system constant at the baseline level.

These changes in retirement decisions fall into the range of retirement elasticities typically found in other empirical studies. For example, Moulton and Stevens (2015) follow the methodology of Coile and Gruber (2007) and obtain similar responses in the retirement probability for an increase in Social Security wealth. Brown (2013) uses a quasi-experimental design to provide estimates of the price elasticity of lifetime labor supply. The results in that paper indicate that a 10% change in retirement benefits lead to an adjustment of retirement age by less than two months. Burtless (1986) and Krueger and Pischke (1992) also investigate by how much a change in the Social Security benefit level affects retirement behavior and find a minor role for policy changes.

In terms of order of size, the calibrated retirement elasticity is also similar to studies on policy experiments applying structural models. For example, Kimball and Shapiro (2003) use a model of consumption and labor supply to study retirement behavior and find that reducing benefits by 10% induces a postponement of retirement by between 0.1 and 0.5 for workers around age 50. This corresponds to an elasticity of total years worked relative to a change in retirement benefits of −0.025 and −0.125.

Figure 2.4. Retirement age distribution in the data and in the baseline model

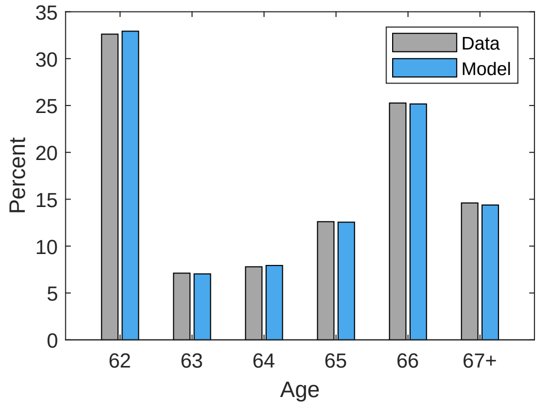

Notes: This figure shows the distribution of retirement age in the data for the United States in 2019 and in the baseline calibration of the model. The data source is Table 6.A4 of Social Security Administration (2020), Annual Statistical Supplement, available at https://www.ssa.gov/policy/docs/statcomps/supplement/.

2.5 Heterogeneity in employment stability and life-cycle dynamics

Using the calibrated model, this section studies how heterogeneity in employment stability shapes inequality in labor market outcomes and shows the life-cycle implications of heterogeneity in employment stability for the baseline calibration of the model.

I simulate life cycles for a large population of workers and compute the mean job-destruction probability of each worker in employed periods after the labor market entry at age 20 up to the last period before retirement. Taking the inverse of the mean job-destruction probability for each worker yields a distribution at the end of the working phase which I define as the distribution of average employment stability over the life cycle. By taking agents below the first quartile and above the third quartile of this distribution, I compare these two groups of workers in terms of their average life-cycle profiles of human capital, labor earnings, consumption, and wealth. Workers in the top quartile are those who had on average the most stable jobs over the life cycle and the bottom quartile refers to workers who had on average the most unstable jobs. This approach allows to analyze the relationship between the extent of job instability a worker has to cope with in the working phase and the life-cycle profiles of key economic variables.

2.5.1 Inequality in labor market outcomes

Figure 2.5a illustrates how the two groups of agents differ in terms of their average job-separation rate over the life cycle. Comparing the two profiles in Figure 2.5a, on average, the top quartile already finds more stable jobs at labor market entry. More-

Figure 2.5. Inequality in labor market history

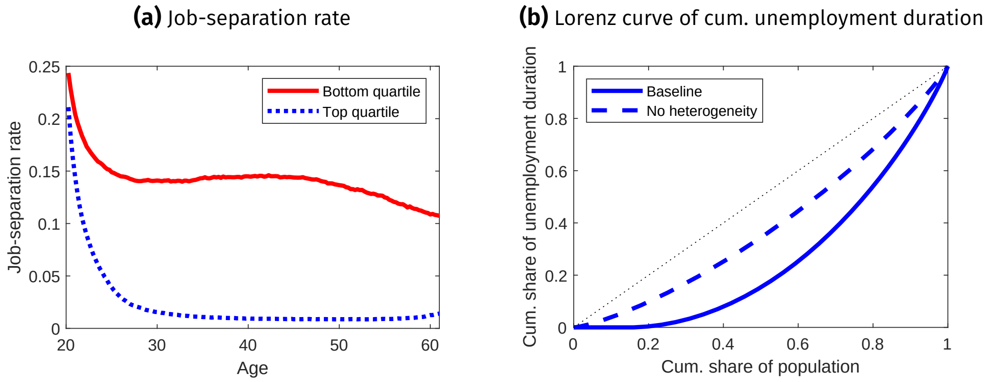

Notes: Panel 2.5a shows the average life-cycle profile of job-separation rate. The solid line displays the average life-cycle profile of job-separation rate of workers with most unstable jobs (bottom 25% of employment-stability distribution). The dashed line represents workers with most stable jobs (top 25% of employment-stability distribution) throughout working life. Panel 2.5b shows the Lorenz curve of cumulative unemployment duration where the solid line displays the baseline economy and the dashed line represents the Lorenz curve for a counterfactual economy in which there is no heterogeneity in employment stability conditional on age.

over, the average separation rate of the top quartile remains close to the most stable job which represents lifetime jobs with a separation probability of 0.006 per quarter. For agents in the bottom quartile of the distribution, the average separation rate drops from 0.25 to below 0.15 in the first five years. But then, during the working phase, there is no significant improvement in employment stability and the profile remains flat until age 50. Towards the retirement age, the average separation rate declines to 0.11, but still, the gap in employment stability between agents in the top and bottom quartile of average employment stability over the life cycle remains sizeable. Compared to the top quartile, workers in the bottom quartile fail on average to find stable jobs over the life cycle. Overall, the difference in employment stability is significant for the two groups of workers and remains persistent until the end of the working phase.

This heterogeneity in employment stability translates into a large inequality in employment history. Figure 2.5b shows the Lorenz curve for cumulative unemployment duration of prime-age workers between ages 25 and 55. The solid line displays the Lorenz curve for the baseline economy and the dashed line a counterfactual economy without heterogeneity in employment stability. More specifically, the counterfactual economy is constructed by assuming that all workers have the same age-dependent separation rate from their jobs. The age-dependent separation rate corresponds to the average separation rate conditional on age in the baseline economy with heterogeneity in employment stability. That is, in the counterfactual economy, the separation rates of workers vary over the life cycle, but there is no cross-sectional heterogeneity in employment stability conditional on age. Comparing the baseline economy and the counterfactual economy, the Lorenz curves indicate that heterogeneity in employment stability is a key driver of inequality in employment history. In the baseline economy, about 40% of all workers account just for 8% of total cumulative unemployment duration of all workers, whereas the Lorenz curve of the counterfactual does not exhibit a conspicuous curvature.

While there are no fixed worker types in this model, there is a high persistence of the initial employment stability over the life cycle. More specifically, I find that a worker with an unstable job at age 30 (below the median of the employment-stability distribution at age 30) will remain in the unstable group at age 54 (below the median of the employment-stability distribution at age 54) with a probability of 69%.[4] Conversely, a worker with a stable job at age 30 will remain in the stable group at age 54 with a probability of 70%. If the model does not produce any persistence of the initial employment stability for workers, the transition probabilities between the types would be 50%. In Appendix 2.A.5, I show that these transition probabilities are consistent with empirical transition probabilities using the Panel Study of Income Dynamics. While the model of this paper does not assume fixed worker types regarding employment stability, the model generates stable and unstable career paths that remain persistent over workers' life cycle. The model implications are therefore consistent with the empirical literature that considers heterogeneity in employment stability as worker types rather than job types (Hall and Kudlyak, 2019; Gregory, Menzio, and Wiczer, 2021; Ahn, Hobijn, and Şahin, 2023).

2.5.2 Life-cycle consequences of employment stability

Figure 2.6a shows the profiles of average human capital level for workers in the top and bottom quartile of the average life-cycle employment-stability distribution. The profiles are normalized by the initial levels and expressed in log deviations. A higher average job-separation rate directly affects the human capital accumulation process since workers lose the opportunity to invest in human capital upon job loss and the expected return on human capital decreases in job instability. As agents with low employment stability have fewer opportunities to invest in human capital, the profile of the bottom quartile exhibits a lower growth than the profile of the top quartile. Consequently, the gap in the average human capital between the two groups of agents rises over the life cycle. In particular, starting from the same initial level of human capital, the top quartile has on average 22% higher stock of human capital at the end of the working phase.

A similar pattern is observed in Figure 2.6b for average earnings and in Figure 2.6c for average consumption. Initially, both groups start from the same level of aver-

4. The unstable group at age 30 refers to workers below the median of the employment-stability distribution at age 30. Equivalently, the unstable group at age 54 refers to workers below the employment-stability distribution at age 54.

Figure 2.6. Average life-cycle profiles by employment stability

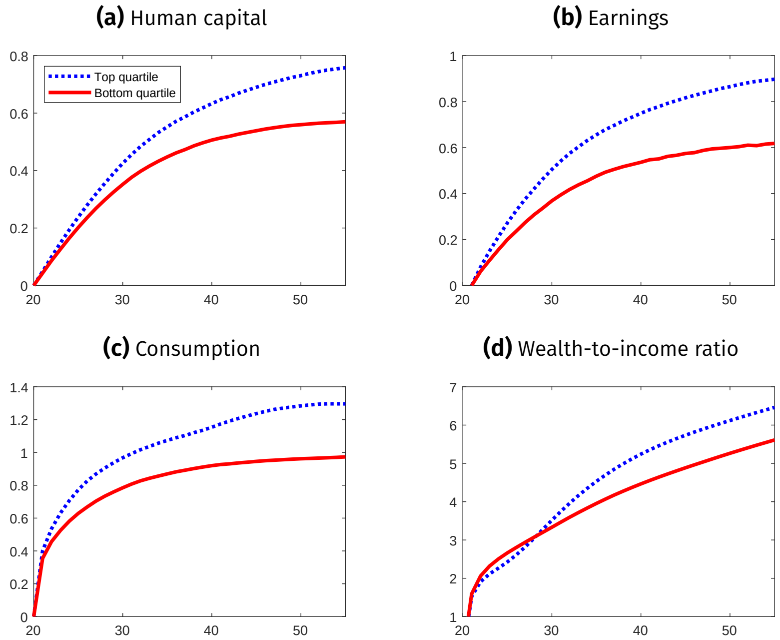

Notes: This figure shows the average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), and consumption (lower left panel), and wealth (lower right panel). The profiles are the log deviations from the initial values of the respective profiles. In all plots the solid line represents workers with most unstable jobs (bottom 25% of employment-stability distribution) and the dashed line represents workers with most stable jobs throughout working life (top 25% of employment-stability distribution).

age earnings, but the profiles quickly diverge during the first five years in which the gap in average separation rate also widens. The earnings and consumption profiles of the bottom quartile features a smaller growth over the life cycle.

Figure 2.6d shows the average life-cycle profiles of wealth. For workers in the top quartile, wealth starts to grow strongly 10 years after labor market entry. The profile of the bottom quartile also increases throughout life, but the growth in wealth is dampened by low earnings growth. Towards the end of the working phase, workers in the bottom quartile have significantly lower assets than agents in the top quartile. Since a high job-separation probability leads to fewer opportunities to invest in human capital and to more frequent job losses, workers in the bottom quartile have on average lower earnings which lead to lower savings. Moreover, workers dissave upon job loss in order to smooth consumption, which in turn decreases asset accumulation. The qualitative results remain unchanged when workers are assigned to the stable (top quartile) and unstable group (bottom quartile) using the employment-stability distribution at age 30. In Appendix 2.A.5, I show that due to persistence of employment stability, using the labor market history in the first 10 years after labor market entry yields a similar result as in Figure 2.6.

In sum, the risk of becoming unemployed constitutes a significant source of earnings uncertainty for individuals in the labor force and heterogeneity in employment stability is a key driver of inequality in lifetime earnings. Consistent with the empirical literature on employment stability, the model at hand produces persistence of employment stability where workers with unstable jobs when young tend to have unstable jobs over the whole life cycle. In stable jobs, workers have the opportunity to invest in human capital and to climb the job ladder, which leads to high earnings growth over the life cycle. Employment stability significantly alters the precautionary saving motive of young workers. Stable employment mitigates the necessity of accumulating precautionary savings as earnings uncertainty is low and therefore, these workers have a better capability to engage in life-cycle consumption smoothing. In addition to the incomplete financial markets which restrict workers from borrowing, having unstable jobs in the early career further constrains workers' ability to engage in life-cycle smoothing of consumption and leads to poor life-cycle outcomes.

The analysis of this section shows that the distribution of employment stability affects the distribution of lifetime earnings and consumption. Therefore, the degree of heterogeneity in employment stability has crucial implications on the desired level of redistribution in the economy and shapes the optimal design of pension systems. A progressive pension system implies decreasing replacement rates in average lifetime earnings and redistributes earnings from high-income workers to low-income workers, reducing consumption inequality across workers after retirement. Progressive pension benefits alleviate earnings shocks accumulated over the working life which would be otherwise fully carried over to the retirement.

2.6 Optimal progressive pension system

Starting from the baseline economy described in the previous section, I derive the optimal pension system. Holding the total expected pension benefits constant at the baseline level, I search for the optimal progressivity parameter γ and the optimal level of Social Security cap that lead to the highest welfare in the economy. I set up a grid for the parameter γ and the cap and compute the corresponding parameter ϕ and the payroll tax rate that achieve budget balance for the government while the total amount of benefits is equal to the total amount of benefits in the baseline economy. This paper focuses on how to most effectively allocate a fixed pool of pension funds from the baseline economy. While the long-term sustainability of pension financing is undoubtedly important, the central concern here is the optimal design of benefits in the presence of employment risk. Given the significant role that employment stability plays in shaping lifetime earnings and savings, the key question is how pension benefits can be structured to provide insurance against persistent labor market risk. The next subsection explains the welfare measure that I apply for the welfare analysis.

2.6.1 Welfare measure

In order to evaluate the welfare effects of alternative pension policies, I compute the consumption-equivalent variation (CEV) which makes workers indifferent between the baseline economy and the economy with an alternative pension system in terms of expected lifetime utility. More precisely, this welfare measure indicates how much additional consumption agents require in the baseline model in order to get a change of expected lifetime utility equal to the change generated by an alternative pension system. Formally, I derive

$$\text{CEV} = \exp\left(\frac{V_a - V_b}{\widetilde{\beta}}\right) - 1$$

where $\widetilde{\beta} = \frac{1-\beta^{T^W+T^R+1-j}}{1-\beta}$ and V_b and V_a denote the expected lifetime utility in the baseline economy and in the economy under an alternative pension system, respectively. For $j = 1$, this welfare measure compares the ex-ante expected lifetime utility at labor market entry in the baseline economy and an economy under an alternative policy. Hence, this welfare measure incorporates the expectation about all possible states and all relevant information over the life cycle.

2.6.2 Measure of progressivity level

The progressivity level of a pension system depends on multiple components: the benefit formula, the contribution via payroll taxes, and the distribution of lifetime earnings in the economy. The shape parameter γ in the retirement benefit formula indicates the extent of variation of the replacement rate across individuals as a function of the average lifetime earnings. However, this parameter alone is not fully informative about the degree of progressivity of the pension system. In fact, the progressivity level depends on the net replacement rate, which is the ratio between the level of retirement benefits and the net average lifetime earnings after deducting payroll taxes, and the Social Security cap, which enters both the contribution and benefit side of the pension system. To capture all components that affect the progressivity level of the pension system, I develop the following measure of progressivity.

Let η_{25} denote the average net replacement rate of workers in the first quartile and η_{75} the fourth quartile of the average lifetime earnings distribution. Then the measure for the degree of redistribution through the pension system ζ is given by

$$\zeta = \frac{\eta_{25} - \eta_{75}}{\eta_{25} + \eta_{75}}.$$

Figure 2.7. Welfare and changes in progressivity level

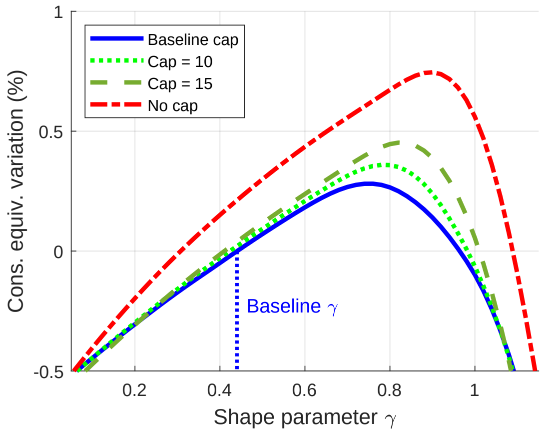

Notes: This figure shows the change in welfare in consumption-equivalent variation induced by changes in the shape parameter γ and the Social Security cap. The solid line displays the welfare change as a function of the shape parameter γ when the Social Security cap is at the baseline level (7.72). The dotted line and dashed line show the welfare changes when the Social Security cap is increased to 10 (30% higher cap than the U.S. Social Security cap) and 15 (94% higher cap than the U.S. Social Security cap), respectively. The dashed-dotted line displays the welfare change when the Social Security cap is removed from the pension system.

Note that, if all individuals are assigned the same net replacement rate, then $\zeta = 0$. If the net replacement rate of the first quartile of the average lifetime earnings distribution is strictly higher than the net replacement rate of the fourth quartile and hence, the pension system is progressive, then $\zeta > 0$.

2.6.3 Optimal pension system

Figure 2.7 displays the welfare change induced by varying the shape parameter γ and the Social Security cap of the pension system in consumption-equivalent variations. Recall that for a given level of the Social Security cap, a pension system with $\gamma = 0$ implies that retirement benefits linearly increase in lifetime earnings up to the cap. An increase in γ to $\gamma > 0$ makes the pension system more progressive. The figure shows that welfare increases with the Social Security cap and a repeal of the cap is optimal. The concave shape of the welfare as a function of the shape parameter γ reveals the trade-off between redistribution and incentive distortions for the optimal design of the pension system. Progressivity offers insurance against unstable employment histories and low lifetime earnings, but it comes at the cost of distorting human capital investment and retirement decisions.

Table 2.1 displays the policy parameters associated with each progressivity level γ. To give a manageable overview of the results, Table 2.1 displays the results only for selected values of parameters. The parameters in the baseline economy are displayed in the first line where $\gamma = 0.44$ and $\phi = 0.72$ with a payroll tax rate of 9.17%. Comparing the welfare change across all parameter specifications, the optimal policy is a pension system with no Social Security cap and the shape parameter set to $\gamma = 0.9$. The corresponding progressivity level is 42.23% and this parameter combination leads to a welfare gain of 0.75% in terms of consumption-equivalent variation for a labor market entrant. The payroll tax rate that achieves budget balance for the pension system decreases to 8.9% and the corresponding level of parameter ϕ increases to 1.12. To understand where this welfare gain comes from, in the next subsections I investigate how retirement decisions and life-cycle dynamics change under the optimal pension system compared to the baseline economy.

The results remain robust when the model is extended to incorporate mortality risk, a bequest motive, and early pension penalties in the Social Security system (see Appendix 2.A.7.1). In the baseline analysis, the income tax system and the unemployment insurance policy are held constant. To evaluate the interaction between the pension system and other policy instruments, I extend the analysis in Appendix 2.A.7 by varying the degree of income tax progressivity and the generosity of unemployment insurance benefits. The results show that the welfare gains from implementing the optimal pension system decline as income tax progressivity and the generosity of the unemployment insurance systems increase. Nevertheless, the qualitative conclusion remains robust: increasing pension progressivity is always welfare-improving relative to the current system and the welfare gains from a more progressive pension system remain substantial.

Note that the model assumes that all workers are ex-ante identical in their wealth level, human capital, and the initial labor market status. The optimal pension system therefore induces the same welfare change for all workers in the model. Ex-ante heterogeneity in the state variables, however, may have important welfare implications on the optimal pension system. In Appendix 2.A.6, I discuss the welfare changes of the optimal pension system as functions of job-separation rate, wealth level, and human capital. Finally, I consider a model with human capital depreciation in Appendix 2.A.8. Overall, all qualitative results regarding the optimal pension system remains unchanged under the alternative model specifications.

How does heterogeneity in employment stability affect the optimal level of progressivity of the pension system? In Section 2.5, I have shown that heterogeneity in employment stability generates differences in employment histories and accounts for a large fraction of inequality in lifetime earnings. In order to analyze how heterogeneity in employment stability affects the optimal pension system, I repeat the welfare analysis for the counterfactual economy without heterogeneity in employment stability. In the counterfactual economy, there is no cross-sectional heterogeneity in employment stability conditional on age, but the separation rates vary over the life cycle. Figure 2.8 displays the welfare changes induced by variations in the retirement benefit formula in the baseline economy and the economy without heterogeneity in employment stability. In Figure 2.8a, the Social Security cap is kept constant at the baseline level. In the economy without heterogeneity in employment stability and a fixed Social Security cap, the current U.S. pension system is close

Table 2.1. Welfare change in the baseline economy under alternative pension systems

Model	Parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline	0.44	0.72	7.72	25.32	9.17	-
No cap	0	0.43	–	5.20	8.74	−0.63
	0.2	0.54	–	14.56	8.77	−0.20
	0.4	0.68	–	23.17	8.81	0.15
	0.6	0.84	–	31.18	8.83	0.44
	0.8	1.02	–	38.67	8.87	0.67
	0.9	1.12	–	42.23	8.88	0.75
	1	1.22	–	45.69	8.89	0.55
	1.2	1.37	–	53.27	8.91	−1.14
	1.4	1.49	–	58.99	8.93	−3.14

Notes: The results for the model "No cap" are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 presents the degree of progressivity achieved by the combination of policy parameters and Column 6 the payroll tax rate in the economy that achieves budget balance for the pension system. Column 7 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

to optimal. Any deviation of the shape parameter γ from its baseline specification leads to welfare losses. Therefore, in the absence of heterogeneity in employment stability, the optimal pension system implies a much lower level of progressivity.

The abolishment of the Social Security cap is, however, welfare increasing in both the counterfactual and the baseline economy. Figure 2.8b shows the welfare changes in the shape parameter in combination of the abolishment of the Social Security cap. In the absence of the cap, an increase in the shape parameter leads to welfare gains in both economies. Yet, the optimal level of progressivity remains lower in the counterfactual without heterogeneity in employment stability. In the baseline economy, the optimal progressivity level is 42.23%, whereas in the counterfactual economy, the optimal pension system implies a progressivity level of 32.45%. Hence, a disregard of heterogeneity in employment stability decreases the optimal level of progressivity by a quarter.

Figure 2.8. Comparison of welfare changes in the baseline economy and the economy without heterogeneity in employment stability

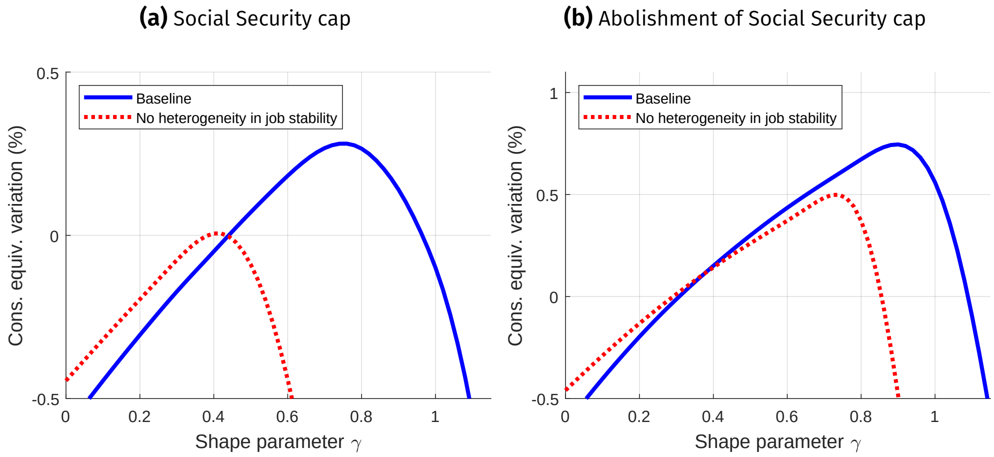

Notes: The graphs show the welfare changes in the baseline economy (solid lines) and the economy without heterogeneity in employment stability (dotted lines) induced by variations in the shape parameter of the retirement benefit function. In the left panel, the Social Security cap is fixed to its initial level. In the right panel, the policy changes include the abolishment of the Social Security cap.

2.6.4 Welfare decomposition

The welfare change under the optimal policy arises from various components. In the following, I decompose the welfare gain into a tax component, an insurance component, a retirement component, and a distortionary component. The tax component is derived as follows: I construct a counterfactual economy in which the optimal pension parameters are implemented, but workers face the same tax rate as in the baseline economy. The welfare change from comparing the baseline economy to the counterfactual economy (CEV_{tax}) is then subtracted from the actual welfare gain (CEV), which yields the tax component ($CEV - CEV_{tax}$). In the next step, I endow workers with the average level of human capital at the end of the working phase. At this average level, the replacement rates are identical in the baseline economy and in the economy under the optimal pension system. However, workers starting with the average human capital do not require insurance against bad labor market outcomes. By subtracting this welfare change (CEV_{ins}) from the tax component, I obtain the insurance component. Lastly, I construct an economy in which workers are, in addition to the previous specifications, not allowed to choose their retirement age. The welfare change under this assumption (CEV_{ret}) is then subtracted from the insurance component which yields the retirement component. The welfare loss arising from distortionary effects is then given by the difference between the total welfare effect and the sum of all components.

The decrease in the payroll tax rate under the optimal pension system induces a welfare gain of 0.27%. The abolishment of the Social Security cap increases the

Table 2.2. Welfare decomposition

Effects	Welfare change (%)
Tax rate	0.27
Insurance	1.78
Endogenous retirement	0.07
Distortions	−1.37

Notes: The first column specifies the effects that affect the welfare gain in the economy. The second column presents the welfare change of each effect. The welfare changes are derived by gradually shutting down each effect.

contribution of high-income workers so that the tax rate declines. A lower tax rate is particularly beneficial for young workers as payroll taxes depress consumption when earnings are low and workers have a high motive for precautionary savings. The choice of the retirement age allows workers to adjust their labor supply at the extensive margin in response to the policy change and induces a welfare gain of 0.07%. The adjustment of labor supply allows to react to the changes in pension wealth. The insurance effect of the optimal pension system leads to a welfare gain of 1.78% as workers have a higher insurance against bad labor market outcomes. The distortinary effects arise as workers invest less into human capital and reduce on average their labor supply under the optimal pension system.

2.6.5 Retirement incentive distortion

The optimal pension system changes the retirement decision of workers and thus entails a distortionary effect. This section aims to examine this distortionary effect more closely. Table 2.3 compares the average retirement age in the baseline economy and the economy under the optimal policy. The first row shows the mean retirement age of all workers for each economy and indicates that the mean retirement age decreases from 64.33 to 64.16. This result implies that on average, workers choose to retire earlier in response to the increase in pension progressivity.

Yet, the change in the retirement age differs for workers depending on their stock of human capital. The second and third rows of Table 2.3 show the mean retirement age of two groups of workers: one group comprises workers who end up with a stock of human capital that is below the median of all workers; the other group refers to workers who achieve a human capital stock above the median of all workers. Workers below the median choose to retire earlier in the economy under the optimal policy than in the baseline economy, whereas workers above the median choose to stay longer in the labor force.

