Three Essays on the Interconnectedness of Labor Markets and Household Finance

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch die

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

vorgelegt von

Gašper Ploj

aus Šempeter pri Gorici, Slowenien

2021

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Moritz Kuhn
Zweitreferent:	Prof. Petr Sedláček, Ph.D.
Tag der mündlichen Prüfung:	11. Juni 2021

Acknowledgements

No man is an island, even less so during doctoral studies. This dissertation would likely not have seen the light of day without the support and advice of many people whom I risk not properly mentioning here.

First of all, I would like to thank my supervisors Moritz Kuhn and Petr Sedláček for their guidance and support over the past few years. I have grown as a researcher and an economist as a result of the many stimulating discussions we had over the course of my doctoral research. Special thanks go also to Thomas Hintermaier and Christian Bayer, whose technical expertise I drew upon at crucial steps while working on the third chapter of this dissertation.

My doctoral studies in Bonn were enabled by the generous support I received from the DAAD, the Bonn Graduate School of Economics, the Institute for Macroeconomics and Econometrics, and the DFG Research Training Group 2281 »Macroeconomics of Inequality«. I am deeply grateful to all those involved in making this possible, among others Silke Kinzig, Britta Altenburg, Vanessa Pollari, and Heide Baumung.

I have also been extremely lucky to be surrounded by a wonderful group of friends and colleagues during my years in Bonn. We not only persevered through challenges but also built many joyful memories together. Jana, Thomas, Marta, Marek, Lucas, Si, Carl, Axel (and many others) I thank you all for making this journey so much more memorable and enjoyable! Additionally, I also want to thank all my friends and colleagues from the institute and the BGSE with whom I had the pleasure to share countless intellectually stimulating lunches in the past years.

I thank my longtime friends Polona, Lucija, Špela, Tilen, and Veronika for the many Skype talks and visits, especially in those times when I needed most a distraction from the daily Ph.D. life. You have been there through the highs and lows of this experience and I am lucky to count you among my best friends. My deep gratitude goes also to Žiga who provided many encouraging words while I was trying to finish the dissertation in the midst of a pandemic.

Finally, I am indebted to my family who, unlike me, never doubted that I would at some point reach the stage when I could write these final words.

Contents

List of Figures

List of Tables

Introduction

This dissertation consists of three self-contained essays that investigate the interconnectedness of labor markets and household finance. The first two chapters focus on differences in workers' labor market experience and on how this heterogeneity influences their financial decisions. The first chapter shows that differences in unemployment risk produce long-lasting effects on lifetime earnings, consumption, and wealth, while the second chapter shows that an important part of the observed heterogeneity in households' portfolio choices can be rationalized by the empirically observed differences in labor market risk. In the third chapter the perspective changes, as this final chapter analyses how shocks in financial markets propagate through household balance sheets to produce macroeconomically relevant effects on the labor market. Using the evidence from the Great Recession, I show that households' limited ability to obtain mortgage credit in the aftermath of the financial crisis of 2007 resulted in fewer firms being created, which consequently negatively affected aggregate job creation.

Chapter 1, which is joint work with Moritz Kuhn, explores the economic consequences of large differences in job stability that can be observed in the labor market. While many workers hold lifetime jobs, some cycle repeatedly in and out of employment. Motivated by this empirical fact, we explore the labor market, financial, and welfare consequences of heterogeneity in job stability and document them at the individual and macroeconomic level. We use the data from the Survey of Consumer Finances (SCF) to document a systematic positive relationship between job stability and wealth accumulation. Per dollar of income, workers with more stable careers hold more wealth. We also develop a life-cycle consumption-saving model with heterogeneity in job stability that is jointly consistent with empirical labor market mobility, earnings, consumption, and wealth dynamics. Using this structural model, we quantify the consequences of heterogeneity in job stability. At the individual level, we explore the life-cycle consequences of early-career heterogeneity in job stability and find that a bad start to the labor market leaves long-lasting scars. The income and consumption level for a worker who starts working life from an unstable job is, even 25 years later, 5 percent lower than that of a worker who starts with a stable job. For the macroeconomy, we explore the welfare consequences of changes in job stability in the context of declining U.S. labor market dynamism. We find welfare gains of 1.6 percent of lifetime consumption for labor market entrants from a secular decline in U.S. labor market dynamism.

Chapter 2, which is joint work with Christian Bayer, Thomas Hintermaier, and Moritz Kuhn, studies the heterogeneity in labor market risk and its implications for wealth accumulations and portfolio choices of households. Using granular data from German retirement accounts we document large and persistent heterogeneity in life-cycle income risks. We find risk to be concentrated in a small group of high-risk individuals and, at the same time, highlight a negative correlation between risk and income growth. Combining labor market and wealth data, we further document that low-risk, high-growth individuals build up more wealth and that portfolio differences amplify income differences as low-risk individuals invest more in illiquid, high-return assets. We propose a life-cycle model that qualitatively and quantitatively explains the empirical patterns of risk, wealth accumulation, and portfolios based on persistent risk types and incomplete, frictional asset markets with systematic return and liquidity differences.

Chapter 3 studies the importance of home equity as a source of funding for new firms and analyzes whether changes in the ability of entrepreneurs to use home equity for startup capital can explain movements in the aggregate firm and job creation. Using detailed mortgage loan data, I find that the contraction of mortgage credit availability experienced in the aftermath of the 2007 financial crisis can account for a statistically and economically significant portion of the overall decline in firm creation during the Great Recession. In line with the presented survey evidence and the existing literature, the effect is bigger for smaller firms, in sectors that are more reliant on home equity financing, and in areas where housing provides better collateral. At the same time, evidence from the Survey of Consumer Finances suggests that entrepreneurs were able to partially compensate for the decline in mortgage credit availability by modifying their funding mix towards other costlier types of credit, which mitigated the decline in firm creation.

Chapter 1

Job stability, earnings dynamics, and life-cycle savings*

Joint with Moritz Kuhn

1.1 Introduction

Labor markets are characterized by large heterogeneity in job stability. Some workers hold lifetime jobs while others cycle repeatedly in and out of unstable employment (Hall, 1982). Motivated by this empirical fact, we explore the labor market, financial, and welfare consequences of such heterogeneity in job stability and document them at the individual and macroeconomic level. At the individual level, we explore the life-cycle consequences of early-career heterogeneity in job stability. For the macroeconomy, we explore the welfare consequences of changes in job stability in the context of declining U.S. labor market dynamism (Molloy, Trezzi, Smith, and Wozniak, 2016).

The paper offers an empirical and a theoretical contribution. First, we provide empirical evidence from the Survey of Consumer Finances (SCF) for a systematic relationship between job stability (tenure) and accumulated wealth. We document that households with more stable jobs accumulate, controlling for income, more wealth. Life-cycle savings are an important driver of this correlation, but even after controlling for age, we find a systematic positive relationship between wealth-to-income ratios and job stability. To quantify the extent of heterogeneity in job stability

* 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 universities in Manchester, Liverpool, Erasmus University Rotterdam, 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. Kuhn thanks the DFG for financial support (DFG No. 433368816). Ploj gratefully acknowledges financial support from the DAAD, the Bonn Graduate School of Economics, and the DFG research training group 2281 "The Macroeconomics of Inequality". The usual disclaimer applies.

in the data, we propose a measure of employment inequality. Using this measure, we find employment inequality to be increasing with age, and during the middle of working life, the average job in the data lasts three times longer than expected in the absence of heterogeneity. Regarding the sources of job-stability heterogeneity, we provide empirical and Monte Carlo evidence that point to employer differences as the important source of such heterogeneity.

The second contribution of the paper is theoretical. We incorporate a frictional labor market model with human capital investment into an otherwise standard life-cycle model of consumption-saving behavior. We demonstrate that this model is jointly consistent with life-cycle earnings, consumption, and wealth dynamics. Given the model's empirical success, we use it to explore the economic consequences of heterogeneity in job stability. At the individual level, we find that differences in job stability at labor market entry leave long-lasting scars on income and consumption. On average, a worker with a bad start to the labor market (i.e., an unstable job at age 25) will have a persistent income gap with 5 percent lower income compared to an otherwise identical worker starting from a stable job. The difference in consumption between these workers is almost 20 percent at age 25, and this gap closes only slowly to the level of the income gap at the end of working life. The welfare effects from changes in job stability are therefore large. For the average 25-year-old worker, the transition to the least stable job is associated with a welfare loss of 1.4 percent of lifetime consumption, but welfare costs can also exceed 10 percent for workers in stable and high-paying jobs early in life. At the macroeconomic level, we explore the decline in U.S. labor market dynamism as a combination of lower job-to-job mobility and a shift in the distribution of job stability, in line with the empirical evidence (Fallick and Fleischman, 2004; Fujita, 2018; Molloy, Smith, and Wozniak, 2020). A less dynamic labor market leads to a welfare *gain* for labor market entrants. In terms of lifetime consumption, labor market entrants are willing to forgo 1.6 percent of consumption to start working life in a less dynamic economy. The most important reason for the welfare gain is better opportunities for human capital investment from higher job stability.

Our model combines a life-cycle labor search model with human capital investment and a consumption-saving model with incomplete financial markets. In the labor market, workers search on and off the job and jobs are heterogeneous with respect to wages and separation rates to nonemployment. Separation rate differences are one determinant of differences in job durations; on-the-job search with workers climbing the job ladder constitutes a second source for differences in job durations. Human capital investment opportunities exist only for employed workers who can exert effort to invest in their human capital. Thus, unstable careers with low employment rates perpetuate low incomes by offering fewer opportunities for human capital investment. The consumption-saving part of the model is standard, with agents facing incomplete financial markets where they can 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, 1993b), the distribution of earnings growth (Guvenen, Karahan, Ozkan, and Song, 2019), wealth dynamics in Panel Study of Income Dynamics (PSID) data, and the joint distribution of income and wealth in SCF data. Most importantly, the model also matches the empirically documented relationship between job stability and wealth accumulation so that we can interpret this empirical correlation through the lens of the model.

To explore the individual consequences of heterogeneity in job stability, we first decompose life-cycle wealth accumulation. The life cycle is an important dimension of heterogeneity in job 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 job stability and low income when young creates an additional tension between the precautionary and life-cycle savings motive for young workers in unstable jobs. We find that the saving rates of 25-year-old workers in unstable jobs are up to ten times higher compared to workers in stable jobs and that, in the absence of the risk of job loss, workers would not save at all at that age. A bad start to the labor market therefore substantially mitigates workers' ability to engage in life-cycle consumption smoothing. For the typical 40-year-old worker, we find that roughly one out of three dollars is saved for precautionary reasons, but for older workers, the importance of precautionary savings quickly diminishes. Workers age 50 hold less than 7 percent of their wealth for precautionary reasons.

To further trace out the consequences of a bad start to the labor market, we compare two identical young workers who differ only in job stability: one worker starts working life in a stable job (25th percentile of age-specific separation rates), whereas the second worker has a bad start to the labor market and starts working life in an unstable job (75th percentile of age-specific separation rates). Comparing income and consumption dynamics between these workers shows that a bad start to the labor market leaves large and long-lasting scars. We find that after one year, incomes between the two workers already differ on average by 13 percent and that this difference remains significant over their entire working life. The consumption difference starts larger at almost 20 percent and closes to 6 percent at the end of working life. The difference between income and consumption dynamics results from the interplay between job stability and earnings growth. Starting from an unstable job subsequently leads to less stable employment and 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. Most job offers represent dead-end jobs that create high average labor market mobility for the macroeconomy. 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, Karahan, Ozkan, and Song (2019), who emphasize the importance of heterogeneity in nonemployment in accounting for life-cycle earnings dynamics in U.S. data. For savings behavior, these income dynamics imply that a bad start to the labor market ties workers to a kind of "Sisyphus cycle" of buffer stock savings where they build up and run down their buffer stock of wealth while cycling in and out of unstable employment (Carroll, 1997). By contrast, starting working life in a stable lifetime job allows workers to engage in life-cycle consumption smoothing from the start. Through the lens of the model, we see that such Sisyphus cycles also account for the empirically observed relationship between job stability and wealth accumulation. The underlying mechanism is the labor market dynamics that intertwine earnings growth and the volatility of earnings.