When workers become eligible for retirement benefits, they have the option to retire or continue working. This decision is shaped by two opposing effects. On the

Table 2.3. The mean retirement age in the baseline economy under the ex-ante optimal pension system

Mean retirement age	Baseline	Optimal pension system
All	64.33	64.16
Human capital below median	64.04	63.69
Human capital above median	64.99	65.29

Notes: Each column presents the mean retirement age in the baseline economy and the economy under the optimal pension system for all individuals and individuals above and below median human capital.

one hand, remaining in the labor force leads to higher earnings and offers the possibility to accumulate additional units of human capital for workers, which may raise the level of retirement benefits. On the other hand, the number of periods of benefit receipt decreases which reduces total pension wealth. These effects are different in terms of their magnitude for workers with different employment status, human capital, and wages. Therefore, the net effect of remaining an additional period in the labor force on lifetime income is ambiguous.

The increase in pension progressivity induces different wealth effects on workers with different levels of human capital. The policy change reduces the pension wealth of workers with large human capital stocks and the relative difference between labor earnings and retirement benefits increases under the optimal pension policy. Therefore, these workers delay their retirement to continue working and receiving labor earnings. By contrast, the increase in pension progressivity raises the pension wealth relative to labor earnings for workers who do not achieve a high human capital stock before retirement. Since their expected value of labor earnings from staying in the labor force is low, it becomes optimal for these workers to retire earlier.

Overall, these results imply that the total distortionary effect of the optimal pension system on retirement decision remains small. As workers with a high stock of human capital decide to retire later, this effect partly offsets the retirement incentive distortion on workers with low human capital.

2.6.6 Life-cycle dynamics under the optimal progressive pension system

What are the life-cycle consequences of implementing the optimal pension system? As in the previous section, I simulate the economy under the optimal pension system for a large number of agents and derive the distribution of average employment stability over the life cycle. Taking agents up to the first quartile and above the third quartile of the distribution, I investigate the life-cycle consequences of the optimal progressive pension system for the two groups of agents.

Figure 2.9. Relative change in average life-cycle profiles by employment stability

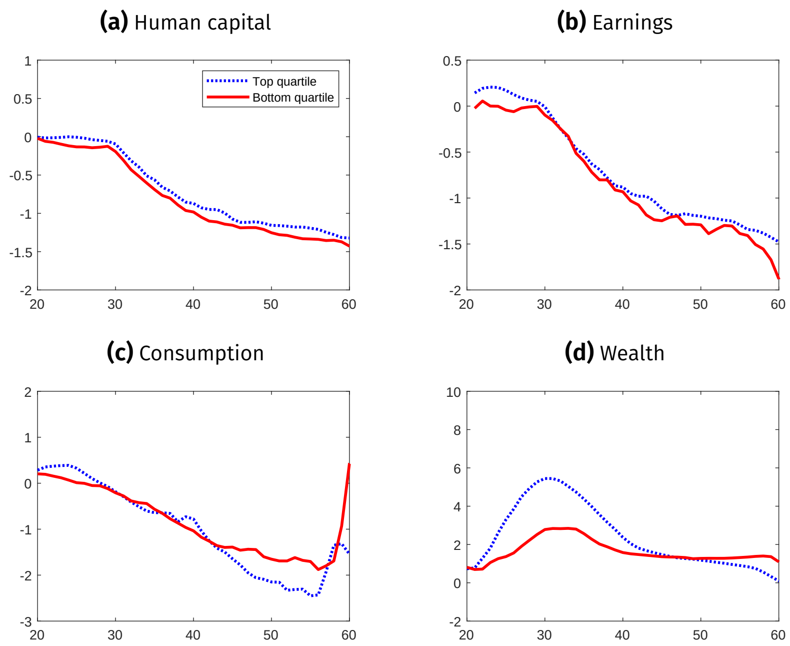

Notes: This figure shows the percent deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy under the optimal pension system compared to the baseline economy. In all plots the solid line represents workers with most unstable jobs (bottom 25% of employment-stability distribution) and the dashed line represents workers with most stable jobs throughout working life (top 25% of employment-stability distribution).

Figure 2.9 shows the relative deviation from the baseline economy of human capital, labor earnings, consumption, and wealth in the economy under the optimal pension system. The first finding is that on average, the increase in pension progressivity discourages human capital investment over the life cycle. In Figure 2.9a, starting at age 30, the relative change in the profiles of average human capital declines steadily until the early retirement age, implying that workers accumulate less human capital under the optimal policy than under the baseline specification.[5] The gap in human capital level grows over life and amounts to −1.4% for agents in the top and −1.5% for agents in the bottom quartile at the early retirement age. Because achieving a high level of human capital is associated with a lower replacement rate in the retirement benefit formula than in the baseline economy, the return on human capital decreases for workers in the top quartile. Consequently, agents employed in stable jobs for whom human capital investment is large otherwise strongly

5. The early retirement age refers to the earliest age at which workers become eligible for retirement benefits. In the model, this corresponds to age 62.

decrease their effort provision for human capital accumulation. A similar reason accounts for lower human capital investment of workers in the bottom quartile.

Interestingly, during the first 10 years after labor market entry, workers do not reduce their investment in human capital. This is because human capital investment is highly productive in the early working phase. Young workers strive to accumulate human capital to increase prospective labor earnings growth and lifetime earnings such that human capital accumulation mostly happens at the beginning of the life cycle. The decrease in return on human capital investment, which realizes after retirement, has little impact on the incentives for human capital accumulation and labor market outcomes of young workers. As a consequence, the policy change does not impede human capital accumulation in the early working phase. When the retirement age gets closer and the disincentives for human capital investment start to grow and become more relevant in the later stage of life, workers reduce their effort provision in any case because investment in human capital becomes unproductive. Hence, the gap in human capital between the baseline economy and the economy with higher pension progressivity does not grow further and remains almost unchanged close to retirement.

The intuition why the progressivity level of the pension system has little effect on human capital investment of young agents is similar to the idea of Michelacci and Ruffo (2015). They show that an optimal unemployment insurance system should provide higher benefits to younger agents since, in contrast to older workers, the moral hazard problem is small for the young. Michelacci and Ruffo (2015) explain that, in line with the findings in this paper, having a job is important for young workers since employment not only increases current income, but it also provides the opportunity to accumulate human capital and therefore a higher income growth. This small distortion in human capital investment among young workers is consistent with empirical evidence suggesting that many individuals are not fully aware of how Social Security benefits are structured (Bairoliya and McKiernan, 2021). Although the model assumes rational expectations, young workers' behavior remains largely unaffected because their short-term labor market prospects matter more than distant changes in retirement benefits. Importantly, the optimality of the proposed pension system is not based solely on young workers' welfare. As shown in Figure 2.12, the transition also yields welfare gains for older workers nearing retirement.

The change in the shape of the average human capital profile consequently affects the average life-cycle profiles of earnings and consumption displayed in Figure 2.9b and Figure 2.9c. The first observation is that the relative changes in the profiles of average earnings closely follow the shape of the profiles in Figure 2.9a. For both groups of workers, the decrease in the average human capital implies lower average earnings compared to the baseline economy. The average consumption of workers in the top quartile decreases towards the later stage of the working phase. Also for the bottom quartile, the average consumption decreases because the average earnings become lower. However, the average consumption strongly increases

Figure 2.10. The life-cycle profiles for variance of log consumption and Gini coefficient of consumption

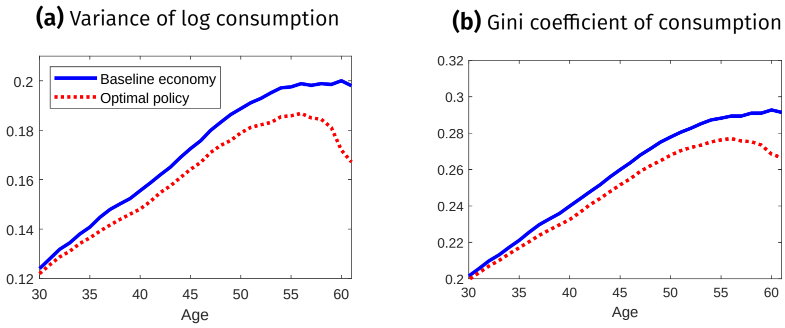

Notes: The left panel shows the variance of log consumption and the right panel the Gini coefficient of consumption over the life cycle. The solid lines display the inequality measure for the baseline economy and the dotted lines for the economy under the optimal pension system. The variances and Gini coefficients are computed using the whole sample.

towards the end of the working phase in the new economy. The reason is that the optimal pension system leads to an increase in pension wealth for workers with unstable jobs. The wealth effect induces these workers to raise their consumption close to retirement. This finding indicates that the increase in pension progressivity provides insurance to workers who face on average the highest level of job instability throughout life by redistributing resources to these workers. Because the pension system becomes more redistributive, agents in the bottom quartile are able to increase their consumption relative to their pre-retirement earnings in the late stage of the life cycle.

Figure 2.9d reflects this change in the consumption-saving behavior. For both groups of workers, the average asset level increases at the beginning of the life cycle, but declines and reaches the baseline level at the end of the working phase. Workers in the bottom quartile of the employment-stability distribution reduce their asset accumulation close to retirement as more resources are available to them in the retirement phase. Life-cycle consumption smoothing implies that these agents reduce their retirement savings such that the life-cycle profile of average assets becomes flatter under the optimal pension system.

Figure 2.10 displays that the optimal pension system reduces consumption inequality over the life cycle. The profile for variance of log consumption in the baseline economy almost linearly increases over life by more than 50% reaching a value of 0.2 before retirement. In the economy under the optimal pension system, the profile becomes flatter and has a concave shape. At the end of the working phase, the variance of log consumption reaches a value of 0.17. The optimal pension system therefore reduces the increase by almost one third of the increase observed in the baseline economy.

Figure 2.11. Transition to the optimal policy

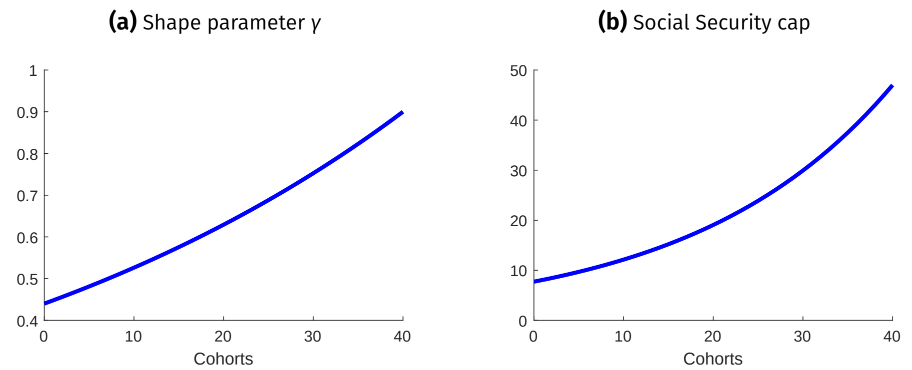

Notes: The left panel shows the shape parameter γ and the right panel the Social Security cap along the transition path to the optimal pension system. Cohort denotes the sequential order of cohorts that face the new policy parameters for the pension system.

Comparing the profiles of the Gini coefficient of consumption in Figure 2.10b, a similar change as for the profile of variance of log consumption is observed. Whereas the Gini coefficient increases linearly in the baseline model and reaches a value of 0.29 before retirement, the optimal pension system dampens the increase in the Gini coefficient over life. In the economy under the optimal pension system, the Gini coefficient of consumption at the end of the working phase decreases to 0.27.

2.6.7 Transition to the optimal pension system

The analysis so far has focused on the steady-state comparison of the baseline economy and the economy with the optimal pension system. The underlying assumption for the steady-state comparison is that workers already face the optimal pension system when entering the labor market. This subsection considers the transition path of the economy between these two steady states and evaluates the differential welfare effects across workers in different cohorts when the pension system gradually moves from the baseline specification towards the optimal system.

More specifically, the shape parameter γ and the Social Security cap are gradually raised by a constant percentage over a time horizon of 40 years. This is depicted in Figure 2.11. In year zero, the baseline pension system is implemented. Starting in year one, the government implements the new policies with higher parameter values of γ and the cap. Similar to Auerbach and Kotlikoff (1987), all workers have perfect foresight along the transition path regarding the pension system that will become effective when their cohort reaches the retirement age.

The welfare changes for the different cohorts are shown in Figure 2.12. These welfare changes are computed in the period the government credibly announces the entire path to the optimal pension system. As the policy parameters are changed gradually for each cohort and the remaining lifetime differs for different cohorts

Figure 2.12. Welfare along the transition path

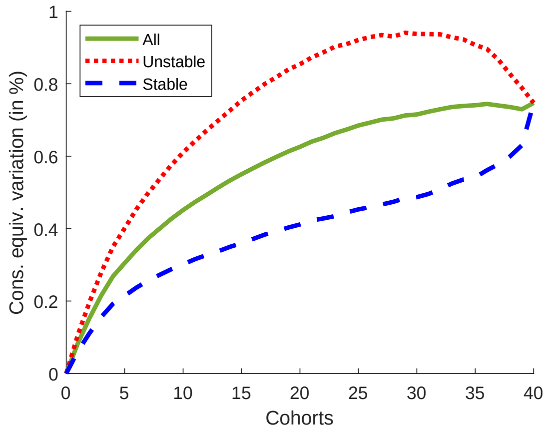

Notes: This figure shows the change in welfare in consumption-equivalent variation along the transition path to the optimal pension system. The solid line displays the welfare change for all workers. The dotted line and the dashed line show the welfare change for workers with unstable (bottom 25% of employment-stability distribution) and stable jobs (top 25% of employment-stability distribution), respectively.

until the retirement age, the transition to the optimal policy induces differential welfare effects across cohorts. The transition to the optimal policy induces for the majority of workers a positive welfare effect which is increasing from zero to 0.75% when reaching the steady state with the optimal pension system.

To analyze the differential welfare effects along the employment-stability distribution, I consider the bottom quartile (unstable) and top quartile (stable) of the employment-stability distribution of each cohort. As in the previous sections, this distribution considers the mean job separation rates faced by each worker after the labor market entry up to the time point of policy change announcement. The welfare gain along the transition path is larger for workers with unstable jobs compared to the average. The gap in welfare gain of unstable and stable workers diverges from the average in the numbering of cohorts, but converges eventually to the average as the policy change is implemented for workers closer to labor market entry. As workers at labor market entry have higher uncertainty regarding their future labor market outcomes, the gap in welfare gain is smaller for younger workers than for workers who already accumulated labor market experience. Workers who have unstable jobs after 10 years of labor market entry are more likely to remain in an unstable job in the future so that the induced welfare gain is larger for these workers. Younger workers, in contrast, have a higher chance to get a stable job. Thus, the welfare gain from a more progressive pension system is smaller for workers at labor market entry.

Note that older workers close to the retirement also have positive welfare gains from increasing the progressivity level. At age 50, workers with low employment stability (bottom 25% of employment-stability distribution) have a welfare gain of 0.6% in terms of life-time consumption. The welfare gain for high employment-stability workers (top 25% of the employment-stability distribution) is at 0.3% which is much lower compared to low-employment stability workers, but still positive. The reason why even older workers benefit from the implementation of the higher progressivity level is due to the abolishment of the Social Security cap. The tax revenue from high-income workers increases as the Social Security cap is abolished so that the vast majority of workers in the economy face a lower payroll tax rate. This induces a positive welfare effect even for most of the older workers close to the retirement.

2.7 Consequences of a shift in the employment-stability distribution in the U.S. labor market

This section analyzes the consequences of a shift in the employment-stability distribution on the optimal design of pension systems. Various reasons may account for a shift in the employment-stability distribution. Changes in the economic environment such as reforms of labor market policies, technological advances, changes in the industrial structure of the economy, or changes in search and matching frictions shape the labor market dynamism in the economy.

Intuitively, a decrease in job-separation rate leads to more stable work histories. Higher employment stability allows workers to invest in human capital, which enhances future career path and earnings growth of workers. From a policy perspective, an important question is how the optimal pension scheme should take into account such a shift in the employment-stability distribution. Should pension systems become more or less progressive in response to an increase in employment stability? I address this question by focusing on the recent development in the U.S. labor market which indicates a shift in the employment-stability distribution towards more stable jobs.

A number of studies have consistently found an increase in employment stability since the 1990s and an important role for short-duration jobs in explaining the observed changes in the labor market dynamism. Among others, Hyatt and Spletzer (2013) show that separations and hires have decreased in the recent years between 10% and 38%, and highlight the importance of the decline in short-duration jobs in explaining this trend. They point out that the decline in short-duration jobs explains nearly half of this decrease in hires and separations. In a later work, Hyatt and Spletzer (2016) show that the U.S. labor market features decreasing employment stability in the 1980s and 1990s, but that the recent data exhibit a reverse trend. Using the Current Population Survey, they show that the job tenure distribution indicates a move from unstable toward more stable jobs since 2000. Pries and Rogerson (2019) use the Quarterly Workforce Indicators and find that the decline in short-duration jobs which last for less than a quarter account much of the decrease in job separations in the U.S. labor market. They show that the decline in short-term employment is not caused by demographic or employer-related changes, but rather a shift in the labor market environment.

Given these empirical observations, I study the consequences of a shift in the employment-stability distribution on the optimal design of pension systems. To this end, I adjust the employment-stability distribution of the baseline economy in order to capture the empirical finding that the change in job-separation rates is largely driven by the extent of short-duration jobs. While holding the job-separation rate of the most stable job constant at the baseline value, I reduce the separation rate of the most unstable job. As a result, the support of the employment-stability distribution becomes narrower such that the degree of employment-stability heterogeneity declines. The average job-separation rate decreases by 6% in the economy with higher employment stability.[6]

2.7.1 Life-cycle consequences of a shift in the employment-stability distribution

How does a lower job-separation rate affect life-cycle dynamics? To address this question, I analyze the relative deviations from the baseline economy of human capital, earnings, consumption, and wealth for the economy with higher employment stability in Figure 2.13. Figure 2.13a shows that in the economy with higher employment stability, workers in both the top and bottom quartile achieve on average larger stocks of human capital at the early retirement age. Because of the increase in average employment stability, workers are less affected by career interruptions and thus make more human capital investments than in the baseline economy. The increase in human capital is larger for workers with most unstable jobs. Whereas human capital increases by 0.5% at the early retirement age for workers with most stable jobs, workers with unstable jobs achieve on average a human capital level that is 1.8% higher than the baseline level. Since employment stability increases for jobs with the largest job-separation rates, the effect of higher employment stability is strongest for workers who hold on average the most unstable jobs over the life cycle.

In Figure 2.13b, the average life-cycle profiles of labor earnings reflect these changes in human capital. Earnings become around 2.7% higher for workers with the most unstable jobs, and around 1% for workers with most stable jobs over life. The increase in average earnings is larger than the relative increases in human capital since higher employment stability not only leads to larger human capital stocks, but also facilitates moving up the job ladder such that workers find jobs

6. In Appendix 2.A.9, I show how the average separation rate changes over the life cycle in the economy with higher employment stability.

Figure 2.13. Relative change in average life-cycle profiles by employment stability

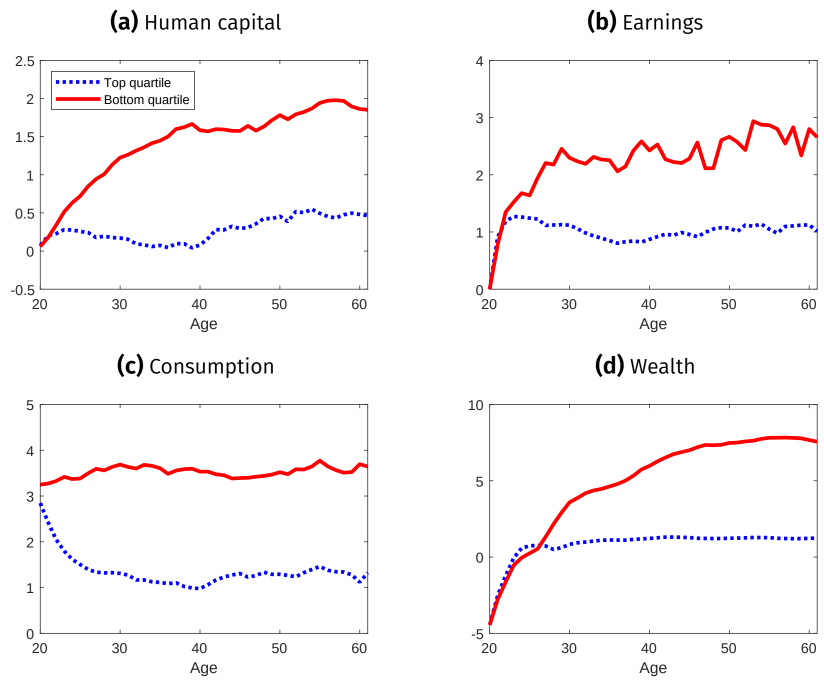

Notes: This figure shows the percent deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy with higher employment stability compared to the baseline economy. In all plots the solid line represents workers with most unstable jobs (bottom 25% of employment-stability distribution) and the dotted line represents workers with most stable jobs (top 25% of employment-stability distribution).

with higher wages. Moreover, both consumption and wealth increase in the economy with higher employment stability. More stable career paths and lower earnings losses due to higher employment stability enable workers to accumulate larger wealth. Before the early retirement age, workers with most unstable jobs have 7.5% higher wealth and workers with most stable jobs 1.2% higher wealth than in the baseline economy.

2.7.2 Distribution of employment stability and optimal pension systems

In order to analyze how these changes in life-cycle dynamics affect the optimal pension system in the economy with higher employment stability, I search for the optimal parameter combination of the parameter γ and the cap in the retirement benefit function. As in the previous section, policy changes are accompanied by a constant payroll tax rate on labor earnings and the parameter ϕ which scales the level of benefits in order to achieve budget balance for the pension system. The results are summarized in Table 2.4.

In the economy with higher employment stability, the optimal policy parameters γ and the cap remain at the same level as in the baseline economy. However, the

Table 2.4. Optimal pension systems in the baseline economy and in the economy with higher employment stability

Model	Optimal policy parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline	0.9	1.12	–	42.23	8.88	0.75
Higher employment stability	0.9	1.12	–	41.85	8.65	0.76

Notes: Column 1 specifies the model. Column 2-4 show the optimal pension system parameters γ and ϕ, which determine the shape and the level in the pension benefit formula, and the cap. Column 5 presents the degree of progressivity achieved by the combination of policy parameters and Column 6 the tax rate in the economy that achieve budget balance for the pension system. Column 7 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

progressivity level decreases in the economy with higher employment stability: The replacement rate decreases from 68% in the baseline to 67% for the median earner and from 88% to 84% for earners at the 25th percentile of the lifetime earnings distribution. At the same time, the replacement rate of earners at the 75th percentile increases slightly from 52% to 53%.

The reason why the optimal policy parameters do not change in the economy with higher employment stability is as follows. Higher employment stability shifts the distribution of lifetime earnings upwards so that the implied progressivity level by a given set of policy parameters becomes lower in the economy with higher employment stability compared to the baseline economy. Hence, even though the optimal progressivity level decreases in the economy with higher employment stability, the magnitude of the decline is small so that the policy parameters have to remain unchanged to achieve the optimal progressivity level. One important question is why the optimal progressivity level does not decrease strongly when employment stability increases. The reason is that higher employment stability does not decrease income inequality in the economy. In fact, the earnings variance does not decline despite the fact that higher employment stability increases earnings for low-income workers. As jobs get more stable, the earnings dispersion generated by the job ladder and stochastic human capital accumulation becomes more important so that on average, the earnings inequality does not decrease with the increase in employment stability at the bottom of the distribution. This implication of the model is compatible with the observation that labor market dynamism has been decreasing over time in the United States and the fact that income inequality has been stable or increasing as documented by a wide range of the literature (Braxton et al., 2021; Guvenen et al., 2021; Heathcote et al., 2023). As higher employment stability does not lead to a decrease in earnings inequality, the optimal progressivity level remains high.

Interestingly, the welfare gain from implementing the optimal pension system is higher in the economy with more stable jobs. This is because higher employment stability reduces the cost of a marginal increase in the progressivity level for low earners. Payroll taxes are welfare-detrimental for young workers who have to cope with a high risk of job loss. Payroll taxes depress consumption and saving, while a low degree of employment stability increases precautionary savings. In the presence of a borrowing constraint, both low employment stability and payroll taxes restrict the ability to smooth consumption over the life cycle, and these effects mutually amplify each other. In the economy with more stable jobs, a decline in unemployment risk and earnings losses reduces the cost of implementing a pension system with higher progressivity. This is reflected in the decrease of the payroll tax rate in the economy with higher employment stability. Compared to the baseline economy, the payroll tax rate associated with the optimal pension system drops by 23 basis points. Due to higher employment stability, the average labor earnings increase in the economy with higher employment stability, especially for workers with unstable jobs over life. Thus, the shift in the employment-stability distribution towards more stable jobs raises the tax revenue for a given amount of tax rate such that an increase in retirement benefits entails little adjustment of the tax rate, making the increase in pension progressivity less costly in the economy with higher employment stability. As a consequence, the welfare gain from implementing the optimal pension system is larger if jobs get more stable.

The implementation of a pensions system with higher progressivity is thus also important for the economy with higher employment stability as it leads to a considerable welfare gain. One key reason is that households are constrained in borrowing. In the absence of the borrowing constraint, agents can perfectly smooth consumption over life and the increase in employment stability makes a lower progressivity level and higher incentives for human capital investment desirable. This finding highlights the importance of incorporating life-cycle components and the role of incomplete markets in analyzing optimal pension systems, which is also a point made by Hubbard and Judd (1987).

2.8 Conclusion

This paper studies how a progressive pension system optimally considers heterogeneity in employment stability and quantifies the welfare gains from implementing the optimal pension system. Using a life-cycle model with heterogeneity in employment stability, endogenous human capital accumulation, and retirement decision, I find that abolishing the U.S. Social Security cap and increasing the degree of pension progressivity relative to the current U.S. pension system is optimal. Progressive pension systems provide insurance against bad labor market outcomes in the presence of heterogeneity in employment stability. A crucial consideration for the design of optimal pension systems is to weight the positive insurance effects of pension progressivity against its distortionary effects on human capital investment and retirement decision as well as the effects of the payroll tax finance of the pension system.

In a realistically calibrated life-cycle model with employment-stability heterogeneity, I show that heterogeneity in employment stability translates into a large inequality in labor market outcome and is a key driver of inequality in lifetime earnings. The numerical analysis in this paper indicates that an increase in pension progressivity for the U.S. economy induces a welfare gain of 0.75% in terms of lifetime consumption. Higher progressivity decreases consumption inequality over the life cycle and offers insurance to workers who suffer from job instability and low lifetime earnings, but distorts human capital investment and retirement decisions.

Motivated by the wide range of empirical findings in the literature that employment stability in the U.S. labor market has been increasing since the 1990s, I study the consequences of a shift in the employment-stability distribution on the optimal design of progressive pension systems. The model implications are consistent with the empirical findings, reconciling that we observe lower labor market dynamism, but at the same time the earnings inequality has been stable or increasing in the United States. The optimal pension system in the economy with higher employment stability implies a lower progressivity level, but the optimal policy parameters remain largely unchanged from the baseline economy. Importantly, the welfare gain from implementing the optimal pension system becomes even higher as employment stability increases. In the economy with more stable jobs, workers still value the insurance provided by the progressive pension system as heterogeneity in employment stability is still large, but the cost of implementing the optimal pension system decreases. Higher employment stability increases average earnings and the payroll tax rate associated with the optimal pension system becomes lower as employment stability increases. This reduces the negative impact of the optimal pension system on young workers for whom payroll taxes are very costly as they are constrained in borrowing. This cost of implementing a more progressive pension system is mitigated in the economy with higher employment stability and hence, the welfare gain from implementing the optimal pension system becomes higher. The importance of the optimal pension system does not diminish in the economy with higher employment stability as the degree of heterogeneity in employment stability is still large in today's economy.

2.A Appendix

2.A.1 Logistic distribution and expected utility

Let ε be a logistically distributed random variable. Let μ and σ denote the location parameter and the scale parameter of the Logistic distribution. Then, the cumulative distribution function of ε is given by

$$F(\varepsilon,\mu,\sigma) \;=\; \frac{1}{1+\exp\left(-\frac{\varepsilon-\mu}{\sigma}\right)}.$$

So as to economize on notation, I define $V_1 := V(a,w,\lambda,h,j)$ and $V_2 := V^r(a,w,h,j)$. A worker's decision problem is described by

$$V_{\max} \;=\; \max\{V_1, V_2+\varepsilon\}$$

The worker first observes the shock ε before making the decision. Let $\delta := V_1 - V_2$. Note that the worker chooses to remain in the labor force if

$$\begin{aligned} V_1 &> V_2+\varepsilon \\ \Leftrightarrow \qquad \varepsilon &< V_1 - V_2 \;=\; \delta \end{aligned}$$

Define $p := F(\delta,\mu,\sigma)$. Ex ante, the worker chooses to continue working with probability $F(\delta,\mu,\sigma)$. Therefore, the ex-ante expected utility is

$$\mathbb{E}V_{\max} \;=\; pV_1 + (1-p)V_2 + \int_{\delta}^{\infty} \varepsilon f(\varepsilon,\mu,\sigma)d\varepsilon \tag{2.A.1}$$

where

$$\int_{\delta}^{\infty} \varepsilon f(\varepsilon,\mu,\sigma)d\varepsilon \;=\; \int_{\delta}^{\infty} \varepsilon \frac{\exp\left(-\frac{\varepsilon-\mu}{\sigma}\right)}{\sigma\left(1+\exp\left(-\frac{\varepsilon-\mu}{\sigma}\right)\right)^2} d\varepsilon.$$

Integration by parts yields

$$\int_{\delta}^{\infty} \varepsilon f(\varepsilon,\mu,\sigma)d\varepsilon \;=\; \delta\cdot\frac{\exp\left(-\frac{\delta-\mu}{\sigma}\right)}{1+\exp\left(-\frac{\delta-\mu}{\sigma}\right)} + \sigma\cdot\log\left(1+\exp\left(-\frac{\delta-\mu}{\sigma}\right)\right). \tag{2.A.2}$$

Now, I use $p = F(\delta,\mu,\sigma)$ to get

$$\begin{aligned} p \;&=\; F(\delta,\mu,\sigma) \;=\; \left(1+\exp\left(-\frac{\delta-\mu}{\sigma}\right)\right)^{-1} \\ \Leftrightarrow \qquad \frac{1-p}{p} \;&=\; \exp\left(-\frac{\delta-\mu}{\sigma}\right) \\ \Leftrightarrow \qquad \delta \;&=\; -\sigma\cdot\log(1-p) + \sigma\cdot\log(p) + \mu. \end{aligned}$$

I use these expressions to simplify Equation (2.A.2) and get

$$\begin{aligned}
& \delta \cdot \frac{\exp(-\frac{\varepsilon-\mu}{\sigma})}{1+\exp(-\frac{\varepsilon-\mu}{\sigma})} + \sigma \cdot \log\left(1+\exp\left(-\frac{\varepsilon-\mu}{\sigma}\right)\right) \\
= \; & \delta \cdot \frac{\frac{1-p}{p}}{1+\frac{1-p}{p}} + \sigma \cdot \log\left(1+\frac{1-p}{p}\right) \\
= \; & \delta \cdot (1-p) - \sigma \log(p) \\
= \; & (-\sigma \cdot \log(1-p) + \sigma \cdot \log(p) + \mu) \cdot (1-p) - \sigma \log(p) \\
= \; & -\sigma\big((1-p) \cdot \log(1-p) + p \cdot \log(p)\big) + \mu \cdot (1-p).
\end{aligned}$$

Using this result, Equation (2.A.1) can be simplified to

$$\mathbb{E}\, V_{max} = pV_1 + (1-p)V_2 - \sigma\big((1-p)\log(1-p) + p \cdot \log(p)\big) + \mu \cdot (1-p)$$

which yields Equation (2.1) in Section 2.3.1.

2.A.2 Model solution

The model is solved on a discretized state space of wage, asset, human capital, and job-separation probability. The upper bound of the grids are chosen so as to not restrict the dynamic optimization of agents. Starting in the final period before death in which all remaining assets are consumed, the model is solved via backward induction applying on-grid search for consumption-saving decision, investment choices for human capital, and acceptance-rejection decision for outside job offers. Based on the computed policy functions I simulate life cycles for a population of $200,000$ agents.

2.A.3 Human capital

In this section, I discuss an alternative specification for the human capital accumulation process. In particular, I consider a human capital accumulation process that corresponds to the Ben-Porath (1967) model. While I leave all other parts of the model unchanged, I assume that human capital investment entails a monetary cost as a function of the investment time instead of a utility cost. The budget constraint of a worker is then

$$c = (1+r)a + y^n(w,h,e) - a' - \kappa \cdot t^2$$

where t is the time invested into human capital.

I solve the model with this alternative human capital accumulation and calibrate the model to the data. This model matches the empirical profiles well. However,

Figure 2.A.1. Earnings variance

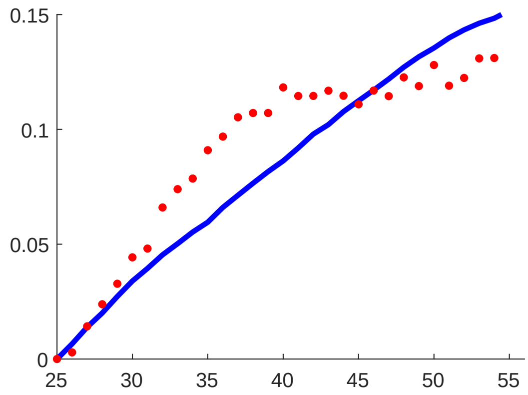

Notes: This figure shows the earnings variance in a model with Ben-Porath (1967) human capital accumulation.

Figure 2.A.1 shows that compared to the baseline model, the increase in earnings variance over the life cycle is slightly overstated as it increases almost linearly in age. The interaction between employment stability and the human capital accumulation explains why the increase in earnings variance is more pronounced if human capital investment entails monetary costs. Workers with unstable employment have on average lower income levels and at the same time, these workers have a large precautionary motive to save due to the high income risk they face. As a result, human capital investment is even costlier for workers with unstable employment if it enters the budget constraint. Lower investment leads to lower human capital accumulation of these workers, and consequently, the increase in earnings dispersion becomes stronger in the model with monetary cost of human capital investment.

While the model with Ben-Porath (1967) human capital accumulation also matches the data moments well, I assume that human capital investment leads to utility costs for the baseline economy so as to match the concavity of the earnings variance profile over the life cycle.

2.A.4 Calibration

2.A.4.1 Job-offer distribution

The number of gridpoints for wages is set to $N_w = 5$ where $\underline{w} = 1$ and $\overline{w} = 1.85$. All gridpoints are equidistant in logs. Concerning the job-separation probabilities, it is assumed that $N_\lambda = 10$ with $\underline{\lambda} = 0.006$ and $\overline{\lambda} = 0.35$. All remaining gridpoints are located non-linearly between these two values. The marginal distributions for wages and employment stability are assumed to have a truncated exponential distribution on the supports $[\underline{w}, \overline{w}]$ and $[1 - \overline{\lambda}, 1 - \underline{\lambda}]$, respectively. The joint distribution of job-separation probability and wage is determined by mapping the supports of these two random variables to the unit interval $[0, 1]$. Let $w^* \in [0, 1]$ denote the standardized wage level and let $1 - \lambda^*$ denote the standardized employment stability. The density of each of these standardized variables is then given by

$$f(x^*) \; = \; (1 - \exp(-\psi_x))^{-1}(\psi_x \exp(-\psi_x x^*))$$

with the shape parameter ψ_x and $x \in \{w^*, 1 - \lambda^*\}$. The job-offer distribution $f(w, \lambda)$ is derived by employing a copula C_θ and the correlation between standardized wage and employment stability is pinned down by the parameter θ. Table 2.A.1 presents the estimated parameters.

Figure 2.A.2. Joint distribution of wages and separation rates

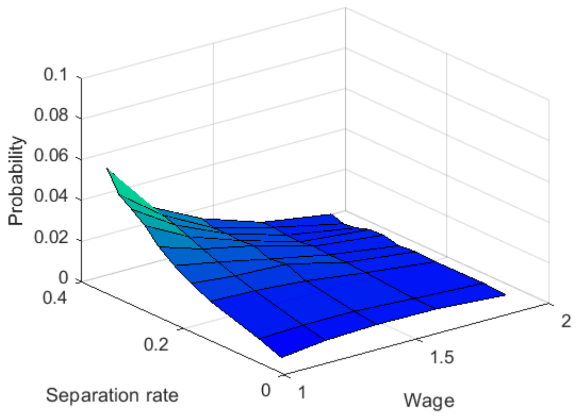

Notes: This figure shows the joint distribution of wages and separation rates.

Figure 2.A.2 shows the estimated distribution of wages and job-separation rates. In general, there is a negative relationship between employment stability and wages. The most stable jobs have high wages and vice versa. The vast probability mass lies in the area of high separation rate, reflecting the scarcity of lifetime jobs.

2.A.4.2 Parameters

I estimate the model parameters using a simulated method of moments. I minimize the sum of squared percentage deviations of the life-cycle profiles produced by the model from the empirical counterparts. The empirical moments include the life-cycle profiles of separation and job-finding rate as well as the transition rate from job to job, mean and variance profiles of earnings, the mean, median, and 75th percentile of the tenure distribution, and the wealth-to-income ratio. Let θ denote the vector of parameters and let a denote age. Then, the objective function is

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

$\pi_s(a,\theta)$ denotes the average separation rate from the model with the underlying parameter vector θ. π_{eo} and π_{ne} denote the job-to-job rate and the job-finding rate, respectively. t_{mean}, t_{median}, and t_{p75} denote the mean, median, and 75th percentile of the tenure distribution, accordingly. Finally, e_{mean} and e_{var} denote the mean and the variance of log earnings, and wti the wealth-to-income ratio. The corresponding empirical profiles are marked with a hat.

Table 2.A.1. Estimated parameters

Parameter	Value	Description
β	0.993	Quarterly discount factor
κ	0.361	Utility cost of effort
π_e	0.437	Probability of a job offer when employed
π_u	0.880	Probability of a job offer when unemployed
ψ_w	0.545	Marginal distribution of w^*
ψ_λ	0.464	Marginal distribution of $1-\lambda^*$
θ	0.435	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.052	Skill upgrading probability
ρ	0.983	Persistence of skill upgrading probability
p_H^*	0.058	Probability to move to h^*

2.A.4.3 Approximation of average lifetime earnings

To evaluate the accuracy of the approximation of average lifetime earnings by the labor earnings at the end of the working phase, I proceed as follows. First, for each worker, I take the 35 years with the highest labor earnings of the worker and compute the average lifetime earnings in these years. Then, I derive the approximations to these average lifetime earnings by regressing them on cubic polynomials of the labor earnings before retirement of each worker. This yields an R^2 statistic of 0.902. The labor earnings in the last period thus approximates the average lifetime earnings well, explaining 90.2% of the variance in the average lifetime earnings.

2.A.5 Persistence of initial employment stability

Using the Panel Study of Income Dynamics (PSID), I show that the persistence of employment stability in the model is consistent with its empirical counterpart. The PSID is a longitudinal survey data on U.S. households and contains information on the labor market situation of households. The PSID surveys first started in 1968 and were conducted annually until 1997 and biennially thereafter.

I restrict the sample to household heads between age 25 and 50 and keep household heads who are not permanently out of labor force. Based on this sample, I compute the cumulative sum of nonemployment duration at each age of household heads starting at age 25. This yields a distribution of cumulative nonemployment duration at each age. Then, using the median of the cumulative nonemployment duration at age 30, I assign individuals to a stable and unstable groups at age 30. Similarly, I use the median of the cumulative nonemployment duration at age 50 to assign individuals to stable and unstable groups. A worker can be either in the stable or unstable group when young (age 30), and either in the stable or unstable group when old (age 50). If employment stability is persistent, workers in the stable (unstable) group at age 30 will remain in the stable (unstable) group at age 50 with a probability $> 50\%$. That is, for a worker in the stable group when young, it is more likely that the worker will remain in the stable group in the future.

The results in Table 2.A.2 indicate that employment stability at age 30 is persistent in the data. A worker with an unstable job at age 30 remains in the unstable group at age 50 with a probability of 70%. The model also produces a high persistence rate of employment stability, while it produces a higher persistence rate for the stable group.

Figure 2.A.3 shows the life-cycle profiles of human capital, earnings, consumption, and wealth-to-income ratios as in Figure 2.6. The profiles are again normalized by the initial levels and expressed in log deviations. In contrast to Figure 2.6, however, the top and bottom quartile of the employment stability distribution are determined based on the first ten years of labor market history rather than the whole

Table 2.A.2. Persistence rate of employment stability at age 30

	Data	Model
Stable	0.61	0.70
Unstable	0.69	0.69

Notes: This table shows the persistence rates of employment stability at age 30, that is, the probability that an individual in the stable (unstable) group at age 30 remains in the stable (unstable) group at age 50. Individuals below (above) the median of the cumulative nonemployment duration at the respective age is defined as the stable (unstable) group. The first column shows the persistence rates in the data and the second column the persistence rates in the model. The underlying data is from the Panel Study of Income Dynamics.

Figure 2.A.3. Average life-cycle profiles by employment stability

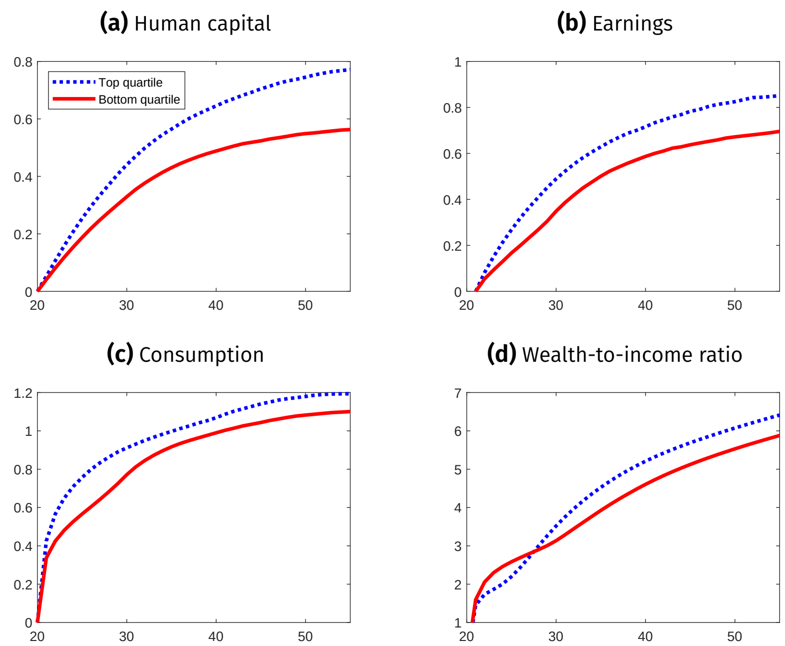

Notes: This figure shows the average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), and consumption (lower left panel), and wealth (lower right panel). The profiles are the log deviations from the initial values of the respective profiles. In all plots the solid line represents workers with most unstable jobs (bottom 25% of employment-stability distribution) and the dashed line represents workers with most stable jobs in the first ten years after labor market entry (top 25% of employment-stability distribution).

working life. Due to persistence of employment stability, worker with unstable jobs at the beginning of the career have on average unstable jobs over the whole life cycle. As a consequence, we observe large gaps in the labor market outcomes as shown in Figure 2.A.3 where the two groups of workers were determined based on the employment stability in the first ten years since labor market entry.

2.A.6 Initial conditions and welfare effects of pension systems

To analyze the importance of the initial state of workers at labor market entry in studying the optimal progressive pension system, I derive the welfare changes of the ex-ante optimal pension policy as functions of separation rate, wealth, and human capital. More precisely, while holding all other state variables constant at the initial state of a worker at labor market entry, I assess how variations in separation rate, wealth, and human capital each affect the welfare gain from the optimal pension policy. The results are displayed in Figure 2.A.4.

Figure 2.A.4a shows that the welfare change increases in job-separation rate, that is, a worker with an initially unstable job gains more from an increase in pension progressivity. The redistributive effect of the optimal policy provides insurance against future poor labor market outcomes and leads to a welfare gain for a worker with an unstable job. For a worker with a very stable job, however, prospective earnings and human capital growth are larger. The expected gain from higher pension progressivity therefore strictly decreases in employment stability.

Figure 2.A.4b presents the welfare change as a function of the wealth level. Interestingly, the function is non-monotonic: the welfare change exhibits a hump close to zero wealth, but then, decreases steadily in the asset level. Starting from zero wealth, higher wealth allows for a better ability of consumption smoothing in the presence of the payroll tax rate. Because young workers have relatively low labor earnings and they save a large proportion of their earnings for precautionary reasons, an increase in wealth leads to higher welfare gains.

However, as the amount of wealth increases further, the total amount of lifetime resources available increases in a way that the insurance effect of higher pension progressivity becomes less important. Still, the welfare change from the optimal progressive pension system remains positive since prospective labor market outcome remains uncertain and pension progressivity provides insurance against this risk.

Finally, Figure 2.A.4c shows the welfare change as a function of human capital. Starting from the initial level, a larger stock of human capital implies a decrease in welfare. This points out that workers with low human capital gain from an increase in pension progressivity, while workers with high human capital lose out. An increase in pension progressivity decreases the welfare gain of workers with high human capital in two respects. Firstly, higher pension progressivity reduces pension wealth and therefore the lifetime resources of workers with high human capital. Secondly, despite the decrease in the payroll tax rate due to the abolishment of the Social Security cap, contributions to the pension system increase relative to benefits. As a consequence, the increase in pension progressivity decreases welfare.

Figure 2.A.4. Welfare change as a function of the initial state variables of separation rate, wealth, and human capital

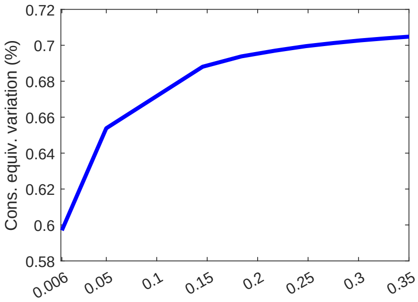

(a) Separation rate

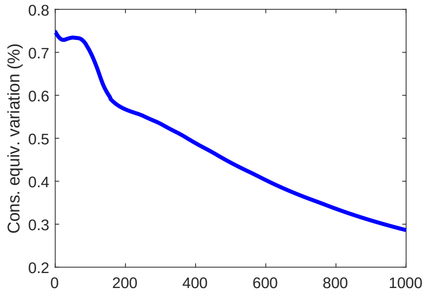

(b) Wealth

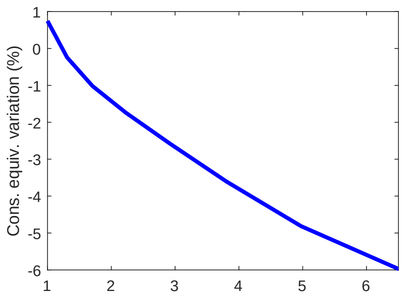

(c) Human capital

2.A.7 Model extensions

2.A.7.1 Bequest motive, mortality risk, and early application penalty

While the baseline model assumes a deterministic lifespan and consumption of all remaining assets by the end of the life cycle, this section extends the framework to include bequest motive, mortality risk, and early application penalty for the retirement benefits. For the bequest motive, I follow De Nardi (2004) and assume that bequests yield utility of

$$\upsilon(a) = \nu \cdot u(a)$$

where the parameter ν captures the strength of the bequest motive and is set to 25. In the extended framework, retired households face death probabilities, which alter their value functions as follows:

$$\begin{aligned} V^r(a,w,h,j) \;=\; & \max_{a' \geq 0} u\big((1+r)a + y_r(w,h,n) - a'\big) \\ & + \beta\big[(1-d_{j+1})V^r(a',w,h,j+1) + d_{j+1} \cdot \upsilon(a')\big]. \end{aligned}$$

Here, d_{j+1} denotes the age-specific probability of death in the next period. The death probabilities are taken from the decennial life tables (1979-1981) published by the National Center for Health Statistics.[7]

For early claiming penalties in the Social Security system, I apply the reduction rates summarized in Table 2.A.3.

Table 2.A.3. Social Security benefit reduction by claiming age

Age at claiming	Total reduction (%)
67	0.00%
66	6.67%
65	13.33%
64	20.00%
63	25.00%
62	30.00%

Notes: Source: https://www.ssa.gov/benefits/retirement/planner/agereduction.html.

Table 2.A.4 repeats the optimal pension system analysis from Section 2.6 for the model with mortality risk and a bequest motive to assess how these extensions affect the design and outcomes of the optimal policy. While the welfare gain from implementing the optimal policy declines slightly from 0.75 to 0.49, the optimal policy parameters in the extended model remain almost unchanged.

While disability insurance is a crucial component of social insurance systems, the current model abstracts from this type of risk. One main reason for this is to isolate the effects of heterogeneity in labor market risk on pension outcomes. It is important to note that disability insurance is designed to offer insurance against health-related income loss, whereas the focus of this paper is on how to redistribute between workers with stable and unstable employment.

7. Table 1, decennial life tables 1979-1981. See http://www.cdc.gov/nchs/products/life_tables.htm.

Table 2.A.4. Welfare change in the baseline economy under alternative pension systems

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.72	7.72	-
No cap	0	0.42	–	−0.41
	0.2	0.54	–	−0.10
	0.4	0.68	–	0.15
	0.6	0.84	–	0.35
	0.8	1.02	–	0.49
	1	1.24	–	0.11
	1.2	1.44	–	−0.72
	1.4	1.60	–	−1.14

Notes: The results for the model "No cap" are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

2.A.7.2 Different levels of income tax progressivity

Next, I examine how changes in the progressivity of labor income taxation influence the design of the optimal policy. To do so, I adopt the tax function specification from Heathcote, Storesletten, and Violante (2017), in which net labor income is given by

$$y_{\text{net}} = \alpha_1 \cdot y^{1-\alpha_2}.$$

The welfare effects of pension reform depend on the level of income tax progressivity. As income tax progressivity increases, the welfare gains from introducing greater progressivity in the pension system are reduced. Nevertheless, the qualitative result remains robust: a more progressive pension system is optimal even in the presence of a highly progressive income tax. For instance, in a scenario where the income tax progressivity parameter is set to 0.45, a value significantly higher than the empirical estimate of approximately 0.18 reported by Heathcote, Storesletten, and Violante (2017), the reform still yields welfare gains exceeding 0.3% of lifetime consumption.

Table 2.A.5. Economy with income tax progressivity α_2 set to 0.15

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.72	7.72	-
No cap	0	0.43	–	−1.52
	0.2	0.55	–	−0.80
	0.4	0.68	–	−0.25
	0.6	0.84	–	0.15
	0.8	1.02	–	0.46
	1	1.21	–	0.90
	1.2	1.34	–	0.69
	1.4	1.45	–	0.12

Notes: The results for the model "No cap" are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

Table 2.A.6. Economy with income tax progressivity α_2 set to 0.3

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.72	7.72	-
No cap	0	0.43	–	−0.58
	0.2	0.55	–	−0.15
	0.4	0.68	–	0.16
	0.6	0.84	–	0.37
	0.8	1.00	–	0.52
	1	1.19	–	0.63
	1.2	1.31	–	0.40
	1.4	1.40	–	0.23

Notes: The results for the model "No cap" are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

Table 2.A.7. Economy with income tax progressivity α_2 set to 0.45

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.72	7.72	-
No cap	0	0.44	–	−0.44
	0.2	0.56	–	0.00
	0.4	0.69	–	0.14
	0.6	0.83	–	0.27
	0.8	0.99	–	0.33
	1	1.16	–	0.35
	1.2	1.25	–	0.00
	1.4	1.31	–	−0.32

Notes: The results for the model "No cap" are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

2.A.7.3 Different unemployment insurance systems

In the following, I search for the optimal pension system in economies with alternative unemployment insurance systems. Specifically, I consider replacement rates of 0.8, 0.6, and 0.2 for the unemployment insurance system. The replacement rate in the baseline economy is 0.4.

Table 2.A.8. Optimal policy in economies with different unemployment insurance systems

Replacement rate of UI	Optimal γ	Welfare change (%)
0.8	0.9	0.68
0.6	0.9	0.73
0.4	0.9	0.75
0.2	0.9	0.77

Notes: The optimal pension systems are obtained by varying the parameters γ and the Social Security cap starting from the baseline economy. For a given level of γ, parameter ϕ and the payroll tax rate are always set to satisfy budget balance for the government. Column 1 specifies the replacement rate of the unemployment insurance system. Column 2 shows the parameter of the optimal pension system. Column 3 presents the welfare change relative to the baseline economy as equivalent variation in consumption in percentages.

Table 2.A.8 summarizes the results. One striking result is that the optimal degree of progressivity does not change for the pension system under alternative unemployment insurance systems. In all cases, it is optimal to remove the Social Security cap and the shape parameter γ of the pension benefit formula takes a value of 0.9. Note, however, that the induced welfare gains are decreasing in the replacement rate of the unemployment insurance system. As workers with frequent job interruptions get a higher replacement rate, the welfare gain from redistribution through the pension system decreases. However, whereas a higher replacement rate of the unemployment insurance system helps workers to smooth consumption upon job interruptions, it does not compensate for the potential loss of human capital. Hence, the optimal pension progressivity remains unaffected by the unemployment insurance system.

2.A.8 Human capital depreciation

In the baseline model, human capital does not depreciate when workers are unemployed. To analyze the potential consequences of human capital depreciation on the design of optimal pension systems, I assume that unemployed workers face a probability of 0.12 of human capital depreciation. This probability implies a yearly depreciation rate of approximately 10 percent which is in range of human capital deprecation rate in the literature. For example, using the NLSY data, Keane and Wolpin (1997) find a yearly depreciation rate of 9.6 percent and 36.5 percent for human capital. Jacobson, LaLonde, and Sullivan (1993) estimate depreciation rates between 10 percent and 25 percent.

I calibrate the model to the empirical moments as for the baseline economy. To find the optimal policy parameters, I again vary the shape parameter γ of the benefit function and the Social Security cap where the pension system achieves budget balance and the total amount of pension benefits correspond to those in the economy with initial pension parameters.