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 out by running down wealth. With respect to previous job stability, 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.

What are the welfare costs of a bad start to the labor market? We derive welfare costs as a consumption-equivalent variation for 25-year-old workers that we move from their current job to the job with the lowest job stability. We find that the welfare costs from such a change in job stability can be as large as 11 percent of lifetime consumption for workers in stable and well-paying jobs. For the typical worker, welfare costs are still large with 1.4 percent of lifetime consumption. We decompose welfare costs into components related to human capital accumulation, incomplete financial markets, and labor market frictions. For the median-wage worker, the effects from lower human capital and additional labor market search account for 40 percent of the welfare effect. Worse consumption smoothing accounts for the remaining 20 percent. For low-paying but stable jobs (e.g. apprenticeships) we find that the opportunity to invest in human capital by far accounts for the largest part of the welfare costs of lowering job stability.

At the macroeconomic level, we study the consequences of the secular decline in U.S. labor market dynamism (Molloy, Trezzi, Smith, and Wozniak, 2016). We interpret the decline in labor market dynamism as a combination of fewer opportunities for job-to-job mobility and a shift in the aggregate job stability distribution, in line with the empirical evidence in Fujita (2018), Fallick and Fleischman (2004), and Molloy, Smith, and Wozniak (2020). Lower labor market mobility induces two counteracting forces for welfare: reducing job offer rates on the job directly reduces wage-ladder dynamics, whereas a shift toward more stable jobs leads to better opportunities for human capital investment. Matching parameter changes to the observed decline in labor market mobility, we find for labor market entrants that the job stability effect dominates resulting in a welfare gain of 1.6 percent. The key reason for the welfare gain is higher earnings growth. On average, we find that earnings grew almost 3 percent more at the end of working life. When decomposing earnings growth, we find that two-thirds of the additional growth results from higher human capital, but we also find that wages grow 1 percent more. Higher job stability therefore not only offers better human capital investment opportunities but also makes the wage ladder more stable, resulting in higher life-cycle wage growth. Hence, our results suggest that the decline in labor market dynamics has had a significant *positive* welfare effect for young American workers.

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 individual consequences of heterogeneity in job stability. In Section 1.5, we study the macroeconomic decline in U.S. labor market dynamism. Section 1.6 concludes.

1.1.1 Related literature

Our work relates to two large strands of literature: models of consumption-saving behavior in the presence of idiosyncratic income risk and market incompleteness (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, Mukoyama, Rogerson, and Şahin, 2017; Hubmer, 2018; Larkin, 2019; Cajner, Güner, and Mukoyama, 2020). We add to this literature by exploring the consequences of heterogeneity in job stability. Our paper connects the part of the literature that focuses on macroeconomic dynamics, as in Krusell, 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. 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, Mukoyama, Rogerson, and Şahin (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.[1] Guvenen, Karahan, Ozkan, and Song (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

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

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) and Karahan and Ozkan (2013). 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.

1.2 Heterogeneity in job stability and wealth accumulation in the data

Our empirical analysis consists of two steps. In the first step, we explore the relationship between job stability and wealth accumulation using 25 years of data from the Survey of Consumer Finances (SCF). In the second step, we combine empirical evidence from the Business Dynamics Statistics (BDS), Current Population Survey (CPS), and Monte Carlo simulations to corroborate the large heterogeneity in job stability in the U.S. labor market and the important role of job heterogeneity in accounting for this heterogeneity.

1.2.1 Job stability and wealth accumulation

The SCF is a triennial household survey providing detailed information on income and wealth for a cross section of U.S. households. It has become the key resource on the distribution of income and wealth for the United States (see, for example, Kuhn and Rios-Rull, 2016; Bricker, Dettling, Henriques, Hsu, Jacobs, et al., 2017; Kuhn, Schularick, and Steins, forthcoming). Besides the detailed information on household income and wealth, the SCF also offers information on household members' labor market situation. Exploring the relationship between the labor market situation and wealth accumulation is the focus of the first step of our analysis. For our analysis, we pool data across survey waves from 1992 to 2016 and restrict the sample to households with employed household heads ages 20 to 60. As our model will abstract from self-employment, we drop households with self-employed household heads and households with extreme wage observations, defined as wages lower than 75 percent of the minimum wage.[2] Additionally, we exclude the top 1 percent

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

of households by wealth and earnings as we do not provide a theory of the very right tail of these distributions, and, similarly, we exclude households in the bottom 1 percent by earnings and households with negative wealth.

Regarding the construction of variables, we consider household wealth as the difference between household assets and debt. Household income is gross income from all sources including transfers, and earnings is income from wages and salaries.[3] We control for income differences nonparametrically by always considering wealth-to-income ratios. Job stability itself is unobserved in the data, and we only observe retrospectively whether an employer-employee relationship has been stable by looking at a worker's employer tenure or the number of employers during a worker's career.[4] 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 1.1 provides the estimates of the empirical measures of job stability and their relationship to wealth-to-income ratios. In Figure 1.1a, 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 1.1b, 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. Quantitatively, the observed slope is economically meaningful. The slope implies that having a lifetime job that leads to a 20-year increase in tenure will, on average, lead to additional wealth corresponding to roughly one year of income. Wealth-to-income ratios increase by roughly three over a 30-year life-cycle in the

3. We follow Bricker, Dettling, Henriques, Hsu, Jacobs, Moore, Pack, Sabelhaus, Thompson, and Windle (2017) and Kuhn and Rios-Rull (2016) for the construction of these variables.

4. 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.

data, so that two years of tenure correspond to one year of age when it comes to wealth accumulation.[5]

Figure 1.1. Wealth-to-income ratios, tenure, and number of employers

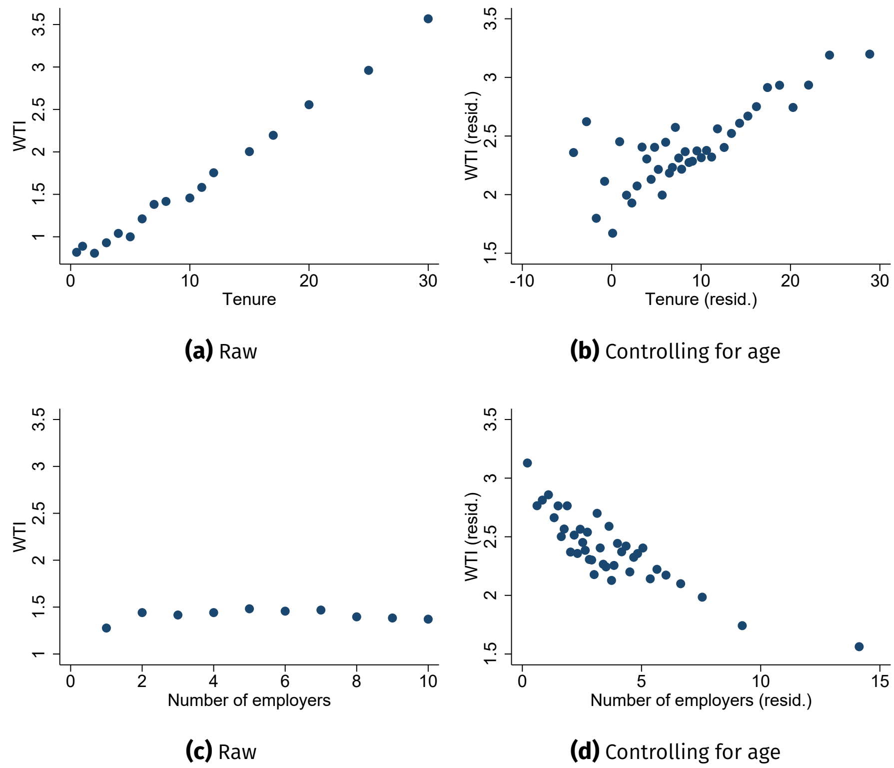

Notes: This figure shows binned scatter plots of wealth-to-income ratios against tenure or number of employers for which a person has worked full-time jobs lasting one year or more. In panels (a) and (c), each dot represents a median wealth-to-income ratio for a given bin. Panels (b) and (d) show binned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

Figure 1.1b offers a second interesting observation. While the relationship in the raw data appears almost perfectly linear (Figure 1.1a), the relationship turns into a U-shape for low tenures after controlling for age. This U-shape relationship means that workers who have low tenure relative to their age group tend to have higher wealth-to-income ratios. As we will see below, this is a characteristic property of the model when job losers accumulate precautionary wealth before layoffs and then get a negative shock to income, so that wealth-to-income ratios increase. Consumption smoothing of job losers will lower wealth-to-income ratios over time and allow them

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

to converge back to their target wealth-to-income ratio if no further job loss will occur.

Figures 1.1c and 1.1d 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 1.1c), 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. One concern might be other confounding factors, most prominently, education. In Appendix 1.A.1, we therefore repeat the analysis with additional controls for education, occupation, industry, and risk attitude. We find that the documented relationship is qualitatively and quantitatively robust.

1.2.2 Heterogeneity in job stability

The seminal paper by Hall (1982) documents large heterogeneity in job stability and the existence of lifetime jobs in the U.S. labor market. The interpretation of the paper is that the source of this heterogeneity stems from job differences. Guvenen, Karahan, Ozkan, and Song (2019) explore life-cycle earnings dynamics in high-quality Social Security data for the United States and document large heterogeneity in life-cycle nonemployment. They emphasize the importance of incorporating heterogeneity in nonemployment to account for observed life-cycle earnings dynamics. Recent work by Morchio (2020) further corroborates large differences in separation rates over the life cycle. 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.

Figure 1.2 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 1.2a, we observe a spreading out of the distribution as workers age. As pointed out in Hall (1982), the typical U.S. worker has a stable employment history. At age 60, more than 50 percent of workers have been with their employer for 10 years, and almost a quarter of workers at age 60 have been at the same employer for at least 25 years. The life-cycle profiles for the number of employers in Figure 1.2b 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 1.2. Tenure and number of employers over the life cycle

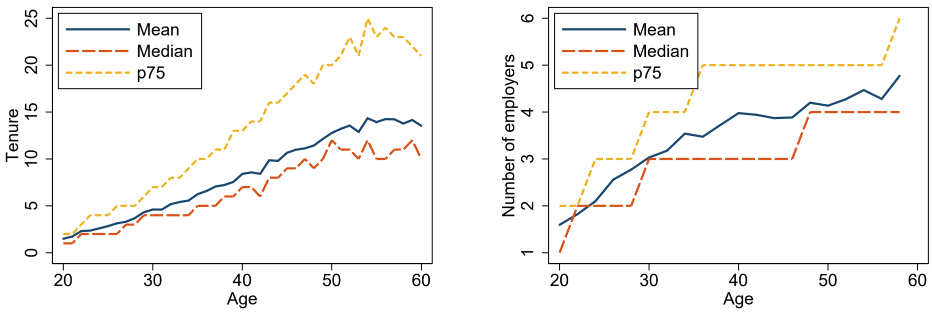

(a) Employer tenure

(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.