Table 2.A.9. Optimal pension systems in the economy with human capital depreciation

Model	Optimal policy parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline cap	0.64	0.86	6.67	35.01	9.45	0.19
No cap	0.86	1.02	–	41.93	9.16	0.71

Notes: Column 1 specifies the model. Column 2-4 show the optimal pension system parameters γ and ϕ, which determine the shape and the level in the pension benefit formula, and the cap. Column 5 presents the degree of progressivity achieved by the combination of policy parameters and Column 6 the tax rate in the economy that achieve budget balance for the pension system. Column 7 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

Table 2.A.9 summarizes the optimal policy parameters. Similar to the baseline model, removing the Social Security cap is optimal and the shape parameter γ increases from 0.44 to 0.64 (with cap) and 0.86 (no cap) so that the progressivity level increases under the optimal policy. The inclusion of human capital depreciation does not change the previous results qualitatively.

Figure 2.A.5. Average life-cycle profile of separation rate for the baseline economy and for the economy with higher employment stability

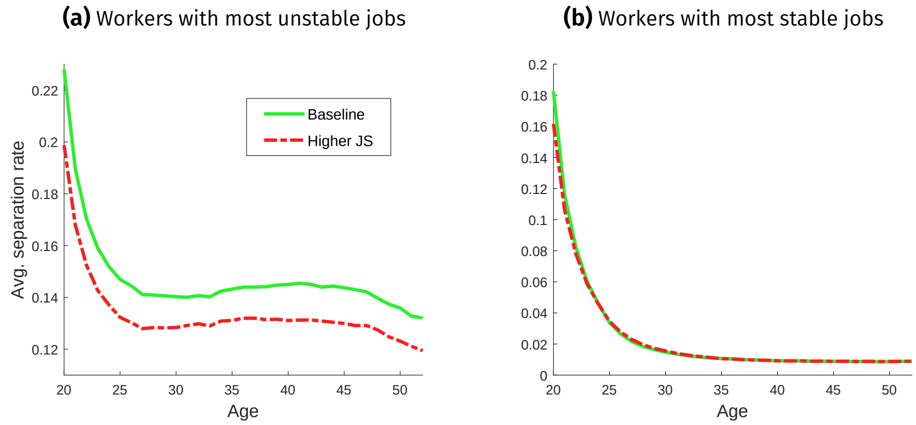

Notes: The left panel shows the average separation rate of workers with most unstable jobs (bottom 25% of employment-stability distribution) and the right panel the average separation rate of workers with most stable jobs (top 25% of employment-stability distribution).

2.A.9 Shift in employment-stability distribution

Figure 2.A.5 displays the average separation rate by age in the baseline economy and the economy with higher employment stability for workers in the bottom and top quartile of the average employment-stability distribution over the life cycle. Relative to the average separation rate in the baseline economy (solid line), the dashed-dotted line shows that the life-cycle profile of average job-separation rate of workers with most unstable jobs shifts downwards in the economy with higher employment stability. At labor market entry, the average separation rate decreases by 3 percentage points for workers with most unstable jobs. This gap declines over the life cycle. A shift in the lower tail of the employment-stability distribution primarily affects the average separation rate of young workers because young workers also accept unstable jobs: Employment offers the opportunity to invest in human capital and to increase prospective earnings growth, and is therefore highly valuable for young workers. Over time, workers climb the job ladder and an average worker finds more stable jobs as the worker spends more time in the labor market. Note that the separation rates of the most stable jobs remain unchanged in the recalibrated economy. By construction, workers with most stable jobs are little affected by the shift in the employment-stability distribution and their profile of average separation rate remains almost unchanged.

References

Ahn, Hie Joo, Bart Hobijn, and Ayşegül Şahin. 2023. “The Dual U.S. Labor Market Uncovered.” Working Paper, Working Paper Series 31241. National Bureau of Economic Research. [65, 69, 84]

Auerbach, Alan J, and Laurence J Kotlikoff. 1987. “Evaluating fiscal policy with a dynamic simulation model.” *American Economic Review* 77 (2): 49–55. [97]

Bairoliya, Neha, and Kathleen McKiernan. 2021. “Revisiting retirement and social security claiming decisions.” *Available at SSRN* 3896031. [95]

Becker, Gary S. 1964. “Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.” University of Chicago Press. [66]

Ben-Porath, Yoram. 1967. “The production of human capital and the life cycle of earnings.” *Journal of Political Economy* 75 (4): 352–65. [66, 72, 106, 107]

Bénabou, Roland. 2000. “Unequal Societies: Income Distribution and the Social Contract.” *American Economic Review* 90 (1): 96–129. [77]

Bénabou, Roland. 2002. “Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?” *Econometrica* 70 (2): 481–517. [77]

Bilal, Adrien. 2023. “The Geography of Unemployment.” *Quarterly Journal of Economics* 138 (3): 1507–76. [69]

Bloom, David E., David Canning, Richard K. Mansfield, and Michael Moore. 2007. “Demographic change, social security systems, and savings.” *Journal of Monetary Economics* 54 (1): 92–114. [70]

Bonhomme, Stéphane, and Jean-Marc Robin. 2009. “Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990–2000.” *Review of Economic Studies* 76 (1): 63–92. [70]

Bönke, Timm, Giacomo Corneo, and Holger Lüthen. 2015. “Lifetime Earnings Inequality in Germany.” *Journal of Labor Economics* 33 (1): 171–208. [70]

Borella, Margherita, Mariacristina De Nardi, and Fang Yang. 2022. “Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labour Supply?” *Review of Economic Studies* 90 (1): 102–31. [70]

Bowlus, Audra J., and Jean-Marc Robin. 2012. “An International Comparison of Lifetime Inequality: How Continental Europe Resembles North America.” *Journal of the European Economic Association* 10 (6): 1236–62. [70]

Braxton, J. Carter, Kyle F Herkenhoff, Jonathan L Rothbaum, and Lawrence Schmidt. 2021. “Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter.” Working Paper, Working Paper Series 29567. National Bureau of Economic Research. [68, 102]

Brown, Kristine M. 2013. “The link between pensions and retirement timing: Lessons from California teachers.” *Journal of Public Economics* 98: 1–14. [81]

Burtless, Gary. 1986. “Social Security, Unanticipated Benefit Increases, and the Timing of Retirement.” *Review of Economic Studies* 53 (5): 781–805. [81]

Coile, Courtney, and Jonathan Gruber. 2001. “Social Security Incentives for Retirement.” National Bureau of Economic Research, Working Paper 7830. [70]

Coile, Courtney, and Jonathan Gruber. 2007. “Future Social Security Entitlements and the Retirement Decision.” *Review of Economics and Statistics* 89 (2): 234–46. [81]

Cremer, Helmuth, Jean-Marie Lozachmeur, and Pierre Pestieau. 2004. “Social security, retirement age and optimal income taxation.” *Journal of Public Economics* 88 (11): 2259–81. [70]

De Nardi, M, S Imrohoroglu, and TJ Sargent. 2001. “Saving and pension reform in general equilibrium models.” *Oxford Review of Economic Policy* 17 (1): 20–39. [70]

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies* 71 (3): 743–68. [114]

Disney, Richard, Carl Emmerson, and Matthew Wakefield. 2006. "Ill health and retirement in Britain: a panel data-based analysis." *Journal of Health Economics* 25 (4): 621–49. [80]

Dwyer, Debra S., and Olivia S. Mitchell. 1999. "Health problems as determinants of retirement: are self-rated measures endogenous?" *Journal of Health Economics* 18 (2): 173–93. [80]

Fehr, Hans, and Christian Habermann. 2008. "Risk Sharing and Efficiency Implications of Progressive Pension Arrangements." *Scandinavian Journal of Economics* 110 (2): 419–43. [70]

Fehr, Hans, Manuel Kallweit, and Fabian Kindermann. 2013. "Should pensions be progressive?" *European Economic Review* 63: 94–116. [70]

Feldstein, Martin. 1974. "Social Security, Induced Retirement, and Aggregate Capital Accumulation." *Journal of Political Economy* 82 (5): 905–26. [70]

Galama, Titus, Arie Kapteyn, Raquel Fonseca, and Pierre- Carl Michaud. 2013. "A health production model with endogenous retirement." *Health economics* 22 (8): 883–902. [80]

Golosov, Mikhail, Ali Shourideh, Maxim Troshkin, and Aleh Tsyvinski. 2013. "Optimal Pension Systems with Simple Instruments." *American Economic Review* 103 (3): 502–7. [70]

Gregory, Victoria, Guido Menzio, and David G Wiczer. 2021. "The Alpha Beta Gamma of the Labor Market." Working Paper, Working Paper Series 28663. National Bureau of Economic Research. [65, 69, 84]

Gruber, Jonathan, and David A. Wise. 2004. *Social Security Programs and Retirement around the World: Micro-Estimation.* University of Chicago Press. [70]

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2021. "Lifetime Earnings in the United States over Six Decades." University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2021-60. [68, 102]

Haan, Peter, Daniel Kemptner, and Victoria Prowse. 2019. "Insurance, Redistribution, and the Inequality of Lifetime Income." Rationality and Competition Discussion Paper Series 188, CRC TRR 190 Rationality and Competition. [70]

Haan, Peter, and Victoria Prowse. 2014. "Longevity, life-cycle behavior and pension reform." *Journal of Econometrics* 178 (3): 582–601. [70]

Hall, Robert. 1982. "The Importance of Lifetime Jobs in the U.S. Economy." *American Economic Review* 72 (4): 716–24. [65]

Hall, Robert, and Marianna Kudlyak. 2019. "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics." Working Paper, Working Paper Series 25625. National Bureau of Economic Research. [69, 84]

Heathcote, Jonathan, Fabrizio Perri, Giovanni L Violante, and Lichen Zhang. 2023. "More Unequal We Stand? Inequality Dynamics in the United States 1967–2021." Working Paper, Working Paper Series 31486. National Bureau of Economic Research. [68, 102]

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *Quarterly Journal of Economics* 132 (4): 1693–754. [77, 117]

Hubbard, R. Glenn, and Kenneth L. Judd. 1987. "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax." *American Economic Review* 77 (4): 630–46. [70, 103]

Hubmer, Joachim. 2018. "The job ladder and its implications for earnings risk." *Review of Economic Dynamics* 29: 172–94. [66]

Hyatt, Henry R., and James R. Spletzer. 2013. "The recent decline in employment dynamics." *Journal of Labor Economics* 2 (5). [68, 99]

Hyatt, Henry R., and James R. Spletzer. 2016. “The shifting job tenure distribution.” *Labour Economics* 41: 363–77. [99]

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan. 1993. “Earnings Losses of Displaced Workers.” *American Economic Review* 83 (4): 685–709. [120]

Jarosch, Gregor. 2023. “Searching for Job Security and the Consequences of Job Loss.” *Econometrica* 91 (3): 903–42. [65, 66, 69]

Jones, Andrew M., Nigel Rice, and Jennifer Roberts. 2010. “Sick of work or too sick to work? Evidence on self-reported health shocks and early retirement from the BHPS.” *Economic Modelling* 27 (4): 866–80. [80]

Jung, Philip, and Moritz Kuhn. 2019. “Earnings Losses and Labor Mobility Over the Life Cycle.” *Journal of the European Economic Association* 17 (3): 678–724. [65, 66, 69]

Keane, Michael P., and Kenneth I. Wolpin. 1997. “The Career Decisions of Young Men.” *Journal of Political Economy* 105 (3): 473–522. [120]

Kimball, Miles S., and Mattew D. Shapiro. 2003. “The Retirement Elasticity: Theory and Implications.” Working paper. [81]

Krueger, Alan B., and Jörn-Steffen Pischke. 1992. “The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation.” *Journal of Labor Economics* 10 (4): 412–37. [70, 81]

Kuhn, Moritz, Leanne Nam, and Gašper Ploj. 2025. “Job Stability, Earnings Dynamics, and Life-Cycle Savings.” Working paper. [65, 66, 69]

Larkin, Kieran. 2019. “Job risk, separation shocks and household asset allocation.” In *2019 Meeting Papers.* 1058. Society for Economic Dynamics. [69]

Michelacci, Claudio, and Hernán Ruffo. 2015. “Optimal Life Cycle Unemployment Insurance.” *American Economic Review* 105 (2): 816–59. [67, 95]

Molloy, Raven, Christopher Smith, and Abigail K. Wozniak. 2020. “Changing Stability in U.S. Employment Relationships: A Tale of Two Tails.” Working Paper, Working Paper Series 26694. National Bureau of Economic Research. [65, 68, 69]

Morchio, Iacopo. 2020. “Work Histories and Lifetime Unemployment.” *International Economic Review* 61 (1): 321–50. [69]

Moulton, Jeremy Grant, and Ann Huff Stevens. 2015. “Reconsidering the Social Security Notch and retirement: Wealth and incentive effects.” *Economics Letters* 132: 65–68. [81]

O’Dea, Cormac. 2018. “Insurance, Efficiency and the Design of Public Pensions.” 2018 Meeting Papers 1037. Society for Economic Dynamics. [70]

OECD. 2019. *Pensions at a Glance 2019.* [78]

Pinheiro, Roberto, and Ludo Visschers. 2015. “Unemployment risk and wage differentials.” *Journal of Economic Theory* 157: 397–424. [66]

Pries, Michael J., and Richard Rogerson. 2019. “Declining Worker Turnover: the Role of Short Duration Employment Spells.” National Bureau of Economic Research, Working Paper. [68, 99]

Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review* 95 (1): 25–49. [76]

Trevisan, Elisabetta, and Francesca Zantomio. 2016. “The impact of acute health shocks on the labour supply of older workers: Evidence from sixteen European countries.” *Labour Economics* 43: 171–85. [80]

Chapter 3

Intergenerational Spillovers: The Impact of Labor Market Risk on the Housing Market*

3.1 Introduction

Recessions lead to significant welfare losses for workers, who suffer persistent earnings losses upon unemployment (Jacobson, LaLonde, and Sullivan, 1993; Davis et al., 2011). Such labor market shocks can spill over into other markets and affect households that do not participate in the labor market. A prime candidate for such spillovers is the housing market, where younger, working-age households often drive housing demand and older, retired households seeking to sell represent the supply side. The housing market can form an "intergenerational hinge," linking the labor market situation of young, working-age households with older households. Despite the importance of such intergenerational spillovers, little is known about how unemployment risk affects housing markets. This paper aims to fill this gap.

In this paper, I provide a quantitative evaluation of the spillover effect of unemployment risk from younger, working-age households to older, retired households through the housing market from both empirical and theoretical perspectives. I empirically document that the spillover effect of unemployment on the housing market is quantitatively important. Motivated by these findings, I develop a theoretical framework to analyze the welfare consequences of the spillover effect of unemployment shocks on different age groups of households. In addition, I examine how the

* I thank Moritz Kuhn for invaluable advice, guidance, and encouragement throughout this project. I also thank Christian Bayer, Thomas Hintermaier, Mariacristina De Nardi, Dirk Krueger, Joseph Mullins, Adam Guren, Loukas Karabarbounis, Kyle Herkenhoff, Jeremy Lise, Anmol Bhandari, and Brigitte Hochmuth as well as participants of numerous conferences and seminars for their helpful feedback and comments. Support by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A03) is gratefully acknowledged. The usual disclaimer applies.

age structure of the unemployment increase and the demographic composition of the population shape the spillover effect and whether increasing the generosity of the unemployment insurance system during recessions can stabilize housing prices.

In the first part of this paper, I analyze the spillover effect of unemployment risk on the housing market using data from the Current Population Survey (CPS). Specifically, I document the existence of the spillover effect in the United States by showing that changes in the unemployment rate have a significant effect on housing prices. To this end, I use an instrumental variable approach with a shift-share instrument which helps to address potential endogeneity and omitted variable bias by identifying exogenous fluctuations in state-level unemployment rates. The CPS data provide information on the state-level unemployment rate and the industry composition of employment needed to construct the shift-share instrument. The results indicate that the effect is also economically significant: a one percentage point increase in the unemployment rate leads to a 1.55% decline in housing prices. An important driver of this spillover is changes in housing demand in response to fluctuations in the labor market. To show that housing demand tends to decline during periods of high unemployment, I use data from the Home Mortgage Disclosure Act (HMDA), which provides data on mortgage applications by U.S. households. I analyze the relationship between mortgage applications and the unemployment rate at the state level using the number of mortgage applications as a proxy for housing demand. The results show a strong negative correlation between the number of mortgage applications submitted and the unemployment rate at the state level. Moreover, I document that financial institutions reject a higher proportion of mortgage applications when the unemployment rate is high, further reducing effective housing demand, in addition to fewer mortgage applications. These results provide evidence that housing demand tends to be substantially lower during periods of high unemployment.

Next, I construct a general equilibrium model with overlapping generations to study the mechanism and the welfare consequences of the spillover effect of unemployment shocks on the housing market. Workers choose their portfolio between housing and liquid assets, where housing prices are determined on the housing market. The life cycle of a household is divided into a working phase and a retirement phase. During the working phase, households face an exogenous income process with unemployment risk. Retired households are not subject to labor market risk because they receive retirement benefits. When buying a house, households take out a mortgage subject to constraints on payment-to-income and loan-to-value ratios. The model is calibrated to the U.S. economy using data from the Survey of Consumer Finances (SCF). The calibrated model fits the data well along many dimensions, including the life-cycle profiles of housing wealth, liquid wealth, and household leverage.

Using the model, I analyze the aggregate consequences of an unemployment shock on the housing market. The spillover effect of an unemployment shock on housing prices in the model is quantitatively similar to the empirically documented magnitude. Specifically, an increase in the average unemployment rate from 4% to 8% leads to a 5% drop in housing prices. Looking at the mechanism behind the price change, I find that more than 60% of this spillover effect is driven by the increase in income uncertainty when the unemployment shock occurs, which leads households to reduce their demand for housing. In contrast, the actual decline in household income due to higher unemployment accounts for 40% of the total spillover effect. When the unemployment shock occurs, households shift their portfolios away from illiquid housing assets to liquid financial assets. Although retired households are not directly affected by changes in unemployment risk, the spillover effect from the labor market on the housing market leads to a decline in housing prices, which reduces the housing wealth of retired households. Retired households lose out from a decline in wealth due to their bequest motive. In response to the decline in housing wealth, retired households reduce their non-durable consumption and increase their savings.

I explore the welfare consequences of unemployment shocks across different age groups of households. The spillover effect of the unemployment shock leads to significant welfare losses for older, retired households, transmitting approximately one-third of the welfare losses of working-age households to retired households measured in terms of remaining lifetime consumption. Young workers face large welfare losses because they are directly affected by the increased risk of unemployment. For a median worker at age 45, the welfare loss amounts to 2.5% of remaining lifetime consumption. At the same time, young workers can partially benefit from depressed housing prices. When housing prices fall, young workers can buy houses at lower prices and benefit from future price appreciation when the economy recovers from the recession. This finding indicates that neglecting the spillover effects of labor market risk on the housing market would ignore the welfare consequences for retired households and generate larger welfare losses for young workers.

Does the age structure of the unemployment shock matter for the magnitude of the spillover effect? While the above analysis assumes a uniform increase in the unemployment risk across all age groups, I study two alternative unemployment shocks to address this question: In the first case, only young workers under the age of 45 experience an unemployment shock, while in the second case, only older workers above the age of 45 are affected. I find that an unemployment shock that disproportionately affects young workers generates a spillover effect on housing prices that is more than three times larger than that caused by a shock to older workers. There are two main reasons for this difference. First, young workers play a significant role in housing demand because many young workers are first-time home buyers and want to move to larger houses over time. Therefore, their economic situation is an important driver of the spillover effect on housing prices. Second, young workers experience larger lifetime income losses when they become unemployed compared to older workers, leading to a larger decline in housing demand.

The overall demographic structure of the economy also plays a key role in the magnitude of the spillover effect of unemployment on housing prices. In light of the fact that the U.S. population is projected to age significantly in the next years, I analyze how an increase in the share of old population shapes the spillover effect. In a simple OLG model, it can be analytically shown that an increase in the share of the old, retired population mitigates the spillover effect if the relative change in the housing demand of working-age households is smaller than the relative change in the housing demand of old, retired households. In the quantitative model, I show that an increase in the old population from 25% to 35% mitigates the spillover effect by 4%. Overall, the demographic composition of an economy significantly affects how unemployment shocks propagate to the housing market.

An important question is whether policy can mitigate the spillover effects of unemployment shocks on housing prices in order stabilize the economy. In the final section of this paper, I show that an unemployment insurance (UI) system can serve as a stabilizer for housing prices during periods of high unemployment. Specifically, I find that raising the UI replacement rate from 60% to 80% in the first two years of a recession mitigates the spillover effect by 12%. However, the UI system is not able to fully counteract the spillover effect. The role of the UI system in mitigating the spillover effect is limited because workers face permanent income losses when they become unemployed and shift their portfolio toward more liquid assets during a recession. As a result, housing demand and prices fall and this effect remains little affected by higher UI generosity. While a higher UI replacement rate provides temporary income insurance and helps smooth consumption, it does not provide insurance against the risk of permanent income losses, so that the size of the spillover effect remains significant.

The remainder of this paper is structured as follows. The following section relates this paper to the existing literature. In Section 3.3, I analyze the empirical relationship between unemployment rates and housing prices. Section 3.4 presents the theoretical model which is followed by its calibration. Section 3.5 analyzes the aggregate dynamics and individual life-cycle consequences of an unemployment shock and explores the welfare consequences across different age groups. In Section 3.6, I examine age-dependent unemployment shocks and the importance of demographic structures for the spillover effect. Section 3.7 examines the role of unemployment insurance as a policy tool to mitigate the spillover effect from unemployment to housing prices. Section 3.8 concludes.

3.2 Related literature

This paper contributes to the extensive literature on the impact of income risk on household portfolio choices by integrating unemployment risk with housing market dynamics. In particular, the mechanism of this paper builds on the approach of

Bayer et al. (2019) who show that higher uncertainty leads to higher precautionary savings of households through the accumulation of liquid assets while reducing illiquid investment and consumption, thereby affecting aggregate activity. My paper replicates this effect in a framework where housing is an illiquid asset and unemployment as a specific type of income risk, providing further support for the portfolio channel of income risk in Bayer et al. (2019).

The literature focusing on income risk and housing demand includes, among others, Attanasio et al. (2012) who use a life-cycle model to analyze how uncertainty about earnings and house prices affect the decision to buy a house. They find that households delay buying a house when they face larger uncertainty, while an increase in income leads to an earlier housing purchase. Paz-Pardo (2024) studies whether changes in homeownership rates across generations can be explained by changes in labor earnings risk. Around half of the decline in homeownership rate is due to the labor market becoming more unequal and volatile. Using age-dependent labor market uncertainty, Chang, Hong, and Karabarbounis (2018) show that unemployment risk and uncertainties regarding job turnover and career path shape the portfolio decision of workers. Changes in labor market risk also affected consumption fluctuations around the Great Recession according to Larkin (2019). The decline in labor market risk before the downturn shifted households' portfolio towards illiquid assets which amplified the consumption drop during the recession.

Regarding the model structure, this paper is closely related to the papers examining housing prices in an OLG economy with a housing market. Kaplan, Mitman, and Violante (2020) analyze the house price fluctuations around the Great Recession in an OLG model with housing and aggregate risks. They find that shifts in beliefs, rather than credit conditions were the key driver of house price fluctuations. Landvoigt, Piazzesi, and Schneider (2015) explore the housing market in San Diego also in a quantitative framework with housing and a credit market. The availability of cheaper credit for poor households is the main reason for the observed changes in housing prices over time. Other papers using an OLG model with housing market are, among others, Corbae and Quintin (2015), Favilukis, Ludvigson, and Van Nieuwerburgh (2017), as well as Chambers, Garriga, and Schlagenhauf (2009).

In terms of methodology, the paper is closely related to Glover et al. (2020) who analyze the welfare consequences of the asset price drop during the Great Recession on different age groups of households. They find that old households experience the largest welfare losses from asset price drops while the younger cohorts gain from buying the assets at a lower price. Whereas the idea of analyzing the differential welfare implications of asset price changes across age groups in this paper is similar to that of Glover et al. (2020), there are two crucial differences: Glover et al. (2020) consider an aggregate shock to all households in the economy and therefore, asset prices are determined by the behavior of all households. In this paper, the unemployment risk only has a direct effect on income of households in the labor market, generating a spillover effect on the housing market: The fluctuations in the house prices are then generated only through the labor market condition of younger households, while old, retired households who are homeowners have little influence on the housing market. Moreover, the model in this paper features financial constraints and heterogeneity within the same age group in contrast to Glover et al. (2020) in which there is no intragenerational heterogeneity. While some households may buy houses at a depressed price and gain from the house price appreciation in future time, most of the young households are subject to credit constraints and cannot afford the required down payment or meet the payment-to-income ratios.

This paper also speaks to the large literature on labor market policies, especially the literature on unemployment insurance systems. Beginning with Baily (1978) and Chetty (2008), many studies have focused on the trade-off between insurance against job loss and incentive for job search in the presence of moral hazard in the optimal unemployment insurance literature. More recent work analyzes how the unemployment insurance system affects the other parts of the labor market: Hagedorn et al. (2013) show that unemployment benefit extensions push up equilibrium wages such that vacancy posting of firms decreases and as a consequence, employment drops in the economy. Moreover, Landais, Michaillat, and Saez (2018) argue that the unemployment insurance changes the labor market tightness and therefore the optimal replacement rate. The results in this paper reveal that the unemployment insurance (UI) system also serves as a stabilizer of the housing market during a recession. In addition to providing insurance for workers, a more generous UI system indirectly benefits retired households by mitigating the spillover effects of unemployment shocks on housing prices. This leads to positive welfare effects for old, retired households, showing that the overall gain from an enhanced UI system during a recession are more significant than usually perceived in the macroeconomic literature.

3.3 Empirical analysis

In this section, I analyze the empirical relationship between unemployment rate and housing prices for the United States. To this end, I exploit the state-level variation in unemployment rate and housing prices from 1978 to 2019. The hypothesis is that fluctuations in unemployment rates influence house prices – specifically, that an increase in the unemployment rate results in a decline in housing prices. To test this hypothesis, I begin by describing the data sets used in the analysis. The empirical analysis then proceeds with an OLS regression. Following this, I apply an IV approach to account for potential endogeneity issues or omitted variable bias in investigating the relationship between unemployment rate and housing prices.

3.3.1 Data

The empirical analysis relies on data from the Current Population Surveys (CPS). The CPS provides official U.S. government statistics on labor market status and is available since 1962. The data set offers comprehensive information on employment and demographics of U.S. households, representing the civilian non-institutional population. The sample focuses on individuals aged 20 and older. From this data set, I extract data on age, labor force status, and the industry composition of employment for each state. Unemployment rates are calculated by using the basic monthly data and applying person-level weights for each month, which are then averaged over the months to obtain annual unemployment rates. As explained later, constructing the shift-share instrument requires data on industry employment shares by state. I use the Annual Social and Economic Supplement (ASEC) samples and apply the March supplement weights to derive the industry shares. Industry classifications follow the 2-digit North American Industry Classification System (NAICS).

Additionally, I use the Home Mortgage Disclosure Act (HMDA) data, published annually by the Consumer Financial Protection Bureau to analyze the relationship between mortgage applications and unemployment rates. The data set provides detailed information on the U.S. mortgage market, based on the reports of financial institutions. Using HMDA data from 2006 to 2020, I analyze the relationship between unemployment rate and the number of mortgage applications. As in the previous step, the data is aggregated at the state level and on an annual basis.

3.3.2 House price and unemployment

As a first step, I examine the relationship between the unemployment rate and housing prices using an OLS regression. Specifically, I regress the natural log of state-level housing price changes on the state-level unemployment rate, while controlling for state-level average household income, national macroeconomic conditions, and state- and year-fixed effects. The results, presented in Table 3.1, indicate a statistically significant relationship between unemployment rates and housing prices, where a 1 percentage point increase in the unemployment rate is associated with a 1.41% decline in housing prices (column 1). This relationship persists even after controlling for additional factors such as housing supply and demographic controls (columns 2 and 3). To ensure the results are not solely driven by the Great Recession, I exclude the period of Great Recession from the data set and repeat the regression analysis. As shown in columns 4-6, the results remain largely unchanged.

A potential concern with the OLS regression analysis is that there might be a problem with endogeneity and omitted variable bias when analyzing the relationship between unemployment rates and housing prices. For example, while unemployment rates influence housing prices, it is also possible that changes in housing prices affect the demand for non-durable consumption and lead to changes in employment in the local labor market. To address this issue, I employ a shift-share

Table 3.1. OLS regression

Δlog(HPI)	All data			Excluding Great Recession		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemp. rate	-1.41^{***}	-1.41^{***}	-1.00^{***}	-1.41^{***}	-1.41^{***}	-1.12^{***}
	(0.17)	(0.17)	(0.16)	(0.19)	(0.19)	(0.17)
Observations	2091	2091	2091	1938	1938	1938
R^2	0.57	0.57	0.63	0.54	0.54	0.60
State- and year-fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
House supply controls	No	Yes	Yes	No	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes

Notes: This table summarizes the results from state-panel regressions of log changes in housing prices on the unemployment rates, state and year fixed effects, and the listed controls. Columns (1)-(3) present the results when the regression analysis considers all data from 1978 to 2019. Columns (4)-(6) exclude the period of Great Recession from the considered data. Standard errors are adjusted for clustering at the state level and reported in parentheses.

instrument, also known as the Bartik instrument, following the methodology developed in Bartik (1991). This instrument helps to identify exogenous movements in the state-level unemployment rates that are not affected by fluctuations in housing prices.

$$Bartik_{s,t} = \sum_{j=1}^{J} e_{sjt-1} \left(\log E_{jt} - \log E_{jt-1} \right) \tag{3.1}$$

Equation (3.1) shows how the instrument is constructed. The idea of the instrument is to measure local labor demand that is not affected by local labor supply. The Bartik instrument is constructed by multiplying the employment share e of industry j in state s, measured one year ahead, by the national growth rate of industry j, excluding state s. This approach uses the variation in industry composition across states to isolate exogenous changes in unemployment rates. To avoid potential bias, the industries classified under "construction" and "real estate and rental and leasing" sectors are excluded from the set of industries J, as these industry categories could affect housing prices through a different channel than through the labor market situation of a state.

$$\Delta \log(HPI_{s,t}) = \beta \hat{u}_{s,t} + \delta X_{s,t-\tau} + \gamma Z_{t-\tau} + \lambda_s + \zeta_t + \alpha_1 + \eta_{s,t} \tag{3.2}$$

The IV regression is shown in Equation 3.2. The dependent variable is the percent change in housing prices in state s between year t and $t-1$. The vector X includes state-level characteristics with up to τ lags, where τ is set to four years.

Table 3.2. First-stage regression

Unemployment rate	All data (1)	Excluding Great Recession (2)
Bartik	−0.38***	−0.37***
	(0.06)	(0.06)
Observations	2091	1938
R^2	0.80	0.79

Notes: This table shows the first-stage regression for the IV analysis. The state-level unemployment rates are regressed on the Bartik shocks including state and year fixed effects, state-level characteristics (average household income, population growth, and ratio of young to old households), and national characteristics (GDP growth, national unemployment rate, stock price index, consumer price index, Federal funds rate, new housing permits, and the supply of new housing). The state-level and national characteristics contain the lags of these variables up to four years. Columns (1) presents the results when the regression analysis considers all data from 1978 to 2019. Columns (2) excludes the period of Great Recession from the considered data. Standard errors are adjusted for clustering at the state level and reported in parentheses.

The state-level characteristics consist of the average household income, population growth, and the ratio of young to old households. The vector Z captures national-level characteristics, such as GDP, national unemployment rate, stock price index, consumer price index, Federal funds rate, new housing permits, and the supply of new housing, along with the lags of these variables, also up to four years. State and year fixed effects are denoted by λ and ζ, respectively.

Table 3.2 presents the results of the first-stage regression. There is a significantly negative relationship between the state-level unemployment rate and the Bartik instrument. This indicates that an exogenous decline in labor supply due to shifts in the national industry composition in a local labor market is correlated with an increase in the state-level unemployment rate.

Table 3.3 reports the main IV estimates of the effect of unemployment rate on housing prices. Similar to the OLS regression, the first three columns show the results using the full data set, while the last three columns exclude the Great Recession period from the analysis.[1] Overall, the IV estimates show a significantly negative relationship between unemployment rate and housing prices. In the final column, which includes all control variables, the IV estimate suggests that a one percentage point increase in the unemployment rate leads to a 1.55% decline in housing prices. Notably, the absolute magnitude of the IV estimates is larger than

1. The underidentification and weak identification statistics are reported following Kleibergen and Paap (2006). The strength of the instrument is not a concern as the p values for the underidentification are less than 0.01 and the Kleibergen-Paap weak identification statistics are larger than the critical value of 16.38 for 10% maximal IV size.

Table 3.3. IV regression

Δlog(HPI)	All data			Excluding Great Recession		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemp. rate	-1.98^{***}	-1.98^{***}	-1.50^{***}	-1.93^{***}	-1.93^{***}	-1.55^{***}
	(0.47)	(0.47)	(0.55)	(0.47)	(0.47)	(0.57)
Observations	2091	2091	2091	1938	1938	1938
R^2	0.43	0.43	0.52	0.38	0.38	0.46
State and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
House supply controls	No	Yes	Yes	No	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes

Notes: This table summarizes the state-panel IV regressions of log changes in housing prices on the instrumented unemployment rates. The controls include state and year fixed effects, state-level characteristics (average household income, population growth, and ratio of young to old households), and national characteristics (GDP growth, national unemployment rate, stock price index, consumer price index, Federal funds rate, new housing permits, and the supply of new housing). The state-level and national characteristics contain the lags of these variables up to four years. Columns (1)-(3) presents the results when the regression analysis considers all data from 1978 to 2019. Columns (4)-(6) exclude the period of Great Recession from the considered data. Standard errors are adjusted for clustering at the state level and reported in parentheses.

those from the OLS regression. This indicates that in the OLS regression, the effects of unemployment rate on housing prices are biased downwards.

3.3.3 Mortgage applications

The results of the IV regression reveal a significantly negative relationship between the unemployment rate and housing prices. One possible explanation for this finding is that higher unemployment rate may reduce housing demand, leading to a decline in housing prices. To further explore this, I analyze the link between housing demand and unemployment rate using data from the Home Mortgage Disclosure Act (HMDA). The HMDA data set, which contains mortgage applications submitted by households to financial institutions, serves as a proxy for housing demand of households.

Figure 3.1 presents a summary of the results. In Figure 3.1a, there is a strong negative correlation between the number of mortgage applications and the unemployment rate at the state level. Additionally, Figure 3.1b displays the relationship between the share of denied mortgage applications and the unemployment rate at the state level, showing a positive correlation. This suggests that financial institutions are more likely to deny mortgage applications during times of high unemployment. These findings point to two key observations during times of high unemployment: first, fewer households apply for mortgage loans, and second, a larger share

Figure 3.1. Mortgage applications and unemployment rate

(a) Number of loans **(b)** Share of denied applications

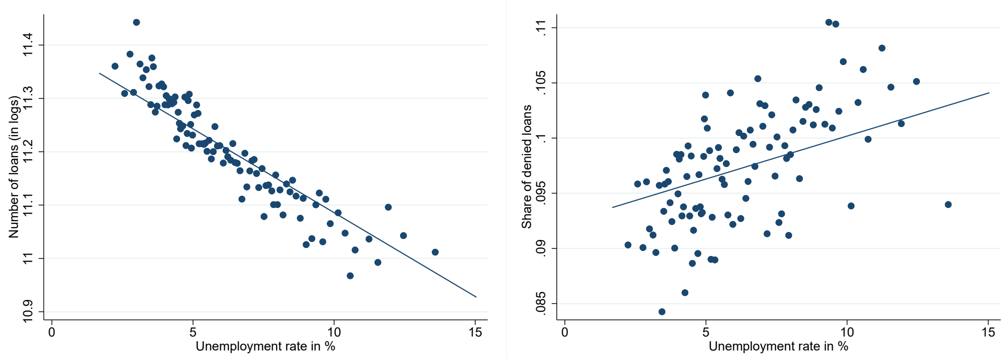

Notes: Panel (a) shows the number of mortgage applications against unemployment rate by each state. Panel (b) shows the share of mortgage applications that have been rejected. The results are after controlling for GDP, nationwide unemployment rate, and state- and year-fixed effects.

of those applications is denied by the financial institutions. Together, both observations indicate that housing demand declines when the economy is experiencing a rise in the unemployment rate.

In summary, this section analyzes the empirical relationship between unemployment rates and housing prices in the United States from 1978 to 2019. The IV regression analysis reveals that higher unemployment rates result in a significant decline in housing prices. An investigation of the HMDA data indicates that high unemployment reduces housing demand, as indicated by fewer mortgage applications and a higher share of denied loans.

3.4 Model

This section begins by outlining the household problem, followed by the equilibrium condition and the calibration of the model to the U.S. data. Finally, I present the properties of aggregate shocks in the model.

3.4.1 Households

The economy is populated by overlapping generations with finitely lived households. Time is discrete and the life cycle of a household is divided into a working $[1, J^{ret}-1]$ and a retirement phase $[J^{ret}, J]$. During the working phase, households provide an inelastic labor supply of one efficiency unit, while facing uninsurable idiosyncratic labor income risk. Households are risk averse and maximize expected lifetime utility by allocating resources to non-durable consumption, housing services, and savings in liquid financial assets. The age of households is denoted by j. The functional form of the utility function is given by

$$u(c,s) = \left([(1-\phi)c^{1-\gamma} + \phi s^{1-\gamma}]^{(1-\vartheta)/(1-\gamma)} - 1\right)/(1-\vartheta)$$

as in Kaplan, Mitman, and Violante (2020), where c denotes non-durable consumption, s denotes housing services, and b is the amount of liquid assets. Households leave bequests which yield a utility of

$$v(b) = \nu \frac{(b+\underline{b})^{1-\vartheta} - 1}{1-\vartheta}.$$

The utility function for the bequest motive follows De Nardi (2004). During the working phase, households are either employed or unemployed. Employed households earn labor income y_j^e given by

$$\log y_j^e = \log \Theta + \chi_j + \varepsilon_j,$$

where Θ represents the aggregate productivity of labor, and χ_j captures the deterministic, age-dependent income profile. The idiosyncratic income component, ε_j, evolves according to a first-order Markov process. Households can experience unemployment for a fraction d^u of a period. While unemployed, they receive a transfer $b^u \cdot y_j^e$, where b^u is the replacement rate of the unemployment insurance system. For the remaining fraction $1-d^u$ of the period, households earn their labor income y_j^e. Consequently, the total income of an unemployed household is given by

$$y_j^u = y_j^e\left[(1-d^u) + b^u d^u\right].$$

Upon becoming unemployed, workers face the risk of persistence earnings losses with probability $q_{j,t}^u$. Specifically, workers experience a decline in their idiosyncratic component ε_j of their labor income process. In this case, there is a persistent reduction in future earnings, leading to long-term effects of unemployment on labor income of workers. The recursive formulation of the household problem, presented below, provides a more detailed account of how these earnings losses are incorporated during unemployment.

3.4.2 Recursive formulation of the decision problem

Households who own a house face the decision between staying in their current house and repaying the mortgage or selling their house and buying a new house. Thus, the value function at the beginning of a period is given by

$$V_{j,t}(b,y,l,m,h) = \max\left\{V_{j,t}^h(b,y,l,m,h),\ V_{j,t}^s(b,y,l,m,h)\right\}. \tag{3.3}$$

$V_{j,t}^h$ denotes the value function of a homeowner repaying the mortgage, and $V_{j,t}^s$ is the value function of a new house buyer. The state variable l denotes the labor market status of the household, with $l = e$ if a household is employed and $l = u$ if unemployed. Households that decide to remain in their current house and repaying the mortgage maximize their utility by choosing consumption c and liquid savings b', given the expectation about future employment and income state and subject to a zero borrowing constraint. The household's optimization problem is given by

$$\begin{aligned} V^h_{j,t}(b,y,l,m,h) &= \max_{c,b'} u(c,s) + \beta \mathbb{E}_{y',l'}\left[V_{j+1,t+1}(b',y',l',m',h')\right] \\ s.t. \quad c+b'/(1+r_f) + \omega_j(m) &\le b + y - \tau(y,h) \\ b' \ge 0, \quad s = h, \quad h' &= h, \quad m' = (1+r_m)m - \omega_j(m) \end{aligned} \tag{3.4}$$

where β is the time discount factor, τ_h is the property tax rate, d_h is the maintenance cost, p_t is the housing price, and $\omega_j(m)$ is the mortgage repayment as a function of the remaining mortgage balance m. The next period's mortgage balance m' is the difference between today's mortgage plus mortgage interest rates r_m and today's mortgage repayment. The mortgage repayment $\omega_j(m)$ follows a constant amortization formula:

$$\omega_j(m) = m \cdot \frac{r_m(1+r_m)^{J-j}}{(1+r_m)^{J-j}-1}$$

The term $\mathbb{E}_{y',l'}\left[V_{j+1}(b',y',l',m',h')\right]$ contains the expectation about the employment and income states in the next period where

$$\begin{aligned} \mathbb{E}_{y',l'}\left[V_{j+1,t+1}(b',y',l',m',h')\right] &= (1-\lambda_{j+1,t+1}) \cdot \mathbb{E}_{y'}\left[V_{j+1,t+1}(b',y',e,m',h')\right] \\ &+ \lambda_{j+1,t+1}\Big\{(1-q^u_{j+1,t+1}) \cdot \mathbb{E}_{y'}\left[V_{j+1,t+1}(b',y',u,m',h')\right] \\ &+ q^u_{j+1,t+1} \cdot \mathbb{E}_{y'}\left[V_{j+1,t+1}(b',y^-,u,m',h')\right]\Big\} \end{aligned}$$

Here, $\lambda_{j,t}$ represents the age-dependent unemployment probability, while $q^u_{j,t}$ is the probability of experiencing persistent earnings losses. When households become unemployed, they may suffer a persistent income loss with a reduction in the income state to y^- with probability $q^u_{j,t}$. For households who decide to sell their house and buy a new one, the budget constraint reflects the sale of the house:

$$b^s = b + (1-\kappa)p_t h - (1+r_m)m \tag{3.5}$$

where κ denotes the transaction cost. Households who purchase a new house obtain a new mortgage m' per housing unit, which is subject to the loan-to-value (LTV) limit ρ_m and payment-to-income ratio ρ_y. The value function of a new house buyer is:

$$\begin{aligned} V^s_{j,t}(b,y,l,m,h) &= \max_{c,h',b'} u(c,s) + \beta \cdot \mathbb{E}_{y',l'}\left[V_{j+1,t+1}(b',y',l',m',h')\right] \\ s.t. \quad c + p_t \cdot h' + b'/(1+r_f) &\le b^s + y - \tau(y,h) + m' \\ b' \ge 0, \quad s = h', \quad m' &\le \rho_m \cdot p_t \cdot h', \quad \omega_j(m') \le \rho_y \cdot y \end{aligned} \tag{3.6}$$

In retirement, workers have the same decision problem as in the working phase, but face an age-dependent death probability of d_j and leave a bequest of b^j upon death. The problem of households staying in the same house is then

$$V^h_{j,t}(b,y,m,h) = \max_{c,b'} u(c,s) + \beta\left[(1-d_{j+1})V_{j+1,t+1}(b',y',m',h') + d_{j+1}\cdot \upsilon(b^{j+1})\right] \tag{3.7}$$

$$\begin{aligned} s.t.\quad & c + b'/(1+r_f) + \omega_j(m) \leq b + y - \tau(y,h) \\ & b' \geq 0, \quad s = h, \quad h' = h, \quad m' = (1+r_m)m - \omega_j(m) \\ & b^{j+1} = b' + (1-\kappa)\cdot p_t h - m' \end{aligned}$$

The combination of a death probability and a bequest motive is similar to the model of Kopczuk and Lupton (2007).

3.4.3 Equilibrium

The individual state variables consist of liquid assets b, income state y, employment l, mortgage balance m, and housing stock h. The individual state vector is denoted by $x := (b,y,l,m,h) \in \mathbb{X}$. A competitive equilibrium is a collection of value functions $\left\{V_{j,t}(x), V^h_{j,t}(x), V^s_{j,t}(x)\right\}$, household decision rules $\left\{g_{j,t}(x), b_{j+1,t}(x), c_{j,t}(x), h_{j+1,t}(x)\right\}$, equilibrium housing prices p_t as functions of time t such that:

(1) Given the housing prices p_t, the decision rules $\left\{g_{j,t}(x), b_{j+1,t}(x), c_{j,t}(x), h_{j+1,t}(x)\right\}$ solve the households' decision problem by solving problems (3.3)-(3.7), with the value functions $\left\{V_{j,t}(x), V^h_{j,t}(x), V^s_{j,t}(x)\right\}$.

(2) The housing market clears at the equilibrium price p_t and the housing stock in the economy satisfies

$$\sum_{j=1}^{J}\int_{\mathbb{X}} h_{j,t}(x)d\mu_{j,t} = \overline{H} \quad \forall t$$

where $\mu_{j,t}$ is the cumulative distribution of individual states in the population and the total housing stock $\overline{H}$ is fixed.

A steady state of the economy is a competitive equilibrium where the distribution of agents is stationary.

3.4.4 Calibration

This section explains the calibration of the model and discusses its empirical fit. The steady-state equilibrium is calibrated to the U.S. economy using data from the Survey of Consumer Finances (SCF) in year 2019.

A period in the model corresponds to one year. Households enter the model and are in the labor force at age 25 and retire at age 67. All households are assumed to

Table 3.4. Model parameters

Parameter	Value	Description
J^w	42	Working life
J^r	24	Retirement
β	0.98	Discount factor
$1/\gamma$	1.25	Elasticity of substitution
θ	2.0	Risk aversion
υ	100	Bequest motive
$\underline{b}$	5	Bequest as luxury
$\mathscr{H}$	{.01, 1.5, 1.92, 2.46, 3.15, 4.03, 5.15}	House sizes
κ	.07	Transaction cost
τ	.75, .151	Income taxation function
χ_j	Kaplan and Violante 2014	Deterministic age profile
ρ_ε	.97	Autocorrelation of earnings shocks
σ_ε	.20	Standard deviation of earnings shocks
r_f	.03	Risk-free interest rate
r_m	.055	Mortgage interest rate

Notes: This table shows the parameter values in the model. One model period corresponds to one year.

die at age 91. The parameter values used in the calibration are summarized in Table 3.4. The estimated annual discount factor β is 0.98, which aligns with values commonly used in the macroeconomic literature. The elasticity of substitution between housing and non-durable consumption is set at 1.25 following Piazzesi, Schneider, and Tuzel (2007). To match an elasticity of intertemporal substitution of 0.5, ϑ is calibrated to 2. The parameters governing the bequest motive are $\upsilon = 100$ and $\underline{b} = 5$, which are consistent with the estimates provided by Kaplan, Mitman, and Violante (2020). In setting up the house size grid $\mathscr{H}$, I follow the discretization approach Kaplan, Mitman, and Violante (2020).[2]

The income process is calibrated using the deterministic age profile from Kaplan and Violante (2014). The idiosyncratic earnings shock ε_j follows an AR(1) process with a persistence of 0.97 and standard deviation of innovations of 0.2. Consistent with Heathcote, Perri, and Violante (2010), this calibration of earnings shocks produces an increase in the variance of log labor earnings of 2.5. The risk-free interest rate r_f is set at 0.03 and the mortgage interest rate r_m at 0.055.

2. See Appendix 3.A.5 for further details.

In order to calibrate the pension system in the model to match the U.S. Social Security system, I follow the 2019 U.S. Social Security legislation. The Social Security cap is set at $132,900, with the first and second bendpoints at $926 and $5,583, respectively. The retirement benefits formula is given by

$$\Omega(\bar{y}) = \begin{cases} 0.9\bar{y} & \text{if } \bar{y} < bp_1, \\ 0.9bp_1 + 0.32(\bar{y} - bp_1) & \text{if } bp_1 \leq \bar{y} < bp_2, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(\bar{y} - bp_2) & \text{if } bp_2 \leq \bar{y} < cap, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(cap - bp_2) & \text{if } \bar{y} > cap \end{cases} \tag{3.8}$$

where $\bar{y}$ denotes the average lifetime earnings, $\Omega(\bar{y})$ denotes the assigned benefit level, and bp_1 and bp_2 denote the two bendpoints.[3]

Figure 3.2 presents the life-cycle profiles for the unemployment rate, wealth, housing wealth, and leverage, comparing the model outcomes to their empirical counterparts. Looking at the unemployment rate by age in Figure 3.2a, the model closely matches the empirical profile. Both in the model and the data, there is a strong decline in the unemployment rate until age 35, after which it stabilizes at approximately 4%. Figure 3.2b demonstrates that the model effectively captures the steep increase in wealth throughout the life cycle as observed in the data. The bequest motive in the model is crucial for matching the empirical wealth profile. In the absence of bequest motive, households would decumulate their wealth as they approach the end of the life cycle, resulting in a steep fall in wealth. The model also matches the profiles of housing wealth and leverage very well. The profile of housing wealth, shown in Figure 3.2c, implies a strong increase in housing wealth during the early stages of the life cycle, peaking around age 40. Afterwards, the profiles remain almost constant. Concerning household leverage, Figure 3.2d shows a gradual decline in leverage over the life cycle both in the model and in the data. Young workers exhibit high leverage as they climb the housing ladder and take on mortgage loans. As they get older, households repay their mortgages, which drives down the leverage level.

3. The retirement benefit formula is taken from http://www.ssa.gov/OACT/COLA/piaformula.html.

Figure 3.2. Life-cycle profiles

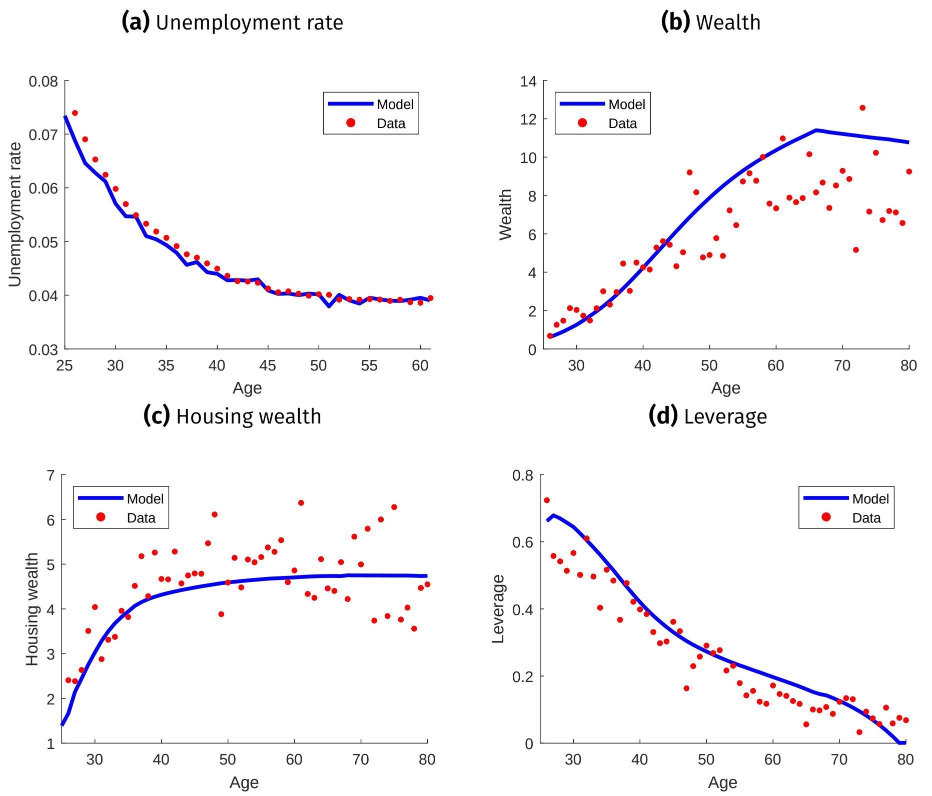

Notes: This figure shows the life-cycle profiles of unemployment rate, wealth, housing wealth, and leverage. The dots show the empirical profiles, while the solid lines show the corresponding model profiles. The empirical unemployment rate profile is computed using data from the Current Population Survey. Other empirical profiles are computed using data from the Survey of Consumer Finances.

3.4.5 Unemployment shock

Table 3.5. Properties of the unemployment shock

Description	Value	
	Steady state	Recession
Average unemployment rate	4%	8%
Unemployment duration	0.25	0.5
Probability of earnings loss	0.2	0.8

Notes: This table compares the model parameters and properties when the economy is in steady state and when an unemployment shock occurs.

The economy is initially in a steady state. The unemployment shock is introduced as a one-time, unanticipated shock at $t = 0$. Households have perfect foresight and know that the economy will revert to the initial steady state after T pe-

Table 3.6. Income losses upon unemployment

State of economy	Model	Davis et al. (2011)
Steady state	−25%	−25%
Recession	−44%	−39%

Notes: This table compares the earnings losses at the time of displacement for workers at age 40 in the model (column 2) and the results in Davis et al. (2011) (column 3).

riods. When the unemployment shock occurs, the age-dependent separation probability and the unemployment duration increase. Moreover, the risk of persistent income losses also increases. Table 3.5 outlines the properties of the unemployment shock for the baseline economy. When the unemployment shock occurs, the average unemployment rate increases from 4% to 8%, while the average unemployment duration increases from a quarter to 6 months.

Table 3.6 compares the average income loss following unemployment relative to pre-displacement income in the model to the findings of Davis et al. (2011). In the steady state, the model produces a 25% decline in earnings in the period of displacement, while during a recession, the income loss is much larger and amounts to 45%. Compared to Davis et al. (2011), the income loss in a recession is slightly larger in the model at hand, but the overall calibration aligns well with the empirical income losses.

3.5 Spillover effect of unemployment on the housing market

This section explores the spillover effect of an unemployment shock on the housing market using the calibrated model. In particular, I show that the impact of the spillover effect is large, both on the aggregate level and for individual households. The calibrated model matches the housing price drop observed in the empirical data documented in Section 3.3. For the aggregate effects, I analyze the dynamics of housing price and average income over time following an unemployment shock to the economy. At the individual level, I show how the unemployment shock affects housing demand and the differential welfare consequences of the spillover effect on households of different age groups.