The increasing life-cycle dispersion of job stability mirrors the widely studied increase in wage inequality with age (Heathcote, Perri, and Violante, 2010b), 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.2 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 1.3 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 1.4 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 1.4a) and job loss due to firm closure (Figure 1.4b). We remove year and industry fixed effects in both cases.[6] In Figure 1.4a, 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

6. Appendix 1.A.3 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 1.3. Employment inequality

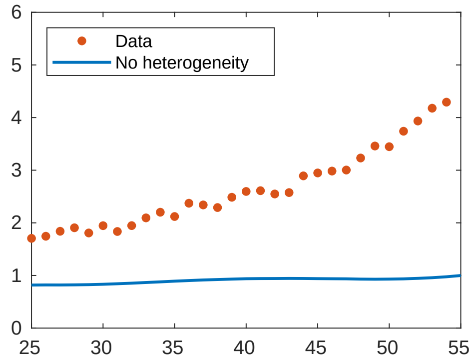

Notes: Estimated life-cycle profile of employment inequality σ_E. Employment inequality is computed as the ratio of observed to expected tenure given average job outflow rates. Underlying data on outflow rates and tenure come from CPS and are taken from Jung and Kuhn (2018). The *no-heterogeneity* case 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.

transitions or worker quits. Figure 1.4b 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.

We provide two Monte Carlo simulations to revisit the extent of heterogeneity in job stability and its sources. Figure 1.5a 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 1.2a. 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 1.3 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 1.5b 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

Figure 1.4. Heterogeneity in job destruction rates by firm age

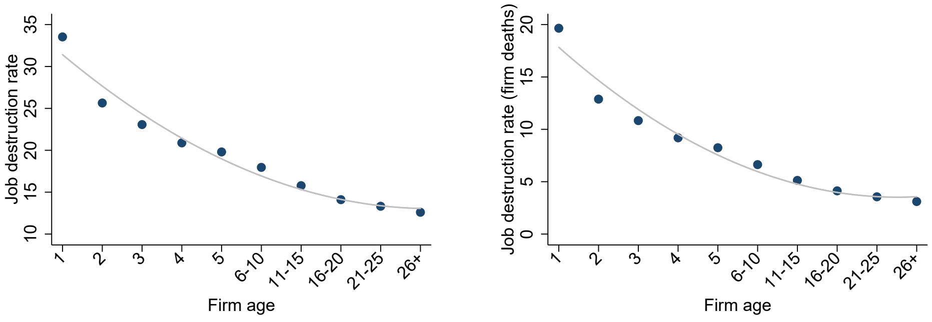

(a) Job destruction rate **(b)** Job destruction rate (firm deaths)

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.

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.4, 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.[7]

To summarize, we document that in the data, job stability is systematically related to wealth accumulation and that the effects are economically significant. We also provide corroborating evidence for large heterogeneity in job stability over the life cycle and in the cross section conditional on age. Finally, we provide evidence in line with job heterogeneity as the source of differences in job stability. The next section develops a model of household saving behavior that explicitly introduces het-

7. 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.

Figure 1.5. Life-cycle profiles and heterogeneity

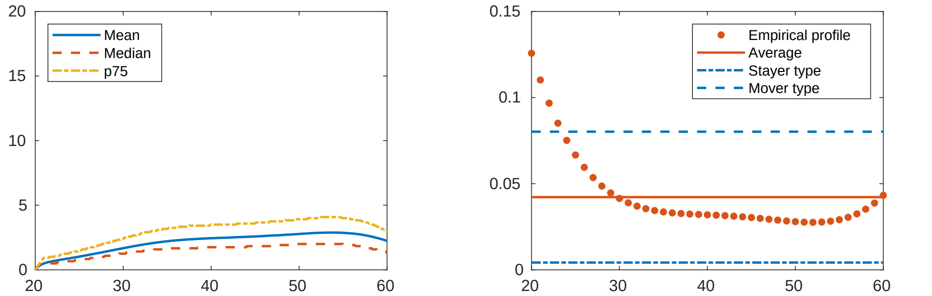

(a) Life-cycle tenure without heterogeneity

(b) Life-cycle mobility with worker 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).

erogeneity in job stability. In a first step, we will use the model to explore the causal effects of heterogeneity in job stability on consumption-saving behavior, earnings dynamics, and welfare at the individual level. In a second step, we interpret the macroeconomic decline in labor market dynamism as a consequence of changes in heterogeneity in job stability and explore the consequences of this macroeconomic phenonemon on the welfare and earnings dynamics of new labor market entrants.

1.3 Heterogeneity in job 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 $y(w_j, h_j, u) = b w_j h_j$ where b denotes the replacement rate and w_j is the wage on the last job. These benefits decline each period if the agent remains 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 $f(w, \lambda)$. 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.[8] One interpretation of this effort provision is as career investment with the current employer (e.g., unpaid overtime, higher work intensity, on-the-job training, or committee work). We assume that human capital levels are discrete and are members of an ordered set with largest (smallest) element h^{max} (h^{min}). We denote by h^+ the immediate successor of human capital level h and by h^- the immediate predecessor of h. Human capital investment is risky. An agent at human capital level h exerting effort t to accumulate human capital has a probability $p_H(t,j)$ of reaching human capital level h^+. We allow for age dependence of $p_H(t,j)$. The law of motion for human capital if the agent exerts effort ($t > 0$) is

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

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

The consumption-saving decision is standard. The agent chooses next period's asset level given her current state and facing a 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 $u(c)$. The working and the transition phase differ only in the possible continuation states. A worker in the transition phase either remains in the transition phase or transits to the retirement phase. A worker in the working phase ages deterministically and transits at the end of prime-age working life to the transition phase. We do not allow workers from the transition phase (retirement phase) to transit back to the working (transition) phase.

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

8. 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.

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

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

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

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

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.5.

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.6.

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 $u(c) = \log(c)$. Human capital takes on discrete values $h_{i,t} \in \{h_1, .., h_{N_h}, h^*\}$, and we partition the support of human capital into two parts. The first part comprises N_h states that we set equidistant in log space between $h_1 = 1$ and $h_{N_h} = 6.5$. The second part is a single high human capital state h^*. We set the human capital state $h^* = 25$ which allows us to match the right tail of the earnings distribution.[9] 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 $\bar{p}_H$. Conditional on reaching human capital level h_{N_h}, a separate age-independent probability p_H^* governs the transitions to state 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.[10] To parametrize the joint 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 mo-

9. This structure is reminiscent of the earnings process in Castaneda, Diaz-Gimenez, and Rios-Rull (2003). The difference here is that the high-income state in our calibration will be much less extreme than the one in Castaneda, Diaz-Gimenez, and Rios-Rull (2003) and much more persistent. We follow similar ideas already proposed in Jung and Kuhn (2018) or Hubmer (2018).

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

ments. 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.[11] In Appendix 1.A.6, 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 1.1 presents the model parameters together with their estimated values.

Table 1.1. Estimated parameters

Parameter	Value	Description
β	0.992	Quarterly discount factor
κ	0.356	Utility cost of effort
π_e	0.404	Probability of a job offer when employed
π_u	0.859	Probability of a job offer when nonemployed
ψ_w	0.532	Marginal distribution of w^*
ψ_λ	0.506	Marginal distribution of $1-\lambda^*$
θ	0.519	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.051	Skill upgrading probability
ρ	0.984	Persistence of skill upgrading probability
p_H^*	0.047	Probability to move to h^*

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 lifetime consumption-equivalent variation between 0.35 percent during the first ten years of working life and less than one-tenth of a percent during the last ten years of working life.[12] 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

11. 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.

12. 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.

(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.4a 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.4b 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. Moving to the highest human capital level h^* is only possible from human capital level h_{N_h} and has a roughly 5 percent probability per quarter. 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. 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 1.6 shows the empirical life-cycle profiles for separation, job-to-job, and job-finding rates and their model counterparts. Looking at the separation rate in Figure 1.6a, we see that its evolution is matched very closely. Model and data show a strong decline up to age 30 and constantly falling separation rates between ages 30 and 50. Figure 1.6b shows that the model also matches the life cycle of job-to-job rates very well, with only a slightly steeper decline of job-to-job rates between ages 20 and 30 compared to the data. 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 1.6c are matched well in level and trend and generally show life-cycle variation. Finally, Figure 1.6d shows the life-cycle profiles of mean, median, and the 75th percentile of the tenure distribution. The model closely matches the empirical increase and heterogeneity in job stability. Importantly, this demonstrates that the model is jointly consistent with high average transition rates (Figures 1.6a and 1.6b) and high job stability for most workers (Figure 1.6d).[13] Appendix Figure 1.A.6 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.[14]

Figure 1.7 turns to the life-cycle dynamics of earnings and wealth. Looking at the life-cycle profile of mean log earnings in Figure 1.7a, we find that the model matches the steep increase in earnings after labor market entry and the flattening out after age 40. It closely matches the large average increase of roughly 0.8 log points over the life cycle but shows slightly less concavity in comparison to its empirical counterpart. Figure 1.7b shows the life-cycle profile of the wealth-to-income ratio as a measure of wealth accumulation. Again, we find a close fit between model and data. Wealth-to-income ratios in model and data rise from slightly above 0 at age 20 to approximately 3.5 at age 55.[15] The empirical profile is slightly less convex than its model counterpart. While the life-cycle profiles are targeted when bringing the model to the data, the relationship between job stability and wealth accumulation in Figures 1.7c and 1.7d is not. In our empirical analysis (Section 1.2), we document a positive correlation between job stability and wealth accumulation after controlling for age effects. Figure 1.7c demonstrates that our model is consistent with this empirical fact. It shows wealth-to-income ratios by tenure in the SCF data and using model-simulated data controlling for age variation nonparametrically. The model data are less dispersed and align well with the SCF data. In particu-

13. For consistency, here we consider transition rates and tenure levels from CPS data. In Appendix 1.A.7, we show that wage and tenure data from the SCF data align closely to the CPS levels.

14. 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.

15. SCF data are at an annual frequency. For comparability, we also aggregate the model to an annual frequency, explaining why wealth-to-income ratios are positive at age 20 despite zero initial asset endowments for labor market entrants in the model.

Figure 1.6. Transition rates and tenure

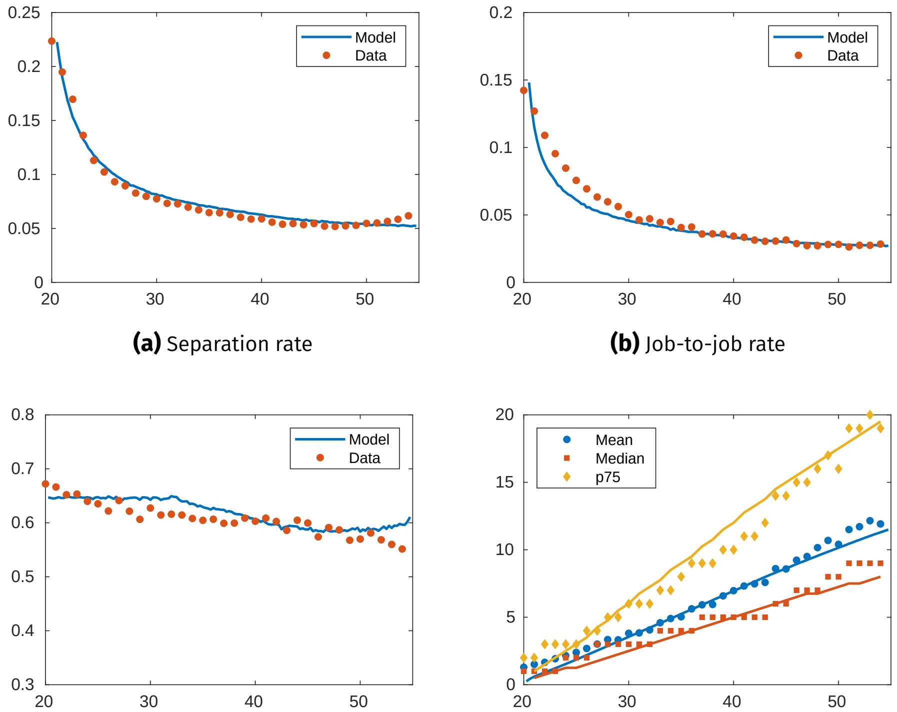

(a) Separation rate

(b) Job-to-job rate

(c) Job-finding rate

(d) Tenure (life cycle)

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 and tenure profiles are computed using data from the CPS.

lar, the U-shaped relationship between tenure and wealth-to-income ratios shows up clearly. The model provides an intuitive explanation for this pattern: workers build up wealth during their employment spell, and upon becoming nonemployed, their tenure and income drop but wealth remains constant and offers a buffer stock to smooth consumption after the job loss and during the recovery phase. This pattern will be more pronounced for more stable jobs as they have higher wealth-to-income ratios during employment and larger income drops. We will return to this heterogeneity in detail in the next section.