3.5.1 Aggregate dynamics

Figure 3.3a shows the average unemployment rate by age for the steady state of the economy and in the first period of a recession. In the baseline economy, the unemployment shock leads to a doubling of the unemployment rate across all age groups, implying a uniform increase in unemployment for workers of all ages. As the unemployment rate is on average higher for younger workers, the absolute increase

Figure 3.3. Unemployment shock and housing price

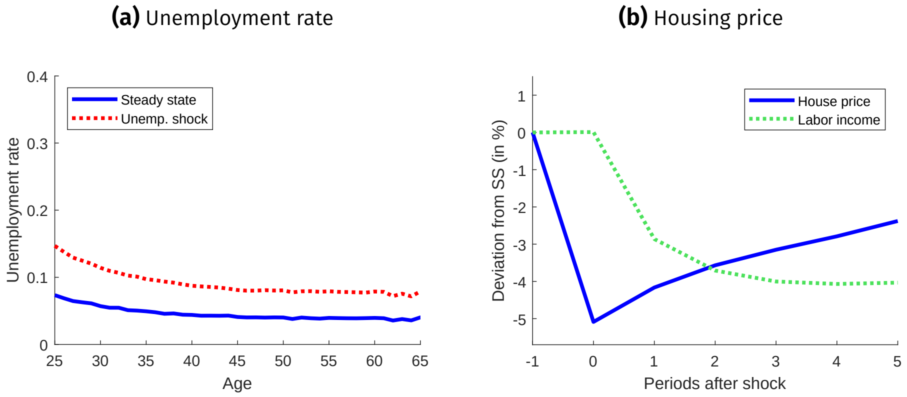

Notes: This figure shows the unemployment shock and the house price dynamics after the unemployment shock hits the economy. Panel 3.3a compares the average unemployment rate by age when the economy is in a normal state (solid line) with the average unemployment rate in a recession (dotted line). Panel 3.3b shows the dynamics of housing price (solid line) and and average labor income (dotted line) when the unemployment shock hits the economy at time zero.

in the unemployment rate during a recession is more pronounced for young workers. To analyze the impact of the unemployment shock on the aggregate economy, I start from the steady state of the model. At time zero, the unemployment shock hits the economy. In the following periods, the economy slowly returns to its steady state, with the unemployment rate and all other parameters reverting to their steady-state levels. The results are displayed in Figure 3.3b. In response to the aggregate shock, both housing prices and average income drop significantly. Specifically, the housing price drops by 5%, while the average income falls by 2.9% in the first period of the shock. The recovery of the average income is slow due to large and persistent income losses of workers who were affected by the shock.

In the first period after the unemployment shock, there is a strong initial recovery of the housing price which can be explained by two reasons. First, the income uncertainty declines over time such that households increase their housing demand. Second, households shift their portfolio during the recession. When the unemployment shock hits the economy, workers reduce their housing demand and shift their savings into liquid assets to insure themselves against the heightened income risk (see Appendix 3.A.7). Once the economy is recovering from the unemployment shock, workers hold more liquid assets for precautionary savings motive than usual, resulting in an overshooting in housing demand. As a consequence, the housing price quickly recovers one period after the shock.

The above results show that the calibrated model matches the observed drop in housing price following unemployment shocks, consistent with empirical data. Two key factors account for the drop in the housing price: first, the average income level declines due to higher unemployment rate and large persistent income losses of workers. As a consequence, workers reduce their demand for housing, and in turn,

Figure 3.4. Counterfactual case with fixed income

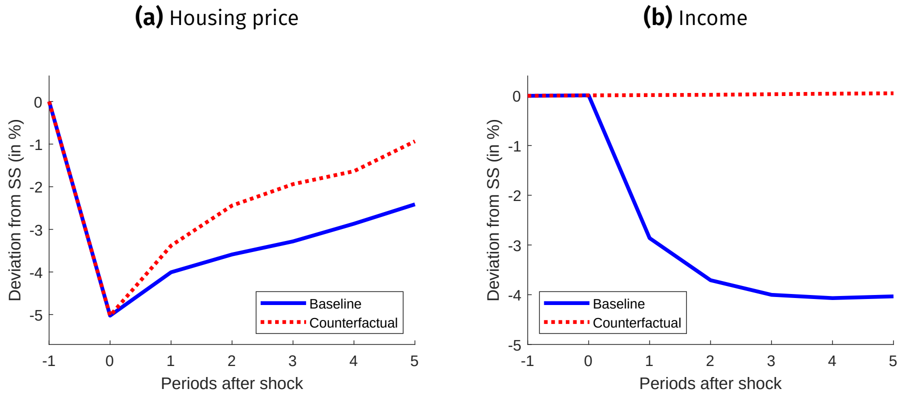

Notes: This figure shows a counterfactual experiment where the economy is hit by the same unemployment shock as in the baseline economy, but the average income level is held constant at the steady-state level. By keeping the spillover effect stemming from the income channel to zero, the counterfactual experiment isolates the risk effect of higher probability of unemployment and permanent income losses when the unemployment shock hits the economy. Panel 3.4a shows the dynamics of housing price and Panel 3.4b the dynamics of income after the unemployment shock occurs at time zero. In both panels, the solid lines refer to the baseline economy and the dotted lines to the counterfactual model.

the general equilibrium effect pushes down the price on the housing market. Second, the persistence of the unemployment shock also plays an important role. Due to the persistence of the shock, workers face a heightened risk of unemployment and income losses in future periods. The increased income risk further suppresses housing demand, leading to an additional decline in housing prices.

One key question is how much the decline in income level and the increase in income uncertainty each contribute to the observed drop in housing prices following an unemployment shock. In the following, I conduct a counterfactual experiment where the economy experiences the same unemployment shock as before, but with the average income level held constant at the steady-state level. In this case, workers are still subject to a higher risk of unemployment and higher probability of persistent income losses, but the actual realized unemployment rate and the average income remain unchanged from the steady state.

The results from the counterfactual experiment are displayed in Figure 3.4. By comparing the housing price drop in the baseline model with the results from the counterfactual economy after the unemployment shock in Figure 3.4a, it becomes evident that the main cause of the housing price drop in the initial period of the shock is the increase in income uncertainty. In the counterfactual experiment, the unemployment shock leads to the same housing price drop of 5% on impact as in the baseline economy, but in the later periods the housing price recovers more quickly. After 5 periods, the housing price in the counterfactual economy is 1% below the steady state compared to 2.4% in the baseline economy. This is because the average income level remains unchanged from its steady state, and as a consequence, the housing price quickly returns to its steady state. Quantitatively, higher income risk explains approximately 60% of the total housing price deviation from the steady state, while actual income losses account for the remaining 40%.[4] The main reason for a lower housing demand is therefore the heightened risk of unemployment. This finding complements the study by Bayer et al. (2019) who show that higher uncertainty leads to a drop in illiquid investment of households, while increasing precautionary savings.

3.5.2 Housing demand and the consequences of the spillover effect

The above results show that unemployment shocks generate large spillover effects on the housing market. Next, I demonstrate that unemployment shocks have a significant impact on housing demand of younger workers compared to older workers.[5] To explore this, I construct a set of counterfactual experiments. In the first experiment, workers never experience unemployment over their life cycle. In the second case, workers become unemployed either at age 26 or 45 during a period without an unemployment shock. The final counterfactual experiment assumes that workers become unemployed during a recession. In the last two counterfactuals, workers become unemployed only at age 26 or age 45 and experience no further unemployment spells thereafter.

Figure 3.5 shows the life-cycle profiles of house size of workers who become unemployed at age 26 and at age 45, respectively. For younger workers who experience unemployment at age 26, the average house size drops by 20% two years after the unemployment shock. In contrast, older workers who become unemployed at age 45 see a drop in house size by only 5%. Even though the gap between young workers unemployed at age 26 and the control group who do not experience unemployment decreases with time, they are not able to fully catch up. Even after 15 years when the unemployment shock occurs, the average house size of workers who become unemployed at age 26 is still almost 15% lower than those who do not experience unemployment. For older workers, the gap in house size does not decrease over time and remains persistent, resulting in a 10% smaller housing size compared to the control group who never become unemployed.

The unemployment shock has more severe consequences for younger households due to several reasons. First, due to losses in permanent income upon unemployment, younger workers who have a longer remaining life span experience larger drops in lifetime earnings compared to the older workers. Another reason is that among young workers, the share of marginal buyers is higher. Younger workers

4. The total house price deviation is defined as the cumulative house price deviation from the steady state across time in both the baseline and the counterfactual scenario.

5. In Appendix 3.A.8, I show that an unemployment shock at the beginning of the life-cycle has substantial and long-lasting consequences on income, housing wealth, and asset accumulation.

Figure 3.5. Housing demand and unemployment shocks at different ages

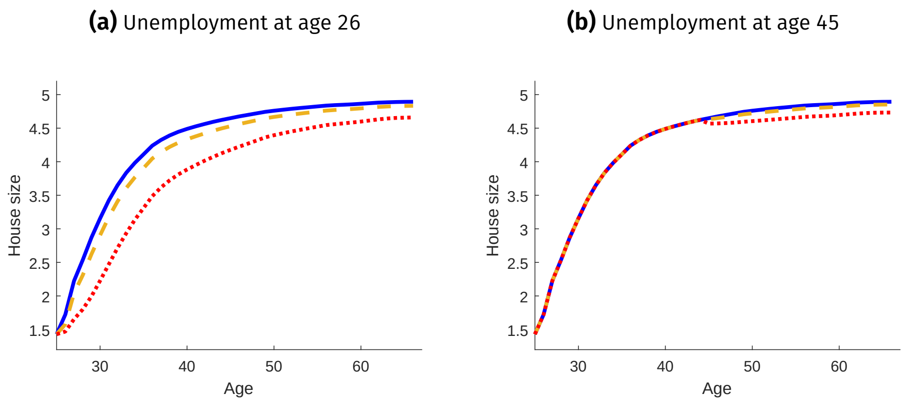

Notes: This figure shows the life cycle profiles of house size in a counterfactual experiment. Panel 3.5a displays the case where workers become unemployed at age 26. Panel 3.5b shows the counterfactual experiment where workers become unemployed at age 45. The solid lines display the life-cycle profiles of workers who never experience unemployment. The dashed and dotted lines show the life-cycle profiles of workers who become unemployed at age 26 or age 45 when the economy is at its steady state and in a recession, respectively.

on average have smaller houses and climb the housing ladder as their incomes grow over time. Hence, young workers have a higher demand for larger houses compared to older workers, who typically own larger houses. Moreover, the demand for housing of older workers is less affected by the unemployment shock because they have on average higher incomes, making it easier to meet mortgage requirements such as the payment-to-income and loan-to-value ratios when buying a new house. In contrast, younger workers subject to unemployment and income drops may struggle to meet credit conditions, so that it is more difficult for younger workers to obtain mortgage loans. Finally, younger workers have little liquid wealth to smooth consumption when they are hit by the unemployment shock. Older workers, on the contrary, have accumulated more liquid wealth for precautionary savings motive and as life-cycle savings, and thus have a better ability to smooth consumption without adjusting their optimal housing size. Overall, the findings imply that unemployment shocks to young workers entail larger spillover effects on the housing market as the labor market condition of young workers is key for housing demand in the economy.

Although housing demand declines when workers face increased income uncertainty, the resulting spillover effect partly stabilizes housing demand. As house prices fall, some workers increase their housing demand because they can buy housing at depressed prices and benefit from the housing price gain in future periods. In the following, I explore the consequences of the spillover effect on working-age households. To quantify the impact of the spillover effect on housing demand of young workers, I construct a counterfactual economy where I shut down the spillover effect. That is, the housing prices are fixed at the steady-state levels and

Table 3.7. Change in housing size relative to a counterfactual without spillover effect

Model	Deviation from steady-state housing size (in %)	
	Unemployment at age 26	Unemployment at age 45
Baseline	−14.79	−1.71
No spillover effect	−16.69	−3.22

Notes: This table compares the steady-state deviation in housing size in the period when an unemployment shock occurs in the baseline economy and in a counterfactual economy where the spillover effect on the housing market is shut down. Column 1 specifies the model. Column 2 shows the results for workers at age 26 who become unemployed. Column 3 considers the case where workers at age 45 become unemployed.

do not move when the economy experiences an unemployment shock. The results are presented in Table 3.7.

In the baseline economy with the spillover effect, the average house size of unemployed workers is 14.79% lower relative to employed workers when the unemployment shock occurs. In the counterfactual economy without the spillover effect, where housing prices remain fixed at steady-state levels, housing becomes more expensive. As a result, housing demand declines further compared to the baseline economy with the spillover effect, with the average house size of unemployed workers at age 26 being 2 percentage points lower than in the baseline model. For workers who are unemployed at age 45, their average house size decreases by 1.71% in the baseline and by 3.22% in the counterfactual economy. These findings suggest that the spillover effect of unemployment shocks on the housing price has strong effects on housing demand of working-age households. The general equilibrium feedback effect mitigates the decline in housing prices, partly stabilizing housing demand during periods of high unemployment.

What are the consequences of the spillover effect of unemployment shocks on housing prices for old, retired households? In the final part of this section, I analyze how the spillover effect of unemployment shocks on the housing market affects old households who are already in retirement and are not directly impacted by the unemployment shock. In particular, I ask how consumption and savings behavior of retired households changes when the unemployment shock hits the economy.

Figure 3.6 shows the impact of the unemployment shock on non-durable consumption and liquid assets for older, retired households. In Figure 3.6a, there is a sharp drop in non-durable consumption at time zero when the unemployment shock occurs. This decline is driven by the drop in housing wealth, which induces retired households to adjust their consumption-saving behavior. As shown in Figure 3.6b, retired households save more into liquid assets to compensate for the loss in housing wealth in response to the unemployment shock. In the later periods, as the unemployment shock diminishes, the consumption level of retired households gradually recovers over time and eventually overshoots the initial steady-state level. This

Figure 3.6. Spillover effect of unemployment shocks on retired households

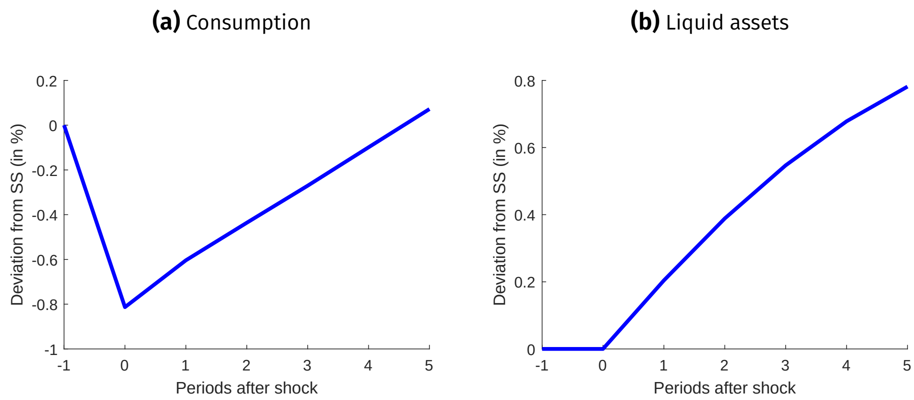

Notes: This figure shows the changes in non-durable consumption (Panel 3.6a) of retired households above age 78 and liquid assets (Panel 3.6b) of retired households above age 67 in response to the unemployment shock which occurs at time zero.

overshooting occurs because retired households accumulate more liquid assets than in the steady state during the recession. As the economy reverts back to its steady state, retired households increase their consumption beyond the steady-state level and drive down their liquid savings. These findings show that even though old, retired households are not directly exposed to unemployment risk, they still experience negative welfare effects when the unemployment shock hits the economy. To further investigate the consequences of the spillover effect on retired households, I analyze the welfare consequences in the next subsection.

3.5.3 Welfare analysis

The previous sections have shown that unemployment shocks generate large spillover effects on housing prices. In response to the shock, housing demand and consumption-saving behavior of households change through its direct effect on household income and also through the spillover effect on the housing market. In this section, I explore the welfare consequences of unemployment shocks across different age groups of households. I also analyze the heterogeneous welfare effects within households of the same age group, asking how employment status and house sizes lead to varying welfare outcomes.

The welfare results are summarized in Figure 3.7. For households in the labor market, Figure 3.7a shows the welfare consequences of a recession, measured in terms of consumption-equivalent variation (CEV), across different ages for workers with median and high income.[6] The welfare losses are large for workers at the be-

6. The consumption-equivalent variation (CEV) is defined as the percentage of consumption a household would be willing to give up in the baseline economy in order to have the same level of lifetime utility as in the economy where the unemployment shock occurs.

Figure 3.7. Welfare consequences of unemployment shocks

(a) Workers **(b)** Retirees

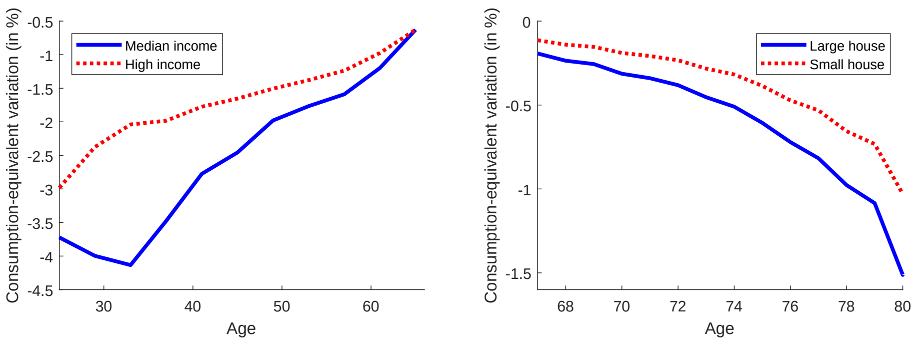

Notes: This figure shows the welfare consequences of a recession for each age in consumption-equivalent variation. Panel 3.7a displays the welfare results for workers who have median income (solid line) and workers with higher income (dotted line). Panel 3.7b shows the welfare results for retirees who own relatively small houses (solid line) and retirees owning a larger house (dotted line).

ginning of the life cycle. At age 25, workers are subject to welfare losses equivalent to 3.7% of remaining lifetime consumption. For older workers, these welfare effects are smaller, declining to 2.5% for workers at age 45 and declining further as workers get close to the retirement age. Households at the beginning of the life cycle bear the largest welfare costs due to following reasons. First, younger workers see larger losses in their lifetime income compared to older workers as younger workers have more years to remain in the labor market and hence, face a larger impact from persistent losses in labor earnings. Second, workers at the beginning of the life cycle gradually increase their housing size in order to save into illiquid asset, and at the same time, draw utility from housing consumption. In a recession, mortgage loans become less accessible to unemployed workers, and as a consequence, this restricts their ability to move up the housing ladder. Workers with high incomes also experience significant welfare losses from the unemployment shock, though the impact is smaller compared to median workers. This is because high-income workers are less constrained in their housing consumption decisions. The difference in welfare effects between high- and median-income workers is particularly large early in life, with a gap of more than 1.5 percentage point at age 30. This disparity gradually decreases in age.

The spillover effect of unemployment shock on housing prices generates large welfare consequences for retired households, as depicted in Figure 3.7b. Households with larger houses at age 70 are subject to welfare losses equivalent to 0.31% of their remaining lifetime consumption. Closer to the end of the life cycle, the welfare losses increase significantly, reaching 1.51% by age 80. Households who own smaller houses experience dampened welfare effects because they see smaller declines in housing wealth. Moreover, retirees in the early stages of retirement face rel-

Table 3.8. Welfare decomposition

Age	Welfare effects (in %)	
	Baseline	No spillover effect
25	−3.72	−5.26
30	−3.99	−4.55
35	−3.65	−3.63
40	−2.91	−2.90
45	−2.46	−2.46
50	−1.95	−1.95
55	−1.66	−1.66
60	−1.31	−1.29
65	−0.63	−0.54
70	−0.19	0.00
75	−0.39	0.00
80	−1.03	0.00
85	−1.06	0.00
90	−7.66	0.00

Notes: The welfare effects are measured in consumption-equivalent variation on remaining lifetime consumption. In all models, the welfare effects are computed for a median household in the economy in terms of assets, employment status, income, housing size, and remaining mortgage balance. The baseline model refers to the results in Figure 3.7. The model with fixed housing price is a counterfactual case where the housing prices are always at the steady-state price of the baseline model.

atively modest welfare losses, as they expect that the economy will recover and the housing prices return to steady-state levels. Overall, the spillover effect of unemployment shocks through the housing market transmits approximately one-third of the welfare losses of workers to retired households, measured in terms of consumption-equivalent variation. It is important to note that retired households are affected by the unemployment shocks only through the spillover effects on housing prices. Retired households are not directly exposed to unemployment risk as they are not taking part in the labor market and their main source of income consists of retirement benefits. As a large share of household wealth consists of housing wealth, the decline in housing prices caused by the spillover effect of unemployment shocks implies that housing wealth declines during a recession. This affects the amount of wealth left as a bequest motive, leading to negative welfare effects on retirees.

How important is the spillover effect of unemployment shocks on housing prices for the welfare consequences reported in Figure 3.7? The above findings suggest that young workers experience significant welfare losses because they lose the opportunity to climb the housing ladder and the decline in income, while old retired households are subject to negative welfare effects due to the spillover effect on housing prices and drops in housing wealth. In the following, I construct a counterfactual model to quantify the importance of the spillover effect for the welfare results.

The counterfactual experiment is a "fixed housing price" scenario, in which the housing price remains constant at its steady-state level. Households have rational expectations, and they know that the housing price is constant and does not change over time. The results are summarized in column 3 of Table 3.8. In this counterfactual model, retired households experience no welfare losses because they are not subject to the heightened earnings risk during a recession. This finding indicates that the welfare losses of retired households in the baseline economy, summarized in column 2, are entirely due to the depressed housing prices during a recession.

In the absence of the general equilibrium effect of housing prices, young households experience even larger welfare losses during a recession. The reason is that when housing prices decline, young households can buy houses at lower prices and gain from the price appreciation in future periods once the economy recovers from the recession. However, when housing prices are fixed, young households are subject to larger welfare losses compared to the baseline economy because they cannot benefit from buying houses at lower prices. This mechanism is in line with the findings of Glover et al. (2020) who show that young households can buy assets at depressed prices, and therefore the welfare losses are smaller than older households.

3.6 Demographic structure and unemployment shocks

The analysis so far has assumed a uniform age distribution of the population and that the unemployment shock leads to a uniform increase in the unemployment rate and income risk for all workers in the economy. In this section, I explore the role of age for the spillover effects of unemployment on housing prices, focusing on the consequences of both age-dependent unemployment shocks and demographic shifts. Unemployment shocks disproportionately affecting younger workers might have larger consequences on the housing market as the housing demand of younger workers are more sensitive to income shocks (see Section 3.5). Moreover, long-term changes in the demographic structure, such as the growing share of older households, change the housing demand patterns and the consequences of recessions. By analyzing both the impact of age-dependent unemployment shocks and the consequences of demographic shifts, this section provides a comprehensive study of the impact of population age structure on the spillover effects of unemployment on the housing market.

Figure 3.8. Different age-group specific unemployment shocks

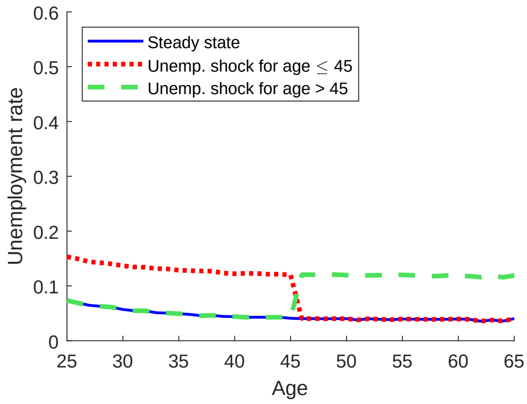

Notes: This figure shows the average unemployment rate in the steady state of the economy and under different structures of the unemployment shock. The solid line displays the average unemployment rate in the steady state. The dotted line shows the average unemployment rate when workers below age 45 are subject to the unemployment shock. The dashed line shows the average unemployment rate when workers above age 45 are subject to the unemployment shock.

3.6.1 Age-dependent unemployment shocks

The structure of the unemployment shock, more specifically, whether the shock disproportionately affects workers in certain age groups, might play a key role for the magnitude of the spillover effect on housing prices. Previous research has found that different recession periods disproportionately affected workers of different age groups. For example, young people suffered most during the Great Recession (Bell and Blanchflower, 2011; Hoynes, Miller, and Schaller, 2012). Moreover, studies covering the post-war period in the United States have shown that the macroeconomic volatility is U-shaped in age, the young experiencing larger labor market volatility than older workers (Clark and Summers, 1981; Ríos-Rull, 1996; Gomme et al., 2004). Also, Jaimovich and Siu (2009) highlight the importance of demographic structures for business cycle volatility.

To investigate whether and how the structure of the unemployment shock shapes the spillover effect, I now examine two distinct scenarios where the unemployment shock is age-dependent: in the first case, only young workers up to age 45 are impacted, with their unemployment rate and their risk of permanent income losses increasing. In the second case, only old workers above age 45 are subject to the increase in unemployment rate and income risk. In both cases, the model assumes an increase of the unemployment rate by 8 percentage points to match the average increase in the baseline model and the risk of persistent income losses increases to 0.8. Other assumptions and parameters remain unchanged from the baseline model. Figure 3.8 shows the changes in unemployment rates under these alternative unemployment shocks.

Figure 3.9. House price and income after different unemployment shocks

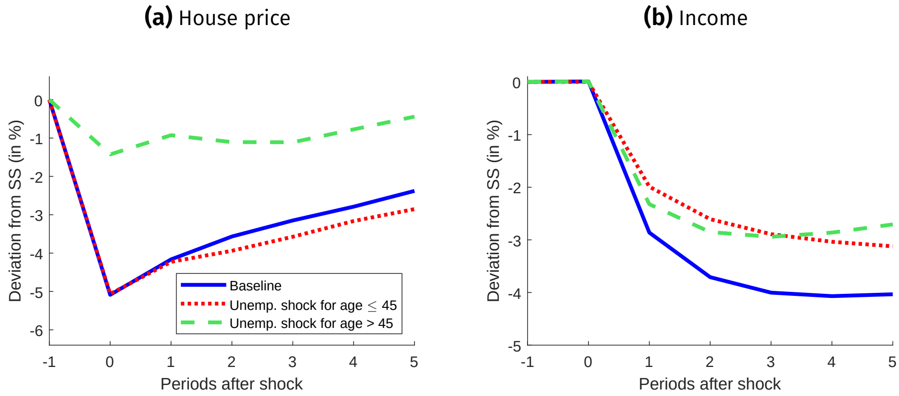

Notes: This figure shows the consequences of different unemployment shocks where either only young workers (below age 45) or only old workers (above age 45) are subject to higher unemployment risk compared to the steady state. Panel 3.9a shows the dynamics of housing price and Panel 3.9b the dynamics of income after the unemployment shock hits the economy. In both panels, the solid lines refer to the baseline economy. The dotted and dashed lines refer to the experiments where only young workers and only old workers are affected by the unemployment shock, respectively.

The results of the counterfactual experiment, summarized in Figure 3.9, reveal important differences in the spillover effects of unemployment shocks on the housing market. The unemployment shock affecting young workers (up to age 45) leads to a large drop in the housing price by more than 5% as in the baseline economy. Additionally, the recovery of the housing price is slower than in the baseline economy, remaining around 0.5 percentage points below the price path of the baseline economy in the following periods. When only older workers (above age 45) are subject to the unemployment shock, the spillover effect on the housing market is smaller and the housing price decreases only by 1.5% when the shock occurs. The impact on the housing price is also less persistent. After the unemployment shock, the housing price quickly recovers to its steady-state level, which indicates that the housing demand of older households remains relatively stable. Hence, the same size of the unemployment shock generates a much larger spillover effect on the housing market when many young workers are subject to the unemployment shock.