Figure 1.7d shows that the model also accounts for the negative relationship between wealth-to-income ratios and number of employers, demonstrating that job loss and job-to-job dynamics and their relationship to wealth accumulation are consistently accounted for by the model mechanism and align quantitatively with the data. By contrast, we show in Appendix 1.A.9 that a model without heterogeneity in separation rates struggles to correctly account for the observed relationship between job stability and wealth accumulation.

Figure 1.7. Earnings and wealth

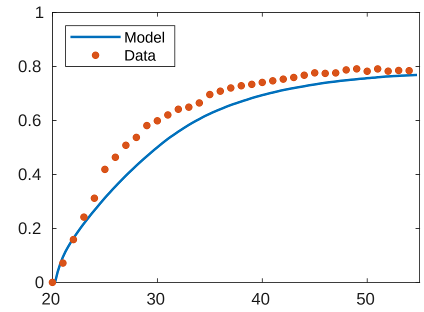

(a) Earnings

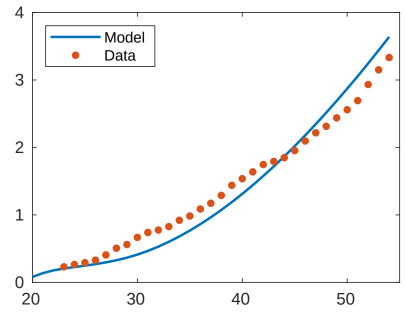

(b) Wealth-to-income ratio

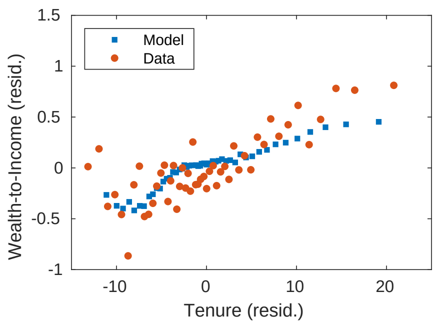

(c) Wealth-to-income ratio and tenure

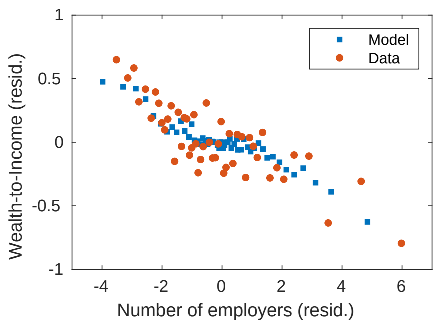

(d) Wealth-to-income ratio and number of employers

Notes: Panel (a) shows the mean of log earnings, normalized to 0 at age 20. Panel (b) shows the mean wealth-to-income ratio, calculated as the end-of-year assets divided by yearly income. Panel (c) shows the relationship between wealth-to-income ratios and tenure, and panel (d) the relationship between wealth-to-income ratios and the number of employers. In both of these cases, we nonparametrically control for age. In all panels, the blue lines/squares are the model profiles, while the red dots show the estimated empirical profiles. In panels (c) and (d), points represent binned scatter plots of wealth-to-income ratios against tenure/number of employers.

After looking at averages, in the next step we explore the model's ability to account for the life-cycle pattern of earnings, consumption, and employment inequality (Figure 1.8). Figure 1.8a shows the life-cycle increase in log earnings variance and the close match between model and data. We see the typical almost linear increase in the variance by age in both the model and the data. Labor market search models oftentimes struggle to account for this increase, as discussed in Lise (2012), Jung and Kuhn (2018), or Hubmer (2018). Augmenting models of job search with differences in human capital accumulation provides one way to account for the observed increase (Hubmer, 2018; Jung and Kuhn, 2018). While we build on this approach, we further refine it by endogenizing the human capital accumulation decision. As we will discuss in the next section, 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 1.8. Earnings, consumption, wealth, and employment inequality

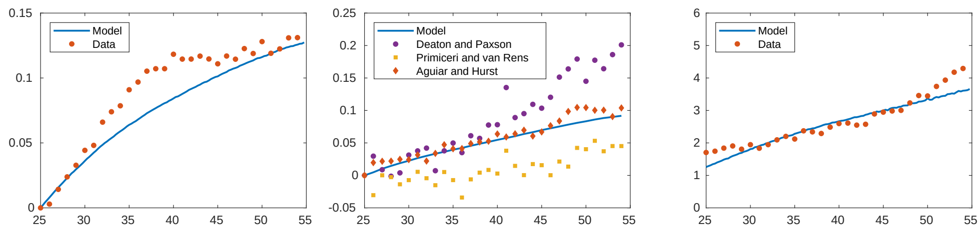

(a) Earnings variance **(b)** Consumption variance **(c)** Employment inequality

Notes: Panel (a) shows the variance of log earnings in CPS data from Jung and Kuhn (2018), normalized to 0 at age 25. 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. Panel (c) shows employment inequality measured as the ratio of empirically observed mean tenure and expected tenure. See Section 1.2 for details.

Figure 1.8b demonstrates that the model also aligns with the empirical estimates of the life cycle increase in the variance of log consumption. Empirical estimates of consumption variance differ across studies (Deaton and Paxson, 1994; Primiceri and Van Rens, 2009; Aguiar and Hurst, 2013), and the model falls in the middle of the range of existing estimates. Comparing Figures 1.8a and 1.8b, we also note that the increase in the variance of consumption is roughly one-third lower than that for earnings. Hence, consumption is partly insulated from earnings dynamics. Finally, we consider employment inequality in Figure 1.8c. We use the measure for employment inequality introduced in the empirical analysis of Section 1.2. We find that the model matches its empirical counterpart in its level and linear increase with age. This close match is a direct consequence of the model's ability to match the average life-cycle profiles of transition rates and the dispersion of tenure distribution by age.

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.10, 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 (2010a). 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, Karahan, Ozkan, and Song (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 consistency of the model with these facts lends support to the calibration of the human capital and wage processes as the two dimensions underlying life-cycle earnings dynamics in the model.

Finally, 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.11 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. We also directly compare the wealth dynamics over the life cycle to wealth panel data from the Panel Study of Income Dynamics (PSID). We find that the model closely matches individual wealth dynamics, lending further support to the economic mechanisms underlying our model framework, its endogenous earnings dynamics, and wealth accumulation decisions.

1.4 Individual consequences of job stability heterogeneity

Using our model framework, this section explores the consequences of differences in job stability on life-cycle earnings, consumption, and wealth dynamics and quantifies the welfare effects of a bad start to the labor market (i.e., the consequences of starting working life in an unstable job).

For our analysis, we will construct counterfactuals for identical workers who differ only in job stability, thereby isolating the causal effect of job stability on economic outcomes. Throughout, we will refer to a job as *stable* if it is at the top quartile of the observed job stability distribution for a worker of that age and as *unstable* if it is at the bottom quartile of the age-specific job stability distribution. We will consider 25-year-old agents as labor market entrants to accommodate the fact that workers start at age 20 ex ante identical and nonemployed to the model's labor market.[16]

1.4.1 Job stability heterogeneity and consumption-saving behavior

Agents in the model accumulate wealth for precautionary and life-cycle reasons. Job stability is a key determinant for how much precautionary savings workers want to accumulate as buffer stock to smooth out income fluctuations. The decision on how much precautionary savings to accumulate is particularly relevant for young workers who have not yet accumulated life-cycle wealth and, typically, have less stable jobs (Michelacci and Ruffo, 2015). The desire of young workers in unstable jobs to accumulate wealth for precautionary reasons is, however, counteracted by

16. For welfare results in Section 1.4.3, we verify that the results remain qualitatively and quantitatively largely unchanged if we consider a 20-year-old worker but use the age distribution at age 25 to define stable and unstable jobs at age 20. We show results in Appendix 1.A.12.

the motive for life-cycle consumption smoothing because young workers expect an increasing earnings profile over the life cycle (Modigliani and Brumberg, 1954). The typical low job stability and low earnings of young workers therefore create a direct tension between the precautionary savings motive and the life-cycle savings motive. By contrast, being a young worker with a stable job resolves some of this tension because a stable job implies fewer earnings fluctuations and requires less precautionary savings when income is low from a life-cycle perspective.

To explore the trade-off and to determine the importance of precautionary savings at different stages of the life cycle, we compare the baseline model to a counterfactual model without the risk of job loss. In this counterfactual model, the probability of entering into nonemployment is zero (i.e., we provide full insurance against nonemployment risk). Figure 1.9a compares life-cycle wealth accumulation between this counterfactual model and our baseline model. To control for the mean income difference between models, we compare wealth-to-income ratios and report the average wealth-to-income ratio of the full-insurance model (no displacement risk) relative to the average wealth-to-income ratio of the baseline model. Numbers below one imply that the wealth-to-income ratio in the full-insurance model is lower than in the baseline model.

Figure 1.9. Job stability, precautionary savings, and consumption growth

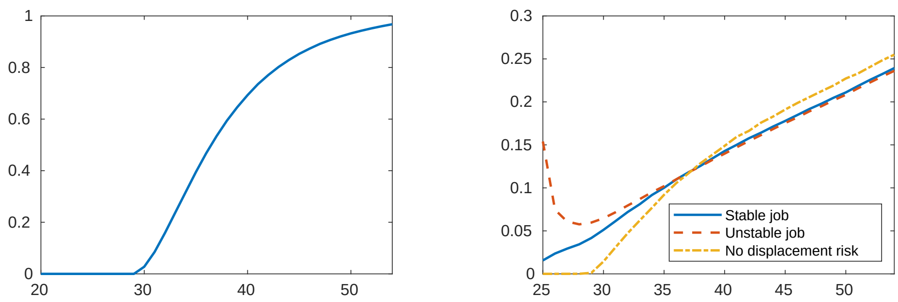

(a) Precautionary savings **(b)** Savings rate

Notes: The left panel shows the amount of wealth that agents accumulate in the case without displacement risk relative to the baseline wealth accumulation. Wealth is in both cases normalized by income to correct for differences in income growth. Panel (b) shows the relationship between the savings rate and job stability for 25-year-old workers in the baseline economy. Workers across the two groups differ only in terms of the separation rate. The stable job corresponds to the 25th percentile of the cross-sectional distribution of separation rates at age 25, and the unstable job to the 75th percentile. Workers have no wealth at the beginning of the experiment. The case with no displacement risk shows the savings rate of workers without displacement risk.

The first striking observation from Figure 1.9a is that with full insurance there is no wealth accumulation up to age 30. Only starting from age 30, workers accumulate wealth for life-cycle reasons, and the gap between the baseline and the full-insurance model closes quickly. The gap in wealth accumulation narrows to less than 0.07 at around age 50. This means that wealth-to-income ratios are just 7 percent lower in the full-insurance model at that point of the life cycle. Figure 1.7b shows that the wealth-to-income ratio in the baseline model is roughly 3 at age 50 so that about 1.3 months of income are held at that age for precautionary reasons. By contrast, all wealth before age 30 is precautionary so that without job-loss risk, agents would not accumulate any wealth. For the typical worker in the U.S. economy who is roughly 40 years old, the value of 0.7 in Figure 1.9a implies that roughly one out of three dollars of wealth is held for precautionary reasons. Between the ages of 30 and 40, life-cycle savings gain strongly in importance, so that the gap between the baseline model and the counterfactual model closes from 100 percent precautionary wealth to approximately 30 percent precautionary wealth. These findings align closely with the savings pattern described in Gourinchas and Parker (2002) and Cagetti (2003) where the precautionary savings motive governs household saving behavior in the first part of the life cycle before life-cycle savings become the dominant savings motive.

To further explore how differences in job stability impair workers' ability to smooth consumption over the life cycle, we look in Figure 1.9b at life-cycle profiles of saving rates for workers starting working life from a stable or unstable job at age 25 but who are otherwise identical. For both workers, we show the share of income that is saved. As a reference, we also show the savings rate for workers without displacement risk. At age 25, workers in unstable jobs save 15 percent of income, whereas agents in stable jobs save only 2 percent. While there are still precautionary savings by agents in stable jobs, the amount is strongly mitigated, and saving rates are much closer to the no displacement risk counterfactual. The additional savings do not translate into corresponding wealth differences as their purpose is to smooth out fluctuating income stemming from job instability. Agents starting from unstable jobs are in a kind of "Sisyphus cycle" of precautionary savings where they cycle repeatedly between nonemployment and unstable employment, building up and running down a buffer stock of savings. When workers find more stable jobs over time, differences in saving rates shrink and disappear in the mid-30s when all agents have started accumulating wealth for life-cycle reasons. Hence, starting working life from an unstable job substantially impairs the opportunity for life-cycle consumption smoothing.

Figure 1.10 turns from wealth accumulation to the life-cycle consequences of a bad start to the labor market on income and consumption. We use the model to construct a counterfactual experiment, where we consider one group of workers to whom we assign a bad start to the labor market (i.e., having an unstable job at age 25) and compare this group to an identical group of workers, but we put these workers in a stable job at age 25. State variables except for job stability are identical, so that income, wage, and human capital are also identical by construction at age

25. We index all life-cycle profiles relative to the level of the variable at age 25 of the worker in the stable job and consider log deviations.[17]

Figure 1.10. Consequences of differences in job stability on income and consumption

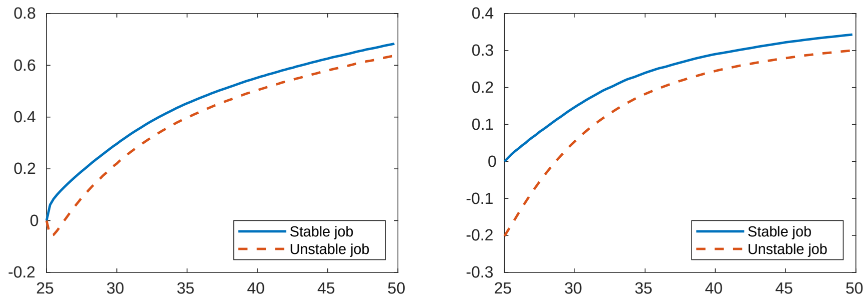

(a) Income **(b)** Consumption

Notes: This figure shows life-cycle profiles for income and consumption for 25-year-old workers with stable and unstable jobs. Workers across both groups initially differ only in terms of their separation rate. The stable job corresponds to the 25th percentile of the cross-sectional distribution of separation rates at age 25, and the unstable job to the 75th percentile. Profiles are normalized by the value of the profile for the stable job in the initial period and expressed in log deviations.

Looking at income in Figure 1.10a, we see that incomes are diverging immediately and already differ substantially after one year. Income combines earnings for the employed and benefit income for nonemployed workers. Differences in employment rates therefore account for part of observed income differences and, as discussed below, are on impact the key driver of the income divergence. Over time, we find not only that the bad start to the labor market leads to quickly diverging incomes but also that the decline remains persistent for the remainder of working life. At age 40, workers who had a bad start still have 5 percent lower incomes compared to the workers who started from a stable job, and this income gap also remains persistent for the following 10 years.

The different income prospects show up directly in very different consumption paths in Figure 1.10b. The key difference between income and consumption paths is that consumption jumps down immediately at age 25 as it is not a predetermined state variable. Starting from a 20 percent gap at age 25, consumption initially converges more strongly than income, so that at age 30, the consumption gap is at 10 percent. Around age 40, the size of the consumption gap converges to the size of the income gap and then remains persistent for the rest of working life. Two drivers account for these consumption dynamics. First, there is a standard level effect as income will be permanently lower starting from the less stable job (Figure 1.10a).

17. For employment rates in Figure 1.11a, we show percentage point differences relative to the average 25-year-old worker.

Second, as employment is less stable and income is more volatile after a bad start to the labor market, agents will engage more in buffer stock saving behavior, which temporarily lowers consumption. When jobs become more stable over the working life, differences in precautionary savings vanish and consumption differences converge to the permanent income differences. Taken together, these consumption dynamics align closely with the idea of the permanent income hypothesis as here workers stabilize consumption relative to more volatile incomes around a permanently lower level (Friedman, 1957). Over the life cycle, accumulated precautionary savings from the first part of working life substitute later on for part of life-cycle savings and allow for some additional convergence of consumption profiles.

This counterfactual model simulation allows us to isolate the causal effect of differences in job stability on income and consumption dynamics. The underlying reason why job stability translates into different consumption dynamics stems from the endogenous labor market dynamics that intertwine average income growth and income volatility after a bad start to the labor market. In this way, a bad start to the labor market leaves long-lasting scars on both life-cycle income and consumption dynamics.

Figure 1.11 decomposes income dynamics into an effect from employment rates, human capital, and wages. Figure 1.11a considers employment rate differences expressed relative to the average 25-year-old worker. Relative to the average worker, workers in stable jobs have higher employment rates so their employment rate initially jumps up by about 10 percent. Over time, the employment rate difference to the average at age 25 is declining, but it stays higher and converges to a 7 percent higher employment rate at age 50. For workers in unstable jobs, the situation is very different. Their employment rates plummet on impact because of low job stability. Afterward, employment recovers only slowly. Only at around age 30, workers starting from an unstable job reach the employment levels of the average 25-year-old worker, and at age 50, employment rates have almost converged between the workers starting from stable and unstable jobs. The fact that employment rates for the unstable job show a positive 6 percentage point level at age 50 highlights that jobs become on average more stable during working life. The comparison between workers starting from stable and unstable jobs highlights that a bad start to the labor market leads to persistently depressed employment for the rest of working life. Put simply, having an unstable job when young implies being less likely to have a job when old. The large initial employment difference is also the reason behind the strong divergence of incomes in Figure 1.10a. The reason for the smaller level of the income gap compared to the employment gap is that nonemployed workers receive benefits so that the income gap is roughly cut in half compared to the employment gap.

Figures 1.11b and 1.11c look only at earnings as the incomes of employed workers and decompose earnings differences between workers starting from stable and unstable jobs into a wage and human capital component. The key observation re-

Figure 1.11. Decomposition income effects from differences in job stability

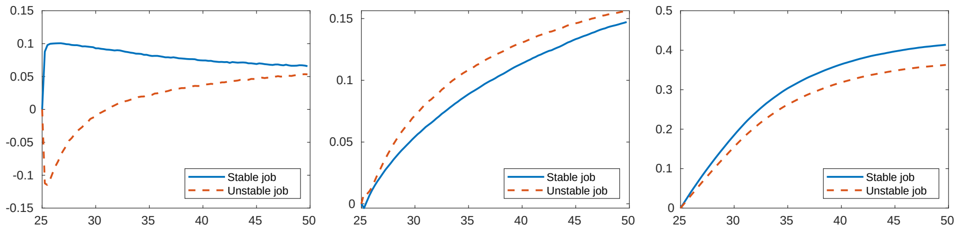

(a) Employment rates **(b)** Wage **(c)** Human capital

Notes: This figure shows life-cycle profiles for employment rates, wages, and human capital for 25-year-old workers with stable and unstable jobs. Workers across both groups initially differ only in terms of their separation rate. The stable job corresponds to the 25th percentile of the cross-sectional distribution of separation rates at age 25, and the unstable job to the 75th percentile. Employment rates are normalized by the average employment rate at age 25 and shown as percentage point difference. Human capital and wage profiles are normalized by the value of the profile for the stable job in the initial period and expressed in log deviations.

garding wages in Figure 1.11b is that they increase more for workers starting from an unstable job (dashed red line). The reason is that starting working life from an unstable job implies that workers will quickly fall down the wage ladder (Figure 1.11a) and will have to start climbing it again. Climbing the wage ladder from scratch leads to additional wage growth compared to a situation in which a worker is already in a stable job because the worker in the stable job has to trade off new wage offers against high job stability of the current job, and this trade-off leads to less wage growth since some better-paying job offers, being less stable, will be turned down. Quantitatively, the resulting wage difference remains small, however, and amounts to only around 1 percent at age 50.

The effect from differences in human capital accumulation in Figure 1.11c is substantially larger and amounts to almost 5 percent at the end of working life. Differences in job stability account for the diverging human capital trends. A worker starting from an unstable job enters on a less stable employment path and will spend less time in employment (lower employment rates) so that she has fewer opportunities to invest in human capital, especially when young and when human capital investment is most productive (Figure 1.11a). The human capital accumulation process therefore provides a propagation mechanism for persistent scars from short-run search outcomes. Differences in initial job stability that are the result of search frictions account for the persistent earnings differences in our model, so the resulting earnings differences should be subsumed under a broader notion of (dynamic) frictional wage dispersion.[18]

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

1.4.2 Job stability heterogeneity and the consequences of job loss

To provide further intuition for the consequences of job stability on life-cycle dynamics, this section explores the consequences of job loss and their relationship to heterogeneity in job stability. To do so, we adapt the approach from the empirical literature on job displacement (Jacobson, LaLonde, and Sullivan, 1993a) 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.[19] Empirical studies document that such job displacements lead to large and persistent earnings losses for workers (Jacobson, LaLonde, and Sullivan, 1993a; Couch and Placzek, 2010; Davis and Wachter, 2011), and heterogeneity in job 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 job stability struggle to account for the persistence in earnings losses (Low, Meghir, and Pistaferri, 2010), a fact we also highlight in Appendix 1.A.9 where we show earnings losses for a model without job stability heterogeneity and how the resulting earnings losses are only transitory.

Figure 1.12a shows that our baseline model with heterogeneity in job stability implies large and persistent earnings losses for the average 40-year-old worker. The initial earnings drop in displaced workers amounts to around 15 percent. Over the subsequent five years, displaced workers are able to cut initial earnings losses by half, but there is little further catch-up. The figure also shows the evolution of the components of earnings (wage and human capital) to uncover the underlying mechanism of the persistent earnings loss. All of the initial loss in earnings comes from the fact than 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 account for 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

19. This approach differs from the empirical approach that conditions on pre-displacement tenure and post-displacement job stability. In the model, we exploit 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 (1993a)).