One reason why the unemployment shock to young workers generates larger spillover effects on the housing market is because of their significant role in driving housing demand. The economic condition of young workers is an important determinant of housing prices, as many young workers are either first-time house buyers or want to climb the housing ladder by selling the current house and purchasing a larger one. Another reason why the unemployment shock to young workers generates larger spillover effects is that they suffer larger lifetime income losses upon unemployment compared to older workers whose remaining years in the labor market are expected to be much shorter. Figure 3.9b shows that an unemployment shock affecting workers under age 45 leads to a more persistent gap in income that remains at −3% even 5 years after the unemployment shock. In contrast, when only old workers are affected by the unemployment shock, the income level recovers more quickly.

3.6.2 Population age structure

In this section, I analyze how shifts in the demographic structure affect the spillover effects of unemployment on housing prices. The demographic composition of the U.S. population is projected to change considerably over the next decades. According to Vespa, Armstrong, Medina, et al. (2018), the population aged 65 and older is projected to double by 2060. These demographic changes will have important implications for the housing market in the economy: first, the demanded housing size is different across age groups and second, households at different stages of life react differently to changes in housing prices and other economic conditions which also affect the demand for housing.

The relationship between demographic structures and the housing and asset markets has been widely studied in the literature. Poterba (2001) finds no significant relationship between the share of population in the prime saving years and the real returns on financial assets, while Leombroni et al. (2020) show that the asset market participation of baby boomers led to a drop in wealth relative to GDP. Focusing on the housing market, Levin, Montagnoli, and Wright (2009) show that aging and shrinking population leads to decreasing housing prices. Similarly, Gong and Yao (2022) show that changes in demographics can explain the housing price growth from 1970 to 2010.

At first glance, the impact of the age structure of the population on the spillover effect is ambiguous. On the one hand, a decline in the working-age population could dampen the spillover effect, because a lower share of households would face the direct consequence of higher unemployment risk, leading to a smaller decline in housing demand and prices. On the other hand, housing demand of younger households may be more sensitive to changes in price. As a consequence, the decline in housing prices could be amplified as fewer young households buy housing to profit from lower prices during a recession.

To gain intuition, consider a simple OLG economy where we have two generations: the young and the old. Let $h_o(p)$ and $h_y(p)$ denote the total housing demand of old and young households, respectively, as a function of the housing price p. Define Δh_o and Δh_y as the housing demand deviation in recession from the steady state of old and young households, respectively. It can be shown that

$$\frac{\Delta h_y}{\hat{h}_y(\hat{p})} \Big/ \frac{\Delta p}{\hat{p}} \quad < \quad \frac{\Delta h_o}{\hat{h}_o(\hat{p})} \Big/ \frac{\Delta p}{\hat{p}} \tag{3.9}$$

is a sufficient condition for the housing price to decline less strongly during a recession when the share of old households increases. Δp denotes the housing price difference in a recession and in the steady state. The inequality (3.9) compares the elasticity of housing demand with respect to prices in a recession and in the steady state. Hence, if the housing demand elasticity of the young households is smaller than the elasticity of the old households, the inequality in (3.9) is satisfied, implying that the spillover effect of unemployment shock on the housing price is mitigated. Appendix 3.A.10 provides the detailed derivation of the above expression.

Now, two conditions are sufficient to satisfy the inequality in (3.9). First, the elasticity of the young households is negative if their housing demand during a recession is lower than in the steady state at the equilibrium housing prices. Second, the elasticity of the older households is positive as they are not directly impacted by the recession and they would increase their housing demand if prices fall during a recession. The first assumption that the total housing demand of young households declines during a recession is plausible given the empirical evidence in Section 3.3. The second assumption regarding the change in housing demand of retired households is more ambiguous. This condition is satisfied in the current model as the credit conditions remain unchanged in a recession and retired households, if anything, increase their housing demand if prices decline. However, if credit conditions become tighter during times of high unemployment, the demand for housing from retired households could also change significantly, and the inequality in 3.9 might not be necessarily satisfied.

In the following analysis, I conduct a steady-state comparison of two economies which are otherwise identical but differ in terms of the demographic structures. More specifically, the new steady state assumes that 65% of the population consists of working-age households, while the remaining 35% of the population are retired households. In contrast, in the baseline model, approximately 75% and 25% of the population consist of working-age and retired households, respectively.

The results in Table 3.9 show that the demographic structure is key to the magnitude of the spillover effect of unemployment shocks on housing prices. In the new steady state with a larger share of retired households, the drop in housing price following the unemployment shock is 4% smaller compared to the baseline economy. Hence, in the model at hand, an increase in the share of retired population mitigates the spillover effect. The reason is that the reduction in housing demand by working-age households, who are directly affected by unemployment shocks, becomes smaller compared to the additional housing demand of retired households. The income of retired households remains unaffected by the unemployment shock, and as a result, their housing demand slightly increases when housing prices decline. These findings suggest that the demographic composition of an economy can significantly affect the magnitude of the spillover effect from unemployment to the housing market. Under plausible assumptions, a larger share of retired population mitigates the impact of unemployment shocks on housing prices.

Table 3.9. Housing price deviation

Time	Model	
	Baseline	Aged economy
0	−5.09	−4.89
1	−4.16	−3.87
2	−3.57	−3.39
3	−3.15	−3.00
4	−2.79	−2.72
5	−2.38	−2.37

Notes: This table shows the percentage deviation of the housing price from its steady-state level following an unemployment shock (uniform across all ages) at time zero.

3.7 Unemployment insurance and housing prices

The previous sections have shown that unemployment shocks have large spillover effects on the housing prices, resulting in negative welfare consequences for old, retired households. An interesting question is which policy tools may mitigate this effect and stabilize housing prices. One natural candidate is to increase the generosity of the unemployment insurance system during times of high unemployment. Unemployment insurance systems provide benefits to workers who become unemployed, offering insurance against temporary income losses and helping them to smooth consumption. This insurance effect could potentially reduce the drop in housing demand and, in turn, mitigate the spillover effect on housing prices and stabilize the housing market.

This section examines whether, and to what extent, the UI system can reduce the spillover effects of unemployment shocks on the housing market. It is important to note that the theoretical framework applied in this paper is not designed to evaluate the optimal design of UI systems. In particular, the model is calibrated on an annual basis, whereas it is necessary to have a model calibrated to higher-frequency data to accurately capture worker flows in and out of unemployment. Moreover, the income process in this model is exogenous, and there is no job search and endogenous unemployment duration of workers. These components play a key role in analyzing the trade-offs when designing optimal UI systems. Therefore, while this section evaluates whether the UI system dampens the spillover effects of unemployment shocks on the housing market, it does not discuss the optimal design of the policy.

In the United States, the generosity of the UI system tends to increase during recessions. A notable example is the Great Recession: although benefits are typically available for a maximum of 26 weeks, programs such as the Emergency Unemployment Compensation (EUC) and Extended Benefits (EB) have provided additional weeks of support during the Great Recession (Mueller, Rothstein, and Wachter, 2016; Hsu, Matsa, and Melzer, 2018). During the Covid-19 pandemic, the Federal Pandemic Unemployment Compensation (FPUC) supplement increased the weekly UI benefits by $600, leading to a substantial increase in the replacement rate for eligible workers (Ganong, Noel, and Vavra, 2020). In the following analysis, two scenarios are considered to examine the effects of the UI system for the spillover effect of unemployment shocks to housing prices. The baseline economy assumes a replacement rate of 60%, while the alternative system provides a higher replacement rate of 80%.

3.7.1 Baseline economy

Figure 3.10. Higher replacement rate of the unemployment insurance system

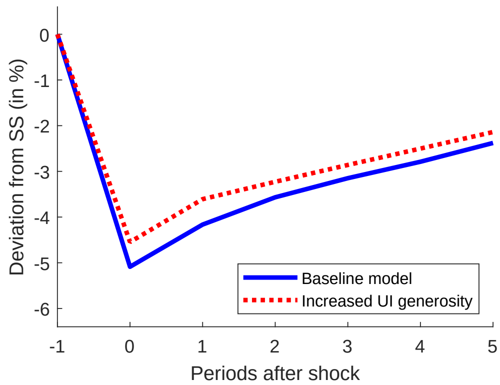

Notes: This figure shows the spillover effects of unemployment shocks on housing prices for the baseline UI replacement rate of 60% (solid lines) and for a UI system with a replacement rate increased to 80% (dotted lines). The results refer to the dynamics of housing prices after an unemployment shock that affects workers at all ages.

Figure 3.10 shows the effect of increasing the UI replacement rate on the spillover effect of unemployment shocks on housing prices. The solid line represents the baseline model discussed in Section 3.5 with a UI replacement rate of 60%, while the dotted line displays the results in an economy where the UI replacement rate is increased to 80% in the first two years of the recession. The increase in the UI replacement rate partly stabilizes the housing market by reducing the spillover effect of unemployment shocks on housing prices by 12%.[7] Hence, increasing UI

7. The effect on the spillover effect is computed by integrating the housing price drop in the baseline economy and in the economy with higher UI replacement rate and comparing these total spillover effects in the two economies.

generosity during recessions partly mitigates the spillover effect, but is not able to fully counteract it. In Appendix 3.A.11, I show that higher UI generosity reduces the welfare losses of retired households from a recession.

However, the impact of UI generosity does not fully mitigate the spillover effect. The reason is that workers experience losses in their permanent income upon unemployment and households shift their portfolios away from housing wealth, which is illiquid, towards more liquid assets. The main driver of the spillover effect is the increased unemployment risk rather than the actual drop in household income, as shown by Figure 3.4 of Section 3.5. While providing insurance against temporary income losses, UI systems do not protect workers from permanent income losses upon unemployment. Hence, the insurance against the temporary income loss provided by higher UI replacement rate does not have a large effect on housing demand, and consequently, the spillover effect on housing prices remains significant.

3.7.2 Model without permanent income losses

In the model at hand, an increase in UI benefit generosity does not significantly change the spillover effect of unemployment shocks on housing prices. The main reason is that higher UI benefits do not provide insurance against permanent income losses. When unemployment results in permanent income losses, households shift their portfolios away from housing toward other more liquid assets.

To confirm that permanent income losses upon unemployment are the main reason why the generosity of UI benefits does not greatly impact the spillover effect, I consider an alternative model that assumes that workers do not experience permanent income losses when they become unemployed. In this alternative framework, workers only lose their incomes during the period of unemployment, after which their incomes fully recover to their pre-unemployment income paths. All other model assumptions remain unchanged from the baseline model, and I calibrate the size of the transitory income drop during unemployment to match the empirically observed drop in housing prices following an unemployment shock. This calibration ensures that the alternative model generates the same spillover effect from an unemployment shock to housing prices as in the baseline model. Next, I repeat the analysis from the previous section by increasing the UI benefit replacement rate from 60% to 80% in the first two periods of the recession, using the unemployment shock that increases unemployment risk for all workers in the economy.

Figure 3.11 displays the dynamics of the housing prices under different UI replacement rates. In the model without permanent income losses, increasing the UI replacement rate from the baseline of 60% to 80% leads to a remarkably smaller decline in housing prices: higher UI replacement rate reduces the decline in housing price by almost one-fifth compared to the baseline economy.

In summary, this section explores the potential of the UI system as a stabilizer of the housing market, reducing the spillover effect of unemployment shocks on hous-

Figure 3.11. Counterfactual without permanent income losses

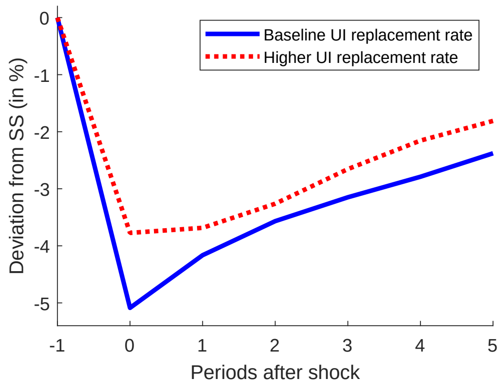

Notes: This figure shows the results of a counterfactual experiment in which unemployed workers are not subject to permanent income losses. Upon unemployment, workers only lose their income in the unemployed period. The figure considers the dynamics of housing prices after an unemployment shock for the baseline UI replacement rate of 60% (solid line) and for a UI system with a replacement rate increased to 80% (dotted line).

ing prices. The analysis shows that increasing the UI replacement rate from 60% to 80% in the beginning of the recession periods reduces the spillover effect by 12% in the baseline model. The spillover effect is not fully mitigated because workers face permanent income losses when unemployed so that they shift their portfolio toward more liquid assets during times of high unemployment risk, which in turn leads to a drop in housing demand and housing prices. While a higher UI replacement rate provides temporary income insurance and helps smooth consumption, it does not provide insurance against the risk of permanent income losses. These findings suggest that labor market policies that stabilize the economy by preventing layoffs during recessions rather than temporary income support, such as short-time work programs, could be more effective in mitigating spillover effects from unemployment on housing prices.

3.8 Conclusion

In this paper, I analyze the spillover effect of unemployment risk for young, working-age households on the old, retired households through the housing market. The paper offers an empirical and a theoretical contribution. First, I empirically show that the size of the spillover effect is large. To this end, I employ the Current Population Survey (CPS) which provides data on state-level unemployment rates and the industry composition of employment in the United States. Using a shift-share instrument to address potential endogeneity and omitted variable bias, I find that an increase in the unemployment rate by one percentage point leads to a decline in the housing price by 1.55%. An analysis in the Home Mortgage Disclosure Act (HMDA) data reveals that there is a strong negative correlation between mortgage applications and unemployment rates, implying that housing demand decreases during times of high unemployment in the economy.

In a next step, I develop an overlapping generations model with a housing market and unemployment risk. In the model, workers choose a portfolio of housing and liquid assets where the housing price is determined in the housing market. The model is calibrated using the data from the Survey of Consumer Finances (SCF). The calibrated model produces spillover effects of unemployment shocks on the housing prices that are similar in magnitude as in the data. Using the model, I study the aggregate consequences of an unemployment shock. More than 60% of the spillover effect of unemployment risk on housing prices is driven by an increase in income uncertainty when the unemployment shock occurs so that households reduce their housing demand. In contrast, the actual drop in income due to higher unemployment rate only accounts for 40% of the spillover effect.

Old, retired households bear significant welfare losses even though they are not directly affected by the increase in unemployment risk: the spillover effect of unemployment risk on housing prices transmits approximately one-third of the welfare losses of working-age households to retired households measured in terms of consumption equivalent variation. Due to the spillover effect, retired households see a drop in housing wealth which constitutes an important share in wealth of old households. Young workers experience significant welfare losses, but they partly benefit from buying houses at depressed prices.

Moreover, I show that unemployment shocks disproportionately affecting younger workers generate a spillover effect that is more than 3 times larger than those affecting older workers. As young workers are typically on the demand side on the housing market, their economic condition is an important driver of the spillover effect on housing prices. The demographic structure of the economy is also key for the magnitude of the spillover effect. I find that a demographic shift towards a larger population share of retired households mitigates the spillover effect.

Finally, I show that increasing the generosity of the unemployment insurance during recession partly stabilizes housing prices. However, the unemployment insurance system is not able to completely counteract the spillover effect as it does not provide insurance against persistent income losses upon unemployment. The channel through persistent income losses limits the ability of unemployment insurance systems to stabilize housing prices.

3.A Appendix

3.A.1 Data sources

- House prices: House Price Index in U.S. states (Source: Federal Housing Finance Agency)
- Stock Market Index: Wilshire 5000 Price Index (Source: Wilshire Associates, Wilshire Indexes)
- Consumer Price Index: U.S. Bureau of Labor Statistics Consumer Price Index
- Housing supply: New housing permits (U.S. Census Bureau, U.S. Department of Housing and Urban Development, New Residential Construction)
- Population growth: U.S. Census Bureau, Annual Estimates of the Population for the U.S. and States
- Federal Funds Effective Rate: Board of Governors of the Federal Reserve System (U.S.), H.15 Selected interest rates
- Real GDP growth: U.S. Bureau of Economic Analysis, Gross Domestic Product

3.A.2 IV regression

The IV approach in Eq. 3.2 uses the Bartik instrument to identify exogenous movements in the state-level unemployment rate. As a robustness check, the IV regression in this section considers the relative changes in state-level unemployment rate over time. While the construction of the Bartik instrument remains unchanged as in Eq. 3.1, the second-stage regression is given by

$$\Delta \log(HPI_{s,t}) = \beta \Delta \hat{u}_{s,t} + \delta X_{s,t-\tau} + \gamma Z_{t-\tau} + \lambda_s + \zeta_t + \alpha_1 + \eta_{s,t} \tag{3.A.1}$$

where $\Delta u_{s,t} = \log u_{s,t} - \log u_{s,t-1}$. Table 3.A.1 shows the first-stage regression results. The second-stage results are summarized in Table 3.A.2.

Table 3.A.1. First-stage regression

Changes in unemployment rate	All data	Excluding Great Recession
	(1)	(2)
Bartik	−2.61***	−2.67***
	(0.64)	(0.66)
Observations	2091	1938
R^2	0.74	0.72

Notes: This table shows the first-stage regression for the IV analysis. The relative changes in state-level unemployment rates are regressed on the Bartik shocks including state and year fixed effects, state-level characteristics (average household income, population growth, and ratio of young to old households), and national characteristics (GDP growth, national unemployment rate, stock price index, consumer price index, Federal funds rate, new housing permits, and the supply of new housing). The state-level and national characteristics contain the lags of these variables up to four years. Columns (1) presents the results when the regression analysis considers all data from 1978 to 2019. Columns (2) excludes the period of Great Recession from the considered data. Standard errors are adjusted for clustering at the state level and reported in parentheses.

Table 3.A.2. IV regression

$\Delta\log(\text{HPI})$	All data			Excluding Great Recession		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemp. rate	-0.32^{***}	-0.32^{***}	-0.21^{***}	-0.29^{***}	-0.29^{***}	-0.21^{***}
	(0.08)	(0.08)	(0.07)	(0.08)	(0.08)	(0.08)
Observations	2091	2091	2091	1938	1938	1938
R^2	0.34	0.34	0.49	0.29	0.29	0.43
State and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
House supply controls	No	Yes	Yes	No	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes

Notes: This table summarizes the state-panel IV regressions of log changes in housing prices on the instrumented changes in unemployment rates. The controls include state and year fixed effects, state-level characteristics (average household income, population growth, and ratio of young to old households), and national characteristics (GDP growth, national unemployment rate, stock price index, consumer price index, Federal funds rate, new housing permits, and the supply of new housing). The state-level and national characteristics contain the lags of these variables up to four years. Columns (1)-(3) presents the results when the regression analysis considers all data from 1978 to 2019. Columns (4)-(6) exclude the period of Great Recession from the considered data. Standard errors are adjusted for clustering at the state level and reported in parentheses.

3.A.3 Home Mortgage Disclosure Act data

Figure 3.A.1. Mortgage applications and unemployment rate excluding the Great Recession

(a) Number of loans **(b)** Share of denied applications

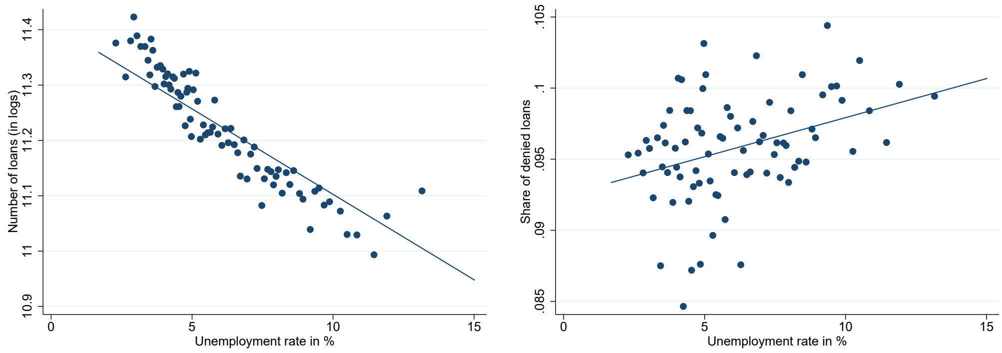

Notes: Panel (a) shows the number of mortgage applications against unemployment rate by each state. Panel (b) shows the share of mortgage applications that have been rejected. The results are after controlling for GDP, nationwide unemployment rate, and state- and year-fixed effects.

The empirical analysis on the relationship between unemployment rates and mortgage applications using the Home Mortgage data in Section 3.3 reveals a strong negative relationship. In order to check that this relationship is not solely driven by the Great Recession, I repeat the analysis excluding the data in the years from 2008 until 2010. Figure 3.A.1 summarizes the results. In Figure 3.A.1a, there is a strong negative correlation between the number of mortgage applications and the unemployment rate at the state level. Regarding the share of denied mortgage applications, Figure 3.A.1b shows a positive correlation between unemployment rate and the share of mortgage applications rejected by the financial institutions. These results indicate that the negative correlation between housing demand and unemployment rate remains strong even when the Great Recession is excluded.

3.A.4 Calibration

Figure 3.A.2. Properties of unemployment shock

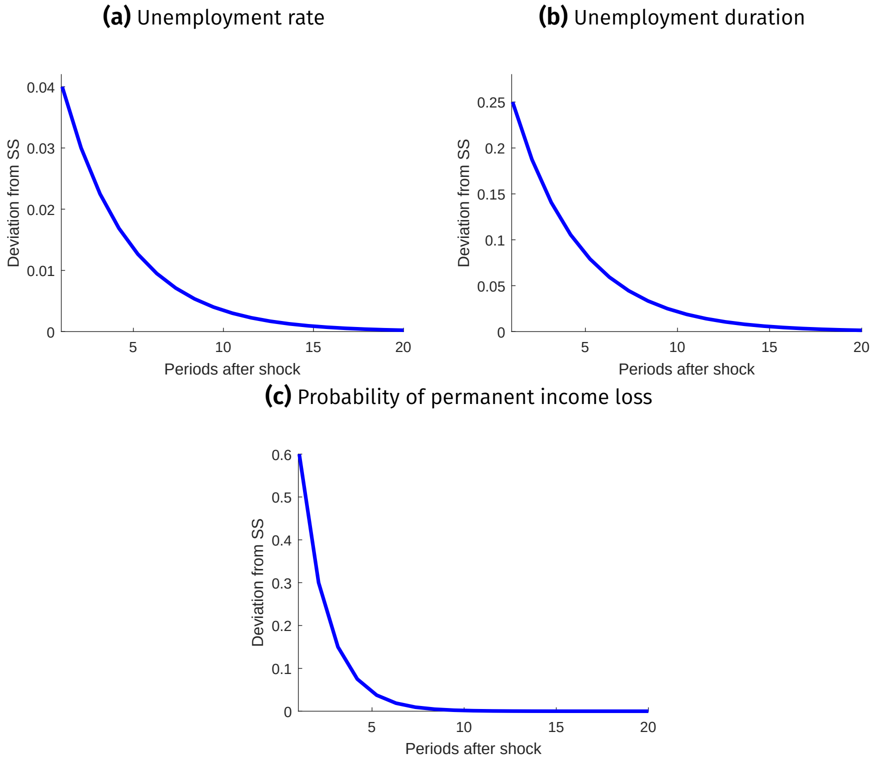

Notes: This figure shows the the evolution of the unemployment shock over time. Panel 3.A.2a and 3.A.2b show the percentage deviation of unemployment rate and unemployment duration from the steady state. Panel 3.A.2c displays the percentage deviation of permanent income loss probability from the steady state.

Figure 3.A.2 shows the evolution of the unemployment shock over time. The unemployment shock affects the unemployment rate (Figure 3.A.2a), the unemployment duration (Figure 3.A.2b), and the probability of permanent income loss (Figure 3.A.2c). When the unemployment shock occurs, the average unemployment rate in the economy increases by 4 percentage points. The unemployment duration increases from one quarter to 6 months. The probability of permanent income losses jumps from 0.2 to 0.8. For the unemployment rate and the unemployment duration, the shock decays at a rate of 0.75, while the probability of persistent income loss is assumed to decay at a rate of 0.5.

3.A.5 House size grid

The house size grid $\mathscr{H}$ follows Kaplan, Mitman, and Violante (2020) who let each grid point represent a stylized housing unit composed of a fixed set of rooms. They normalize the size of a bedroom to one and scale the size of the living room, kitchen, and bathrooms proportionally. Each successive grid point represents a house with one additional room. Comparing the house size changes upon moving between their model and data from the Panel Study of Income Dynamics (PSID) and American Housing Survey (AHS), they show that the percentage changes in housing size in the data and the changes in the housing units in the model are comparable. To accommodate the fact that my model includes only homeowners, I augment the house size grid $\mathscr{H}$ by adding a smallest grid point of 0.01. This allows the model to represent young households who do not yet own a home, for example, those living with their parents, but who still derive a minimal amount of utility from housing services. This extension preserves the structure of the model with homeowners only, while capturing early life-cycle housing situations without formal ownership.

3.A.6 Transitional equilibrium

The aggregate shock is introduced as a one-time, unanticipated shock at $t = 0$. The shock increases the probability of unemployment, the unemployment duration, and the probability of earnings loss as described in Section 3.4.5. Households are initially in the steady state of the economy. When the shock occurs, households have perfect foresight and know that the economy will return to the initial steady state after the shock decays. The solution algorithm is as follows.

(1) Choose a period T in which the economy is assumed to have returned to the initial steady state.

(2) Guess a path for the house prices $(\{\hat{p}_t\}_{t=0}^T)^0$.

(3) Solve the value functions and policy functions backwards from $t = T-1, \ldots, 0$ where households in T have the value and policy functions in the initial steady state.

(4) Starting from the steady-state distribution of households, the economy is simulated forward from $t = 0, \ldots, T$ using the value and policy functions and the exogenous states of the economy.

(5) At each t, the equilibrium house prices p_t is computed using the distribution of households.

(6) Compute the maximum difference between the guess and the equilibrium house price $\xi = \max|p_t - \hat{p}_t|$ in all periods.

(7) If $\xi < 10^{-5}$, the solution has been found.

(8) If $\xi \geq 10^{-5}$, update the guess $(\{\hat{p}_t\}_{t=0}^T)^1 = \alpha(\{\hat{p}_t\}_{t=0}^T)^0 + (1-\alpha)\{p_t\}_{t=0}^T$ and repeat the steps from 3.

3.A.7 Aggregate dynamics

Figure 3.A.3. Ratio of liquid assets to income

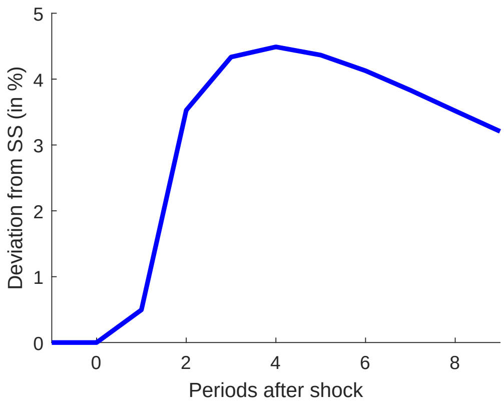

Notes: This figure shows the percentage deviation of the ratio of liquid assets to income from its steady-state level after the unemployment shock (uniform across all ages) hits the economy at time zero.

Figure 3.A.3 displays the steady-state deviation of the ratio of average liquid assets to average household income in the economy. When the unemployment shock occurs, households accumulate additional liquid savings as buffer stocks in response to the increase in income uncertainty. When the unemployment shock abates, the ratio of liquid assets to income decreases as the average income recovers and households face lower uncertainty and households reshuffle their portfolio.