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 1.12. Cost of displacement

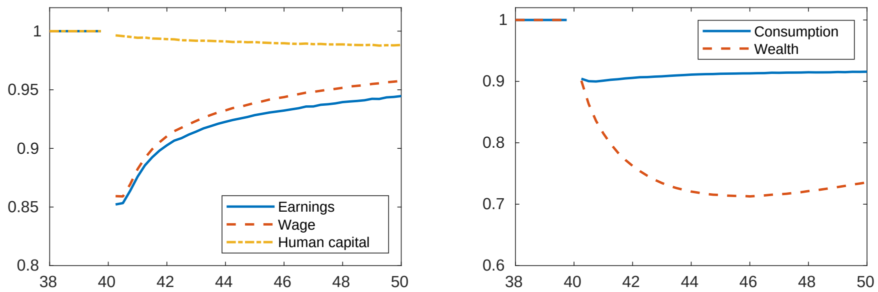

(a) Earnings and Wages

(b) Consumption and Wealth

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

Figure 1.12b turns to the consequences of job loss for consumption and wealth. Looking at consumption, we see a sharp (roughly 10 percent) drop in consumption directly on impact. After the onetime 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 because of lower earnings and lower employment rates, and income will be more volatile because of lower job stability. As a consequence of lower permanent income and higher volatility, agents permanently reduce consumption and increase their precautionary savings to smooth consumption in the future. On impact, consumption drops less than income, and this difference shows up directly in wealth dynamics as wealth is used to smooth the transition to the new, lower permanent income level (Kuhn, 2013). Four years after the job loss, wealth levels stabilize 25 percent below the level of nondisplaced workers and remain persistently lower in line with recent empirical results in Barnette (2020). Three reasons account for this lower average wealth after job loss. First, income is lower so that the wealth level adjusts, too. Second, the job loss has flattened the life-cycle income profile. Current income is now lower relative to income during the rest of the life cycle, which reduces the need for life-cycle savings that aim at reshuffling age-varying income over time. Third, the lower job stability after the job loss sets agents on the kind of Sisyphus saving cycle that we have already described for young workers. 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 tenure implied by the model and observed in the data (Figure 1.7). Workers who lose their jobs experience significant decreases in income, and at the same time, their tenure drops to zero. During the transition, wealth declines, earnings recover so that wealth-to-income ratios fall, and we get the nonmonotonic relationship between tenure and wealth-to-income ratios, as observed in Figure 1.7.

Figure 1.13 explores as the next step how differences in job stability shape the consequences of job loss. It shows the results of the previous displacement experiment but compares workers displaced from initially stable and initially unstable jobs. Except for job stability, workers are again otherwise identical.

Figure 1.13. Effects of displacement by job stability

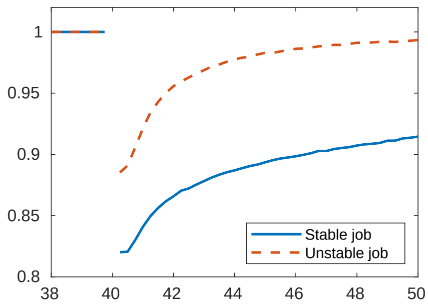

(a) Earnings

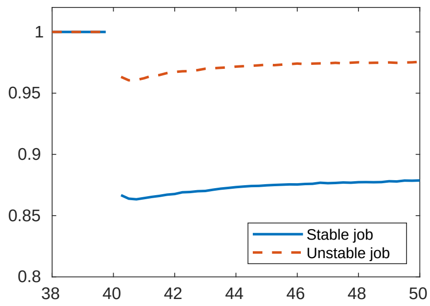

(b) Consumption

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

On impact, losing the stable or unstable job leads to earnings losses of 12 percent and 18 percent, respectively (Figure 1.13a). While these initial earnings losses are similar, the recovery from the initial shock is strikingly different between the stable and unstable job. 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 job, 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 jobs 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 jobs exhibit 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 job stability are only transitory (Appendix 1.A.9). By contrast, the consequences of job loss are strikingly different for workers who lose an initially stable job. Now the same logic applies but with different consequences. If workers in initially stable jobs had not lost their job, the high job stability would imply that they would have been unlikely to lose their job in the future. Hence, high job 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 job 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 1.13b). Workers who lose their stable job 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 percent. 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 lose a stable job as employment rates starting from a stable job 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 percent on average. By contrast, if all job losses were in stable jobs, then the consumption drop would be 13 percent —more than three times as large.

1.4.3 Welfare consequences of heterogeneity in job stability

What are the welfare consequences of a bad start to the labor market? In most search models, bad luck in the search process washes out quickly as workers keep on searching for better opportunities, and those workers who have been lucky at the beginning of the search process enjoy their search outcomes only for a short time as high average separation rates imply a lot of mean reversion (Hornstein, Krusell, and Violante, 2011). Heterogeneity in job stability perpetuates search outcomes, leading to potentially large welfare costs from a bad start to working life. We explore these welfare consequences using our model framework in which we conduct the following counterfactual experiment. We consider workers at age 25 and ask how much consumption they would be willing to give up to keep their current job in terms of separation rate and wage instead of getting the same job but with the lowest level of job stability. This means we keep the job's wage constant and vary the job's separation rate only when determining welfare costs. We derive the consumption-equivalent variation in the differences in job stability considering three wage levels and all possible levels of separation rates from the support of the calibrated job-offer distribution.

Figure 1.14a reports the consumption-equivalent variation for an agent with median wealth and median human capital across all workers of age 25. The dashed vertical line shows the average separation rate for workers of that age, and the three lines show the welfare costs at different wage levels. We see immediately that welfare effects from differences in job stability can be substantial and can exceed 10 percent of lifetime consumption for workers with a high-wage, stable job (dashed yellow line). For a 25-year-old worker with median wage and average job stability, the welfare costs of being moved to the least-stable job are substantially smaller but are, with 1.4 percent of lifetime consumption, still large. By construction, welfare losses at all wage levels disappear the closer we move to the least-stable job. At the other end, welfare losses grow strongly when jobs become more stable. The welfare differences across wage levels are large at the average job, and the differences increase further the more stable jobs become. The reason for the strongly growing welfare costs is that stability of jobs is particularly valuable if wages are high. The more transitory jobs become (i.e., the lower job stability is), the shorter the time period during which workers expect to enjoy a high wage.

To explore the contribution of human capital accumulation, incomplete financial markets, and labor market frictions to these welfare costs, we decompose the welfare effect from Figure 1.14a into a *human capital component*, an *insurance component*, and a *search component*. First, to isolate the effect from human capital investment, we set all workers to the highest level of human capital h^* and then conduct the same comparative statics experiment as for the baseline welfare effect. That is, instead of the average human capital level, we evaluate welfare effects at the highest human capital level. In this case, differences in job stability do not impair workers' ability to accumulate human capital. Subtracting the welfare effects of this experiment from the baseline isolates what we refer to as the human capital component.[20] In the next step, we additionally endow the worker who already has high human capital with high wealth.[21] For this worker, we again conduct the comparative statics experiment of setting her at the least-stable job. Now, financial market incompleteness does not

20. Specifically, let us denote by Δ_b the consumption-equivalent variation for the baseline experiment and by Δ_h the consumption-equivalent variation in the case of high human capital. Then the human capital component is $\Delta_b - \Delta_h$, constituting a difference-in-difference construction.

21. We define high wealth as the highest wealth level attained at the end of working life during our simulations of the baseline model.

Figure 1.14. Welfare costs of job instability for different types of jobs

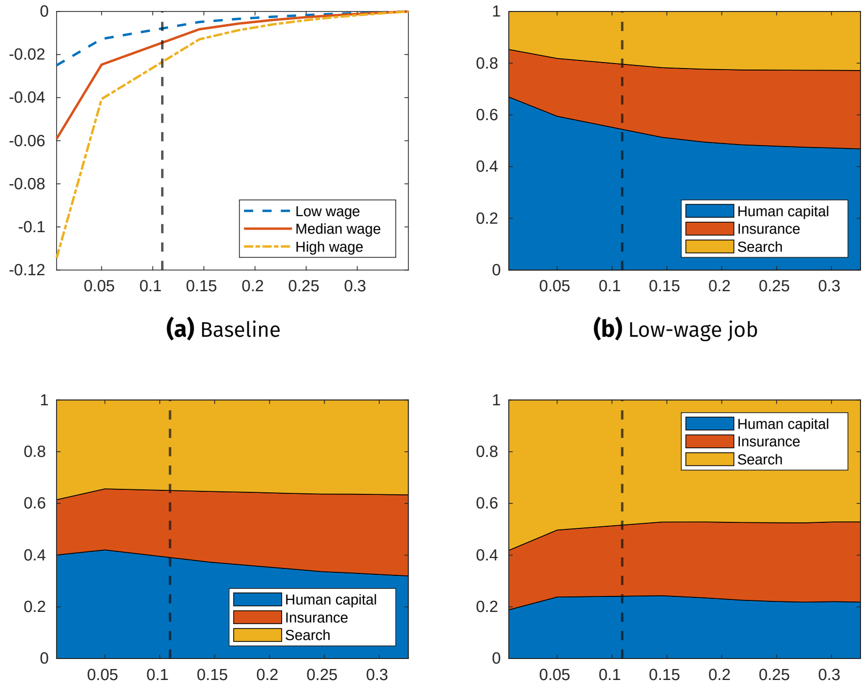

(a) Baseline

(b) Low-wage job

(c) Median-wage job

(d) High-wage job

Notes: Panel (a) shows the welfare costs of job instability as the consumption-equivalent variation for 25-year-old workers in low-, median-, and high-wage jobs for all levels of separation rates λ (on the horizontal axis). Welfare costs are for moving a worker to the least-stable job with the same wage. Panels (b)-(d) show the decomposition of the welfare cost into a human capital component, an insurance component, and a search component. Medium wage corresponds to the median wage at age 25, low wage corresponds to the 25th percentile of wages at age 25, and high wage corresponds to the 75th percentile of wages at age 25. The dashed vertical line shows the average separation rate at age 25. Welfare is evaluated at the levels of median wealth and median human capital at age 25.

impair the consumption-smoothing ability of the agent because of the large buffer stock endowment. By subtracting this welfare effect from the one with high human capital, we isolate the *insurance component*. Finally, we construct the *search component* by subtracting the human capital component and the insurance component from the baseline. The search component captures the welfare effect of moving to the least stable job, so that the current wage level becomes more transitory and workers have to restart their job search sooner.

Figures 1.14b to 1.14d show the decomposition of the welfare effect from Figure 1.14a into the three components at the three different wage levels.[22] We find

22. We construct the search component as the residual so that the three decomposition components always sum to 100 percent of the total effect.

that the insurance effect accounts for roughly 20 percent of the total welfare costs across all wage and job stability levels. The components that vary across wage levels are the relative importance of the human capital and search component. Intuitively, we find that the search component is more important the higher the wage level of the current job is. As explained above, welfare losses are larger at a high wage because moving the worker to a less stable job makes the current high wage level more transitory. For low-wage jobs, the human capital component is most important. For workers in low-wage jobs, high job stability is valuable because it offers them the opportunity to invest in human capital. This effect is slightly nonlinear and accounts for two-thirds of the welfare effect at the most stable but lowest wage jobs (Figure 1.14b). At the median wage (Figure 1.14c), the decomposition into the three components is roughly constant across job stability, with roughly 40 percent for the human capital and search component. We conclude that a bad start to the labor market can be very costly and that even low-paying but stable jobs (e.g., apprenticeships) can be very valuable for labor market entrants as they offer human capital investment opportunities.[23]

1.5 Consequences of the aggregate decline in U.S. labor market dynamism

In this section, we turn to the macroeconomy to explore the consequences of the secular decline in U.S. labor market dynamism (Fallick and Fleischman, 2004; Davis, 2008; Fujita, 2018; Molloy, Smith, and Wozniak, 2020). Declining labor market dynamism captures the widely documented decline in labor market mobility. For example, when focusing on separation rates, Fujita (2018) reports that transition rates from employment into unemployment declined by roughly a quarter, from around 1.7 percent per month in 1976 to around 1.3 percent per month by 2008. For monthly job-to-job transitions, Fallick and Fleischman (2004) report a decline of 0.6 percentage points between 1994 and 2004, starting from around 2.8 percent per month.[24] At the same time, it has been pointed out that a lot of the decline in average transition rates resulted from the disappearance of very short-term jobs (Molloy, Smith, and Wozniak, 2020). Despite the broad consensus regarding declining dynamism in the U.S. labor market, little is still known about its consequences. We will interpret the stylized facts on declining dynamism as originating from a change

23. We repeat the same welfare analysis in Appendix 1.A.12 for a 20-year-old worker where we apply the distribution across jobs at age 25. Results are very similar, but the relative importance of the human capital component in the welfare decomposition increases.

24. Fujita, Moscarini, and Postel-Vinay (2020) recently report adjusted time series for job-to-job transitions that decline by less than job-to-job transitions, following the approach in Fallick and Fleischman (2004). Still, they find a decline of 0.5 percentage points for the longer time period from 1995 to 2020.

in the macroeconomic environment and explore its consequences for workers' careers and provide estimates of associated welfare effects for labor market entrants. Ex ante, it remains ambiguous whether such a secular decline in labor market mobility is welfare increasing or decreasing. On the one hand, lower separation rates are welfare increasing because more frequent job losses are costly, as demonstrated in the last section. On the other hand, lower job-to-job rates reduce welfare because they make the job ladder harder to climb and negatively affect earnings growth over the course of a worker's career.

We rely on a steady-state comparison of two economies that differ in their macroeconomic labor market environment. We compare the results of our baseline economy to an economy that has on average a 1 percentage point lower quarterly separation rate and a 1 percentage point lower quarterly job-to-job transition rate but both economies have the same tenure distribution, following the evidence in Molloy, Smith, and Wozniak (2020). The decline in the separation rates and job-to-job rates follows the estimated declines of monthly rates in Fujita (2018) and Fujita, Moscarini, and Postel-Vinay (2020). We calibrate the model for the new economy by reducing the job offer rate for employed workers by 22 percent to match the empirical decline in job-to-job mobility. To match the decline in separation rates together with a constant tenure distribution, we compress the lower part of the support of job separation rates. Specifically, we decrease the separation rate on the least stable job to $\bar{\lambda} = 0.3$ and adjust the remaining grid points so that average separation rates in the offer distribution decline by 15 percent, and we match the decline in the data.[25] Figure 1.15 shows the life-cycle profiles for separation rates, job-to-job rates, and the mean, median, and 75th percentile of the tenure distribution for the baseline model and the model with a less dynamic labor market ("low mobility").

For the welfare consequences, we compare welfare of labor market entrants across the two economies. It is important to keep in mind that labor market entrants at age 20 start working life as nonemployed with zero assets and the lowest level of human capital, so that they are identical ex ante before entering the two economies. We do not consider a transition phase between economies. Abstracting from the transition phase can be interpreted either as a change in the macroeconomic environment that took place immediately or as comparing two workers several years apart when the macroeconomic environment has changed.

Comparing workers across the two economies, we find that lower labor market mobility leads to a welfare gain. The positive effect of a lower risk of job loss outweighs the negative consequences of reduced on-the-job mobility. A labor market entrant in the economy with lower labor market mobility would be willing to give

25. We set the second grid point at $\lambda_2 = 0.019$ and the remaining grid points according to the same rule as for the baseline economy $\lambda_j = \underline{\lambda} + \left(\frac{j-1}{N_\lambda - 1}\right)^{0.6} \times (\overline{\lambda} - \underline{\lambda})$. See footnote 10 for details of the baseline economy. Figure 1.A.13 shows the marginal distributions of separation rates for the baseline economy and the less dynamic economy.

Figure 1.15. Comparison of labor market life-cycle profiles

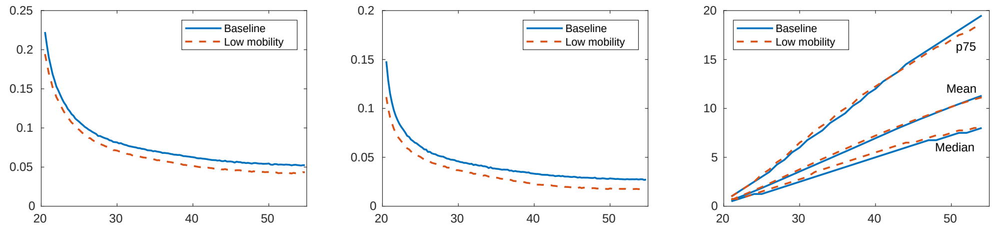

(a) Separation rate **(b)** Job-to-job rate **(c)** Tenure

Notes: This figure shows quarterly life-cycle transition rates and tenure in years by age for the baseline economy and an economy with lower average separation and job-to-job transition rates. Tenure distributions are targeted to be the same. The solid blue lines show the baseline model, and the dashed red lines show the model with lower labor market mobility.

up 1.6 percent of lifetime consumption to avoid entering the dynamic labor market of the baseline economy. After entering the labor market, the magnitude of the welfare gain is decreasing in job stability. A worker employed at age 25 in a stable job would be willing to give up 0.6 percent of the remaining lifetime consumption, and a worker employed in an unstable job would be willing to give up 1.2 percent of lifetime consumption to avoid continuing his or her career in the baseline economy.[26] These findings reflect the fact that for workers in unstable jobs, a significant portion of employment risk has been eliminated in the less dynamic economy. Since the increase in job stability is constructed such that it predominantly takes place in the unstable part of the job-offer distribution, the benefits of transition to the new economy are higher for workers in unstable and low-wage jobs. To isolate the positive welfare effect from declining labor market dynamism, we consider an alternative counterfactual experiment where we reduce separation rates and match the stability of the tenure distribution but keep job-offer rates at their higher level.[27] In this case, we find much larger welfare gains for labor market entrants. The typical young American worker would be willing to give up 2.9 percent of lifetime consumption relative to the baseline economy to start working life in an economy with more stable jobs. The much larger welfare gain in this case indicates that declining job-to-job mobility had a substantial negative welfare effect for young workers, given a combined welfare effect of more job stability and lower job-to-job mobility of 1.6 percent, which is 1.3 percent lower than in the case of only declining separation rates.

26. We compare identical workers across economies in terms of their state variables and interpolate the value function in the job stability dimension to derive welfare effects.

27. If we only reduce separation rates as in the baseline experiment but keep job offer rates constant, then the realized reduction in separation rates will not match the aggregate decline because of changing worker search behavior.

Figure 1.16. Comparison of earnings, wage, and human capital dynamics

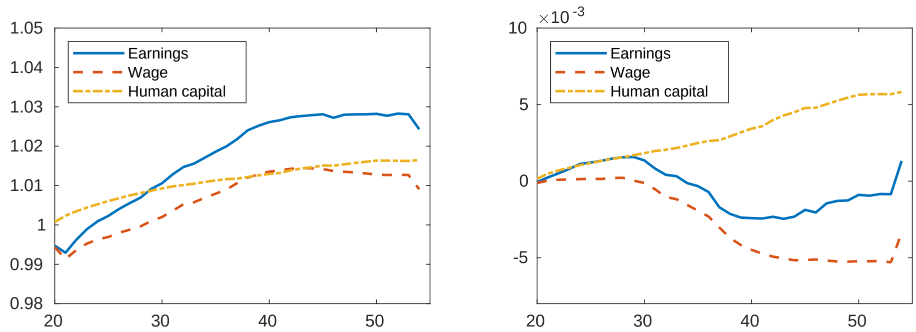

(a) Mean **(b)** Variance

Notes: Panel (a) shows the ratios of average profiles for earnings, wages, and human capital for the baseline economy and the less dynamic economy. Life-cycle profiles are indexed to the baseline economy so that values larger than one indicate higher values relative to the baseline economy. Panel (b) shows the absolute differences in the variance of log earnings, log wages, and log human capital between the baseline economy and the less dynamic economy. Positive values indicate higher variances than in the baseline economy.

Figure 1.16 compares life-cycle profiles for averages and variances of earnings, wages, and human capital between the two economies. To highlight differences in average profiles, Figure 1.16a shows the ratios of the average profiles where a number larger than one implies that the average is higher in the labor market with lower mobility.

For average earnings, we find that more stable careers translate into higher earnings growth. At the end of working life, earnings are on average almost 3 percent higher in the economy with lower labor market mobility. Looking at the decomposition into human capital and wages, we first observe that workers start with lower wages if job stability increases. With higher job stability, workers at the beginning of their career now more often accept low-wage but more stable jobs. Subsequently, they climb the wage ladder more quickly, and wage levels break even at around age 30. Despite fewer job offers, we get that higher job stability and a less slippery wage ladder result at the end of working life in average wages that are 1 percent higher. Fewer opportunities to climb the wage ladder are overcompensated by fewer falls off the wage ladder. Human capital coincides by construction at the beginning of working life between the two economies. In the less dynamic economy, we find human capital to increase more from the start of a worker's career. As discussed in the last section, higher job stability improves opportunities for human capital investment and increases earnings growth. At the end of working life, we find that human capital contributes two-thirds to the higher earnings growth in the less dynamic economy.

Looking at the differences in earnings inequality in Figure 1.16b, we find that differences are small overall and that the difference follows a wave with slightly higher inequality until the mid-30s and lower inequality during the decade between ages 40 and 50. At the end of working life, earnings inequality in the two economies ends up being roughly equal again. Looking at wage and human capital inequality differences, we find two counteracting effects on the life-cycle profile of earnings inequality. First, the additional risky human capital investment arising from higher job stability increases the variance of human capital and earnings. Second, more stable jobs alleviate climbing the wage ladder, which tends to reduce wage inequality as more and more workers end up in high-wage jobs. Lower wage inequality contributes to compressing earnings inequality. The human capital effect dominates during the first part of working life when most human capital investment takes place, and the wage effect dominates during the second part of working life when human capital investment has slowed down.

Lastly, we also find little evidence that lower labor market mobility increases or reduces the life-cycle consequences of heterogeneity in job stability (Section 1.4). When we repeat the experiment from Figure 1.10 in the economy with less labor market mobility, the degree of persistence of labor market outcomes remains in line with the findings for the baseline economy. Starting workers in unstable and stable jobs at age 25, we find that 20 years later, the incomes of the workers with the unstable job are again 5 percent lower than the incomes of workers with the stable job —the same as the 5 percent difference in the baseline economy (Figure 1.10). For the cost of displacement, we find that for the average 40-year-old worker, earnings losses in the short run are lower in the less dynamic economy but the difference in earnings losses shrinks and tends to disappear with time to the initial job loss. Five years after the job loss, the difference in earnings losses is only 0.4 percentage points.

To summarize, we find that lower labor market dynamism is welfare improving for young American workers as it offers better opportunities for human capital investment. Moreover, wage growth is positively affected as the steps of the wage ladder have become more stable.