3.A.8 Life-cycle consequences of unemployment shock

What are the individual consequences of unemployment on income, housing, and wealth accumulation? To address this question, I conduct a counterfactual experiment. The first counterfactual case assumes that workers never experience unemployment over the whole life cycle. In the second counterfactual, workers become unemployed at age 26 when the economy is in a normal state. The last counterfactual assumes that workers become unemployed at age 26 during a recession. In the last two scenarios, workers become unemployed only at the beginning of the life cycle and experience no further unemployment spells thereafter.

Figure 3.A.4. Life-cycle profiles after unemployment shock at age 26

(a) Income **(b)** House

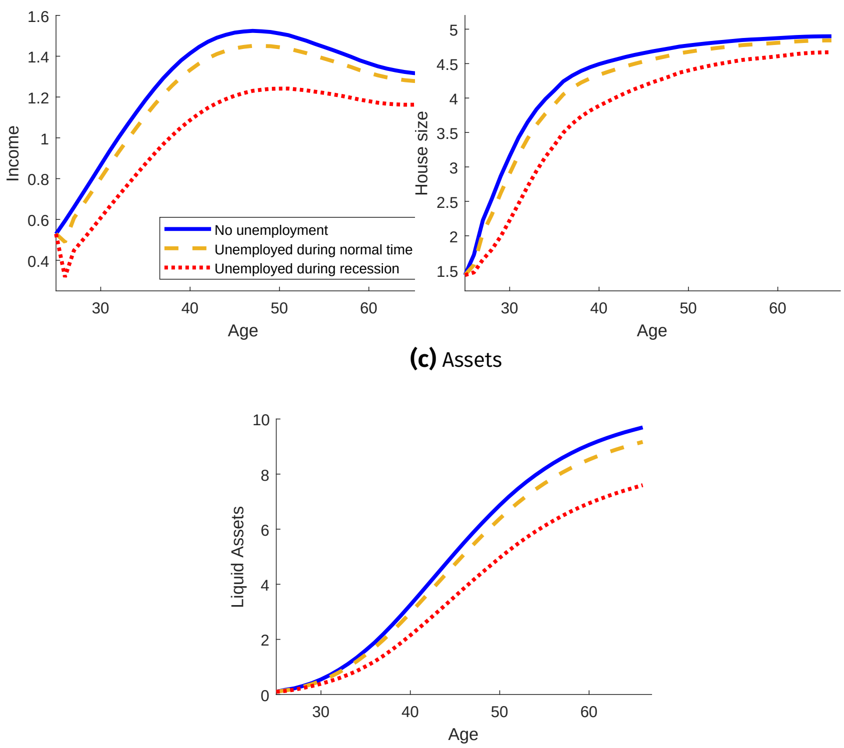

(c) Assets

Notes: This figure shows the life cycle profiles of labor income, house size, and liquid asset. The solid lines display the life-cycle profiles of workers who never experience unemployment. The dashed and dotted lines show the life-cycle profiles of workers who become unemployed at age 26 when the economy is in a normal state and in a recession, respectively.

Figure 3.A.4 presents the life-cycle profiles of income, house size, and assets under the three counterfactual cases. Figure 3.A.4a shows that unemployment when the economy is in a normal state leads to 13% drop in average income, whereas unemployment during a recession leads to a almost 40% drop in average income. In both cases, unemployment results in permanent income losses, the average income level being 6% lower 15 years after being unemployed during a good state of the

economy and 22% lower 15 years after an unemployment spell during a recession. The permanent income losses after unemployment depresses housing demand over the life cycle as shown in Figure 3.A.4b. Compared to the baseline case without unemployment, workers have on average 15% smaller houses when they become unemployed in a recession. On the contrary, being unemployed during a good state has only has a small effect on housing demand of around 8%. Finally, Figure 3.A.4c shows the life-cycle profiles of assets. The permanent loss of income reduces both the housing demand and the life-cycle saving motive of workers. At the end of the working phase, workers who become unemployed in a recession have on average 21% lower assets.

Figure 3.A.5. Life-cycle profiles after unemployment shock at age 45

(a) Income **(b)** House **(c)** Assets

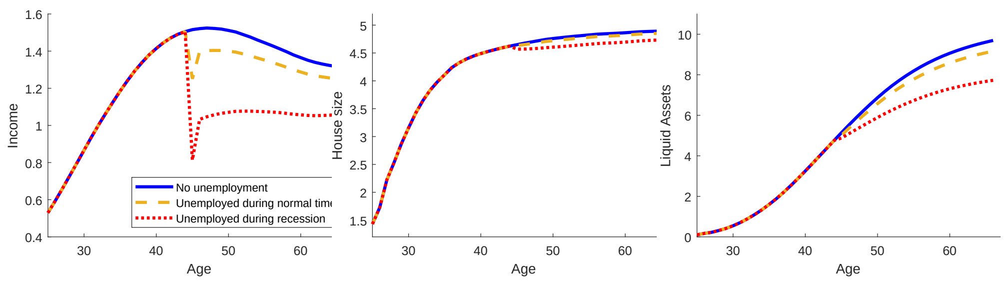

Notes: This figure shows the life cycle profiles of labor income, house size, and liquid asset. The solid lines display the life-cycle profiles of workers who never experience unemployment. The dashed and dotted lines show the life-cycle profiles of workers who become unemployed at age 45 when the economy is in a normal state and in a recession, respectively.

Figure 3.A.5 shows the life-cycle profiles of income, house size, and liquid assets when households become unemployed at age 45. Similar to the results in Figure 3.A.4, households at age 45 experience the largest decline in their income when unemployed during a recession. The average house size and liquid assets also decline following an unemployment shock. The decline in house size is, however, less pronounced than in the counterfactual case where households experience unemployment at age 26.

Figure 3.A.6 compares the change in the average house size between the two counterfactuals where households become unemployed at age 26 and 45, respectively. When becoming unemployed at age 26, the average house size declines by 30% compared to a control group who do not experience unemployment and recovers slowly over time. In contrast, the average house size of households becoming unemployed at age 45 only declines by less than 5%.

The change in the house size is stronger when young households become unemployed due to the following two reasons. First, older households already own large houses, have higher income, and have accumulated more liquid wealth over the life cycle compared to young households. When they become unemployed, many house-

Figure 3.A.6. House size after unemployment shock at age 26 and 45

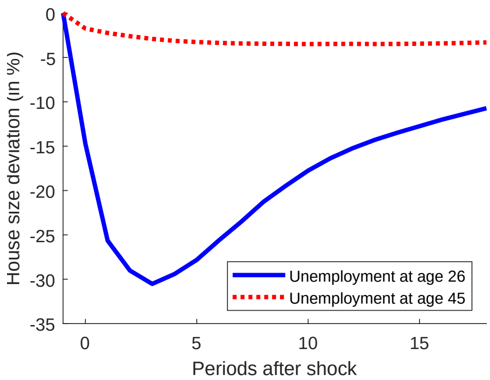

Notes: This figure shows the relative change in house size of workers who experience unemployment at age 26 (solid line) and at age 45 (dotted line) compared to a control group who does not experience unemployment.

holds stay in their house and there is no need to adjust the house size. Younger household who become unemployed own on average smaller houses, and cannot move to larger houses when they are affected by the unemployment shock. The second reason is that the decline in lifetime earnings due to unemployment is larger for young workers than for the older workers. The persistent income losses following an unemployment spell reduces future labor earnings as well, which is more important for the younger workers.

3.A.9 Welfare consequences

Figure 3.A.7. Welfare consequences: max. liquid assets

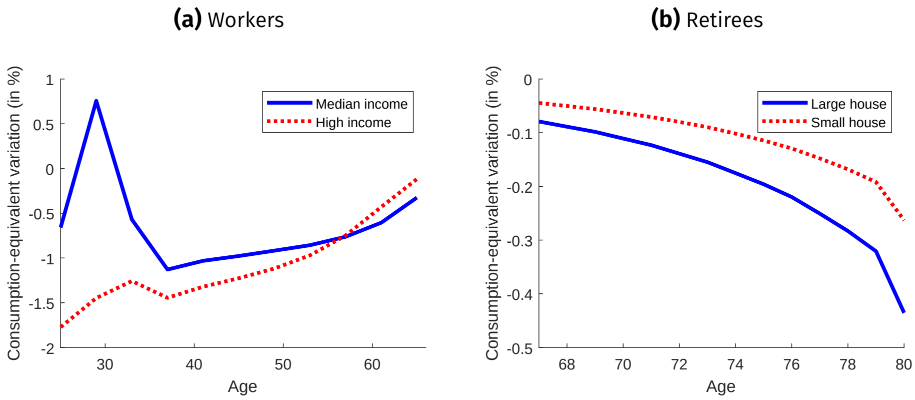

Notes: This figure shows the welfare consequences of a recession for each age in consumption-equivalent variation for households who are rich in liquid assets. Panel 3.A.7a displays the welfare results for workers who have median income (solid line) and workers with higher income (dotted line). Panel 3.A.7b shows the welfare results for retirees who own relatively small houses (solid line) and retirees owning a larger house (dotted line).

This section presents additional welfare results to Section 3.5.3. In Figure 3.A.7, I consider the case where households at all ages are rich in liquid asset. Figure 3.A.7a shows that young workers around age 30 experience a welfare gain from the unemployment shock. As households are not liquidity constraint, they can buy a house at depressed prices and benefit from future house price appreciations, leading to a welfare gain. For households with high income, the risk of permanent income losses is more important such that they still experience a welfare loss from the recession.

Compared to the welfare results in Section 3.5.3, Figure 3.A.7b indicates smaller welfare losses for retired households from the recession. As households are rich in liquid asset, the share of housing wealth in their total wealth is smaller such that the relative decline in housing wealth is muted. As a consequence, a high amount of liquid wealth decreases the welfare losses from a recession for the retired households.

Figure 3.A.8. Welfare consequences: no GE effects on house prices

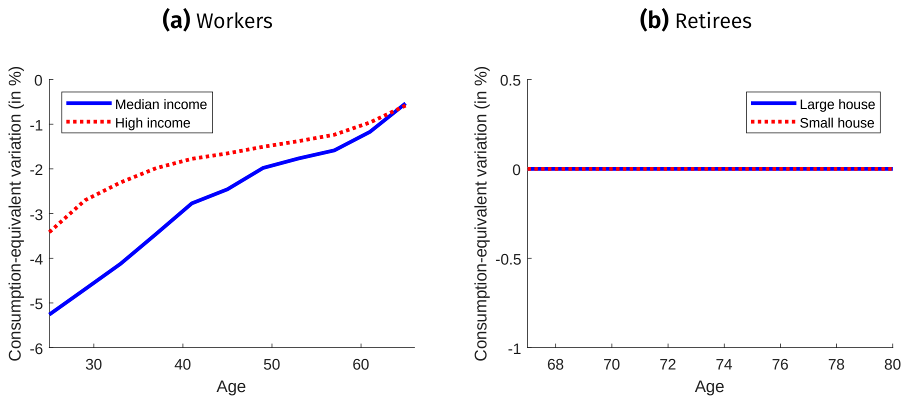

Notes: This figure shows the welfare consequences of a recession for each age in consumption-equivalent variation when house prices are fixed at the steady-state prices. Panel 3.A.8a displays the welfare results for workers who have median income (solid line) and workers with higher income (dotted line). Panel 3.A.8b shows the welfare results for retirees who own relatively small houses (solid line) and retirees owning a larger house (dotted line).

Figure 3.A.8 shows the welfare consequences in the counterfactual where house prices are fixed at the steady-state levels. Young workers bear larger welfare losses compared to the baseline economy as they cannot benefit from the decline in house prices. Old households do not experience any welfare losses as their housing wealth remains unaffected by the recession. As a consequence, the profiles in Figure 3.A.8b remain flat at all ages.

Figure 3.A.9 shows the welfare results for households who are rich in liquid wealth in the counterfactual economy where house prices are fixed at the steady-state levels. Looking at the working-age population in Figure 3.A.9a, the welfare losses are much smaller for households with large liquid wealth: at age 25, the welfare losses are approximately 2% in terms of consumption-equivalent variation. The

Figure 3.A.9. Welfare consequences: max. liquid assets and no GE effects on house prices

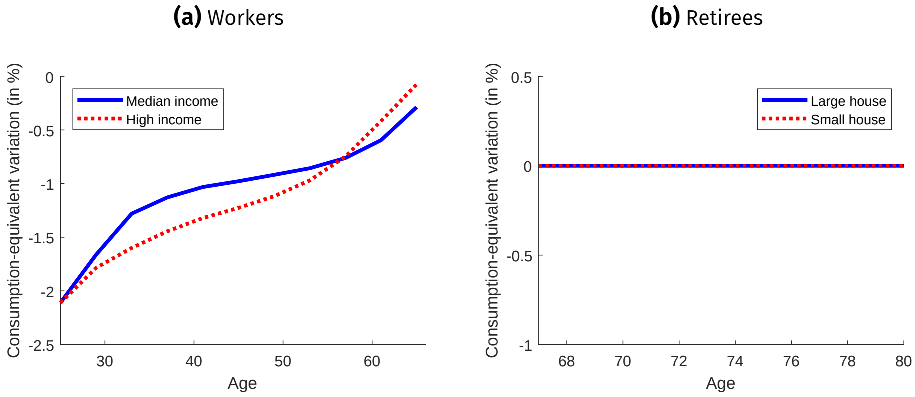

Notes: This figure shows the welfare consequences of a recession for each age in consumption-equivalent variation for households who are rich in liquid assets when house prices are fixed at the steady-state prices. Panel 3.A.9a displays the welfare results for workers who have median income (solid line) and workers with higher income (dotted line). Panel 3.A.9b shows the welfare results for retirees who own relatively small houses (solid line) and retirees owning a larger house (dotted line).

welfare losses from the recession are smaller for workers with median income as these households can afford a house when house prices decline during the recession. In contrast, for high-income workers, the change in house prices is less important than for the median-income workers and the persistent income losses upon unemployment leads to larger declines in income for them. Hence, high-income workers face larger welfare losses from a recession. The welfare results for retired households are displayed in Figure 3.A.9b. Similar to the previous results in Figure 3.A.8, retired households do not bear any welfare costs because the house prices do not move over time and their housing wealth remains unaffected from the unemployment shock.

3.A.10 Population age structure and housing demand

Let $h_o(p)$ and $h_y(p)$ denote the total housing demand of old and young households, respectively, as a function of the housing price p. Housing demand always decreases in price p, i.e. $h'_o(p) < 0$ and $h'_y(p) < 0$. Let λ denote the population share of old households. Then, the total housing demand in the steady state of the economy is given by

$$\hat{H}(\hat{p}) \;=\; \lambda \hat{h}_o(\hat{p}) + (1-\lambda)\hat{h}_y(\hat{p}). \tag{3.A.2}$$

where the hats denote the steady-state values. The relative change in total housing demand when a recession occurs is

$$\eta \;=\; \frac{H^r(p^r) - \hat{H}(\hat{p})}{\hat{H}(\hat{p})} \tag{3.A.3}$$

where $H^r(p^r)$ denotes the total housing demand in the recession at price p^r. In a recession, housing demand at a given price p is assumed to be lower than in the steady state so that $H^r(p) < \hat{H}(p)$. Plugging in the expression from Equation (3.A.2) yields

$$\begin{aligned}
\eta \;&=\; \frac{H^r(p^r) - \hat{H}(\hat{p})}{\hat{H}(\hat{p})} \\
&=\; \frac{\lambda h^r_o(p^r) + (1-\lambda) h^r_y(p^r) - \lambda \hat{h}_o(\hat{p}) - (1-\lambda)\hat{h}_y(\hat{p})}{\lambda \hat{h}_o(\hat{p}) + (1-\lambda)\hat{h}_y(\hat{p})} \\
&=\; \frac{\lambda\left[h^r_o(p^r) - \hat{h}_o(\hat{p})\right] + (1-\lambda)\left[h^r_y(p^r) - \hat{h}_y(\hat{p})\right]}{\lambda \hat{h}_o(\hat{p}) + (1-\lambda)\hat{h}_y(\hat{p})}
\end{aligned}$$

Define $\Delta h_o := h^r_o(p^r) - \hat{h}_o(\hat{p})$ and $\Delta h_y := h^r_y(p^r) - \hat{h}_y(\hat{p})$. Then, we obtain

$$\begin{aligned}
\eta \;&=\; \frac{\lambda \Delta h_o + (1-\lambda)\Delta h_y}{\lambda \hat{h}_o(\hat{p}) + (1-\lambda)\hat{h}_y(\hat{p})} \\
&=\; \frac{\Delta h_y + \lambda\left(\Delta h_o - \Delta h_y\right)}{\hat{h}_y(\hat{p}) + \lambda\left(\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right)}.
\end{aligned}$$

The partial derivative of η with respect to λ yields

$$\begin{aligned}
\frac{\partial \eta}{\partial \lambda} \;&=\; \frac{\left(\Delta h_o - \Delta h_y\right)\cdot\left[\hat{h}_y(\hat{p}) + \lambda\left(\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right)\right]}{\left[\hat{h}_y(\hat{p}) + \lambda\left(\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right)\right]^2} \\
&\quad - \frac{\left[\Delta h_y + \lambda\left(\Delta h_o - \Delta h_y\right)\right]\left(\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right)}{\left[\hat{h}_y(\hat{p}) + \lambda\left(\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right)\right]^2}
\end{aligned}$$

It holds that η is equal to zero at market clearing prices $p^r = \hat{p}^r$ because the total housing supply is assumed to be fixed. Now, consider a new steady state with a higher share of old households compared to the initial steady state. The equilibrium housing price in the initial economy is denoted by p_1^r in recession, while p_2^r denotes the equilibrium housing price in the new economy during a recession. Assume that the housing supply and steady-state housing price remains at the same level in the new steady state. If $\frac{\partial \eta}{\partial \lambda} > 0$, then in the new steady state with a higher share of old population, the housing price p_2^r has to increase in equilibrium. In this case, the change in housing price is $\hat{p} - p_2^r < \hat{p} - p_1^r$, implying that the drop in housing price during a recession becomes smaller in absolute value. We want to assess in which cases we obtain $\frac{\partial \eta}{\partial \lambda} > 0$. Simplifying the above expression yields

$$\begin{aligned} &\left(\Delta h_o - \Delta h_y\right) \cdot \hat{h}_y(\hat{p}) + \lambda\left(\Delta h_o - \Delta h_y\right) \cdot \left[\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right] \\ &-\Delta h_y \cdot \left[\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right] - \lambda\left(\Delta h_o - \Delta h_y\right) \cdot \left[\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right] \quad > \quad 0 \\ \Leftrightarrow &\left(\Delta h_o - \Delta h_y\right) \cdot \left[\lambda \hat{h}_o(\hat{p}) + (1-\lambda)\hat{h}_y(\hat{p})\right] > \left(\lambda \Delta h_o + (1-\lambda)\Delta h_y\right) \cdot \left[\hat{h}_o(\hat{p}) - \hat{h}_y(\hat{p})\right] \end{aligned}$$

Rearranging yields

$$\begin{aligned} \Delta h_o \cdot \hat{h}_y(\hat{p}) \quad &> \quad \Delta h_y \cdot \hat{h}_o(\hat{p}) \\ \Leftrightarrow \quad \frac{\Delta h_y}{\hat{h}_y(\hat{p})} \quad &< \quad \frac{\Delta h_o}{\hat{h}_o(\hat{p})} \\ \Rightarrow \quad \frac{\Delta h_y}{\hat{h}_y(\hat{p})} \Big/ \frac{p^r - \hat{p}}{\hat{p}} \quad &< \quad \frac{\Delta h_o}{\hat{h}_o(\hat{p})} \Big/ \frac{p^r - \hat{p}}{\hat{p}} \\ \Leftrightarrow \quad \frac{\Delta h_y}{\hat{h}_y(\hat{p})} \Big/ \frac{\Delta p}{\hat{p}} &< \frac{\Delta h_o}{\hat{h}_o(\hat{p})} \Big/ \frac{\Delta p}{\hat{p}}. \end{aligned} \tag{3.A.4}$$

If the inequality is satisfied, a higher share of old population leads to, all else equal, a smaller spillover effect on housing prices.

3.A.11 Welfare implications of higher generosity of unemployment insurance

The findings in Section 3.7 indicate that the unemployment insurance system partly stabilizes housing prices during a recession. In the following, I analyze the changes in welfare consequences of an unemployment shock induced by higher generosity of the UI system. In particular, the goal is to assess whether increasing UI benefits mitigates the welfare consequences on old, retired households who suffer from the spillover effect of unemployment risk on the housing prices.

Table 3.A.3 summarizes the results. Comparing the welfare results in the economy with higher UI generosity during a recession (column 3) to the baseline economy (column 2), it becomes evident that retired households also gain from the increase in the UI generosity. As a more generous UI system stabilizes the housing prices, the housing wealth of retired workers is stabilized as well during a recession,

Table 3.A.3. Welfare decomposition

Age	Welfare effects (in %)	
	Baseline	Higher UI generosity
25	−3.72	−3.83
30	−3.99	−3.97
35	−3.65	−3.60
40	−2.91	−2.87
45	−2.46	−2.43
50	−1.95	−1.92
55	−1.66	−1.63
60	−1.31	−1.28
65	−0.63	−0.59
70	−0.19	−0.17
75	−0.39	−0.34
80	−1.03	−0.81
85	−1.06	−0.96

Notes: The welfare effects are measured in consumption-equivalent variation on remaining lifetime consumption. In all models, the welfare effects are computed for a median household in the economy in terms of assets, employment status, income, housing size, and remaining mortgage balance. The baseline model refers to the results in Figure 3.7. The model with a more generous unemployment insurance system refers to the results in Figure 3.10.

and as a consequence, the welfare losses of retired households are mitigated. Interestingly, the welfare losses of the young households (age 25) is slightly larger in the economy with higher UI generosity, as a stabilization of housing prices reduces the asset price gain of young households.

References

Attanasio, Orazio P, Renata Bottazzi, Hamish W Low, Lars Nesheim, and Matthew Wakefield. 2012. "Modelling the demand for housing over the life cycle." *Review of Economic Dynamics* 15 (1): 1–18. [129]

Baily, Martin Neil. 1978. "Some aspects of optimal unemployment insurance." *Journal of public Economics* 10 (3): 379–402. [130]

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute. [132]

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden. 2019. "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk." *Econometrica* 87 (1): 255–90. [129, 145]

Bell, David N. F., and David G. Blanchflower. 2011. "Young people and the Great Recession." *Oxford Review of Economic Policy* 27 (2): 241–67. [152]

Chambers, Matthew, Carlos Garriga, and Don E Schlagenhauf. 2009. "Accounting for changes in the homeownership rate." *International Economic Review* 50 (3): 677–726. [129]

Chang, Yongsung, Jay H Hong, and Marios Karabarbounis. 2018. "Labor market uncertainty and portfolio choice puzzles." *American Economic Journal: Macroeconomics* 10 (2): 222–62. [129]

Chetty, Raj. 2008. "Moral hazard versus liquidity and optimal unemployment insurance." *Journal of political Economy* 116 (2): 173–234. [130]

Clark, Kim B, and Lawrence H Summers. 1981. "Demographic Differences in Cyclical Employment Variation." *Journal of Human Resources* 16 (1): 61–79. [152]

Corbae, Dean, and Erwan Quintin. 2015. "Leverage and the foreclosure crisis." *Journal of Political Economy* 123 (1): 1–65. [129]

Davis, Steven J., Till von Wachter, Robert E. Hall, and Richard Rogerson. 2011. "Recessions and the Costs of Job Loss [with Comments and Discussion]." *Brookings Papers on Economic Activity,* 1–72. [125, 142]

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies* 71 (3): 743–68. [136]

Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium." *Journal of Political Economy* 125 (1): 140–223. [129]

Ganong, Peter, Pascal Noel, and Joseph Vavra. 2020. "US unemployment insurance replacement rates during the pandemic." *Journal of Public Economics* 191: 104273. [157]

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2020. "Intergenerational redistribution in the great recession." *Journal of Political Economy* 128 (10): 3730–78. [129, 130, 151]

Gomme, Paul, Richard Rogerson, Peter Rupert, and Randall Wright. 2004. "The business cycle and the life cycle." *NBER macroeconomics annual* 19: 415–61. [152]

Gong, Yifan, and Yuxi Yao. 2022. "Demographic changes and the housing market." *Regional Science and Urban Economics* 95: 103734. [154]

Hagedorn, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman. 2013. "Unemployment benefits and unemployment in the great recession: the role of macro effects." Working paper. National Bureau of Economic Research. [130]

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13 (1): 15–51. Special issue: Cross-Sectional Facts for Macroeconomists. [139]

Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. "Who Suffers during Recessions?" *Journal of Economic Perspectives* 26 (3): 27–48. [152]

Hsu, Joanne W., David A. Matsa, and Brian T. Melzer. 2018. "Unemployment Insurance as a Housing Market Stabilizer." *American Economic Review* 108 (1): 49–81. [157]

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709. [125]

Jaimovich, Nir, and Henry E. Siu. 2009. "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility." *American Economic Review* 99 (3): 804–26. [152]

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy* 128 (9): 3285–345. [129, 136, 139, 165]

Kaplan, Greg, and Giovanni L. Violante. 2014. "A MODEL OF THE CONSUMPTION RESPONSE TO FISCAL STIMULUS PAYMENTS." *Econometrica* 82 (4): 1199–239. [139]

Kleibergen, Frank, and Richard Paap. 2006. "Generalized reduced rank tests using the singular value decomposition." *Journal of Econometrics* 133 (1): 97–126. [133]

Kopczuk, Wojciech, and Joseph P Lupton. 2007. "To leave or not to leave: The distribution of bequest motives." *Review of Economic Studies* 74 (1): 207–35. [138]

Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2018. "A macroeconomic approach to optimal unemployment insurance: Theory." *American Economic Journal: Economic Policy* 10 (2): 152–81. [130]

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The housing market (s) of San Diego." *American Economic Review* 105 (4): 1371–407. [129]

Larkin, Kieran. 2019. "Job risk, separation shocks and household asset allocation." In *2019 Meeting Papers.* 1058. Society for Economic Dynamics. [129]

Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers. 2020. "Inflation and the Price of Real Assets." Working Paper, Working Paper Series 26740. National Bureau of Economic Research. [154]

Levin, Eric, Alberto Montagnoli, and Robert E. Wright. 2009. "Demographic Change and the Housing Market: Evidence from a Comparison of Scotland and England." *Urban Studies* 46 (1): 27–43. [154]

Mueller, Andreas I., Jesse Rothstein, and Till M. von Wachter. 2016. "Unemployment Insurance and Disability Insurance in the Great Recession." *Journal of Labor Economics* 34 (S1): S445–S475. [157]

Paz-Pardo, Gonzalo. 2024. "Homeownership and portfolio choice over the generations." *American Economic Journal: Macroeconomics* 16 (1): 207–37. [129]

Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. "Housing, consumption and asset pricing." *Journal of Financial Economics* 83 (3): 531–69. [139]

Poterba, James M. 2001. "Demographic Structure and Asset Returns." *Review of Economics and Statistics* 83 (4): 565–84. [154]

Ríos-Rull, José-Víctor. 1996. "Life-Cycle Economies and Aggregate Fluctuations." *Review of Economic Studies* 63 (3): 465–89. [152]

Vespa, Jonathan E, David M Armstrong, Lauren Medina, et al. 2018. *Demographic turning points for the United States: Population projections for 2020 to 2060.* US Department of Commerce, Economics / Statistics Administration. [154]