1.6 Conclusions

Our analysis started from the observation of large heterogeneity of job stability in the U.S. labor market. Using SCF data, we demonstrate that differences in job 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 job stability for earnings, consumption, wealth, and welfare. At the macroeconomic level, we explore the consequences of the declining labor market dynamism on life-cycle dynamics and welfare. We find that a bad start to the labor market with an unstable job early in life leaves long-lasting scars on a worker's career. Job instability leads to less income growth, less consumption, and less wealth. The consequences of job stability for life-cycle dynamics stem from two sources. First, lower job stability leads to less income growth, mainly from less human capital accumulation. Second, lower job 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. The welfare losses from job instability are large and amount to 1.4 percent of lifetime consumption for a typical 25-year-old worker.

When we explore the consequences of the macroeconomic changes in job stability due to the secular decline in U.S. labor market dynamism, we find a net effect of declining separation rates and job-to-job mobility that is welfare improving for labor market entrants. In line with the empirical evidence on the tenure distribution, we model the decline in separation rates as asymmetric so that predominantly unstable jobs disappear. The shift toward less job instability allows for more investment in human capital and outweighs the negative effect of fewer dynamics on the wage ladder. We conclude that declining labor market dynamism has been welfare improving for young American workers.

1.A Appendix

1.A.1 Employment history and wealth: robustness checks

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.1 shows that our findings are not affected by the inclusion of additional controls. In the first column, we nonparametrically control for age, education, occupation, and industry. The relationship between tenure and wealth-to-income ratios (the top row) remains unaffected and significant. The same holds for the relationship between the number of employers and wealth-to-income ratios (the bottom row). In the second step, we additionally control for differences in risk attitudes of workers by nonparametrically controlling for different levels of risk attitudes, as elicited in the SCF survey. As shown in the second column, the relationship between labor market experience and wealth is not affected by the inclusion of all these additional variables.

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

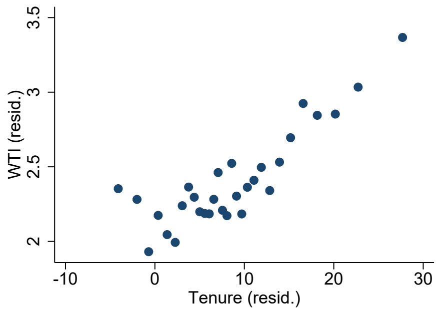

(a) Controls: age, education, occupation, industry

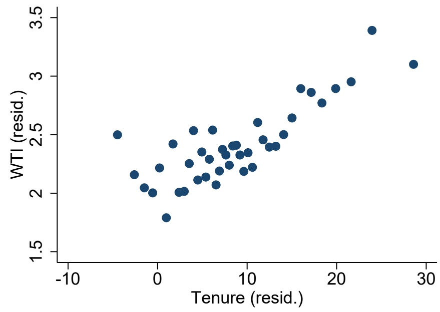

(b) Controls: age, education, occupation, industry, risk attitude

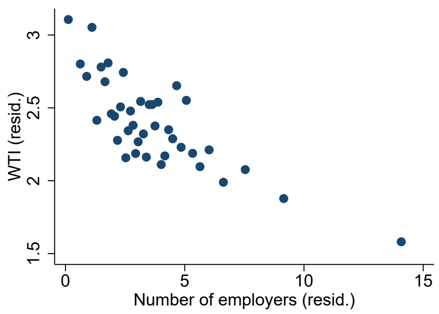

(c) Controls: age, education, occupation, industry

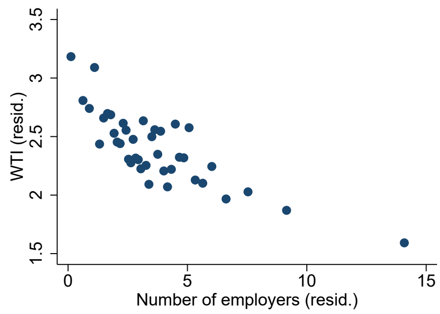

(d) Controls: age, education, occupation, industry, risk attitude

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.

1.A.2 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

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

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

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

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

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

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.3 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.2 shows that the heterogeneity in job destruction rates persists even after controlling for year and Metropolitan Statistical Area (MSA) fixed effects.

Figure 1.A.2. Heterogeneity in job destruction rates by firm age

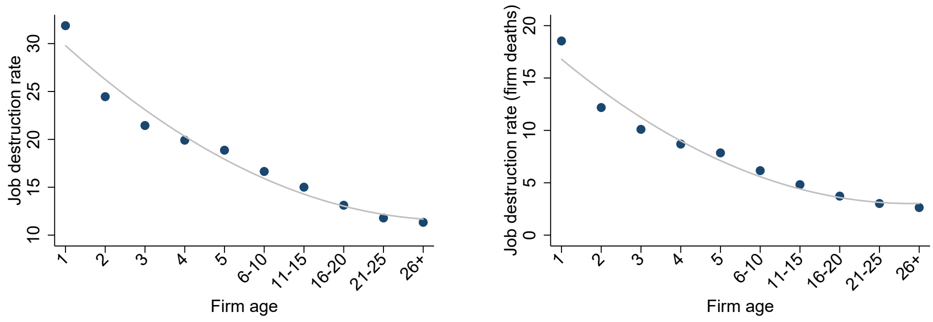

(a) Job destruction rate

(b) Job destruction rate (firm deaths)

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.4 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.3 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 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.3, the stayer type has a transition rate that is 25 percent lower than the average transition rate, whereas the mover type has a transition rate that is 25 percent 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 percent lower, and the mover type has a transition rate that is 90 percent higher than the average transition rate. The left panels show the resulting tenure distribution, and the right panels show the transition rates.

Figure 1.A.3. Worker types and tenure distribution

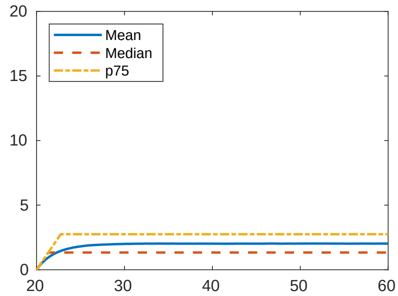

(a) Tenure (low type differences)

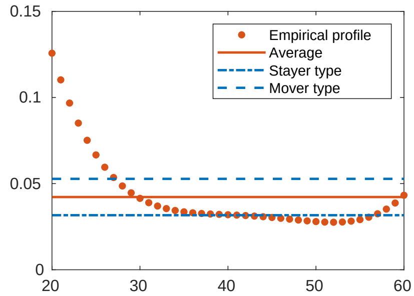

(b) Transition rates (low type differences)

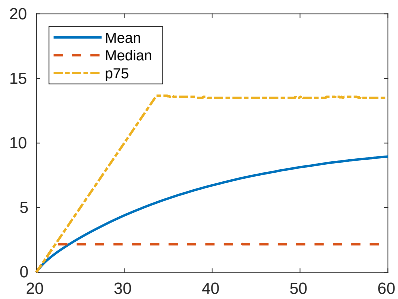

(c) Tenure (high type differences)

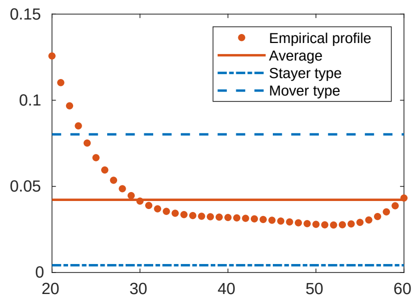

(d) Transition rates (high type differences)

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 percent lower than the average transition rate and workers of the mover type have a 25 percent 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 percent.

Compared to the empirical profiles shown in Figure 1.6, 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.5 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

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

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

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

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

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

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

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.6 Model solution and estimation

1.A.6.1 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 $[0, 1]$.[28] 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.

28. 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.6.2 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 by a hat.

1.A.6.3 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.6.4 The estimated job-offer distribution

Panel 1.A.4a of Figure 1.A.4 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.4b additionally shows the conditional distribution of separation rates for different levels of wages. The distribution of separation rates in low-wage jobs first-order stochastically dominates the distribution of separation rates in high-wage jobs.

Figure 1.A.4. Job-offer distribution

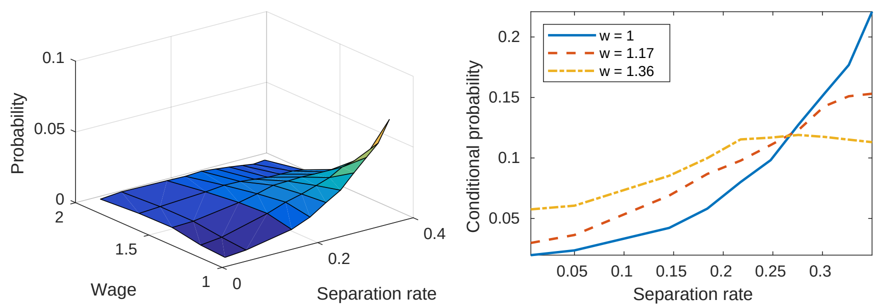

(a) Joint distribution **(b)** Conditional distributions

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.

1.A.7 Comparison of life-cycle profiles in SCF and CPS

Figure 1.A.5 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.5. Earnings and tenure in SCF and CPS

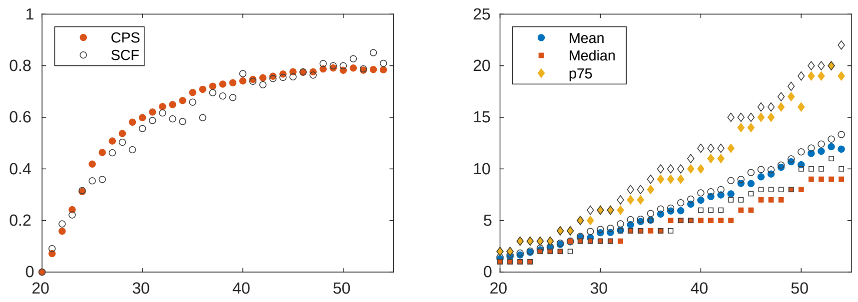

(a) Earnings **(b)** Tenure

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.8 Cross-sectional distributions of tenure and the number of employers

Figure 1.A.6 shows the cross-sectional distribution of tenure and the number of employers for which a worker has worked for at least one year during her working life. We combine all workers and show the corresponding distribution using histograms. When pooling data from the model, we assume that each age group has the same share in the pooled sample.

Figure 1.A.6. Cross-sectional distribution of tenure and number of employers

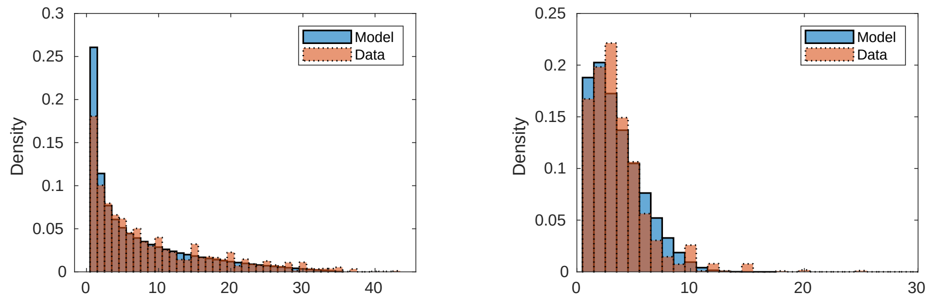

(a) Tenure distribution

(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 with a solid outline are the SCF data; blue bars with a dotted outline 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.

1.A.9 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 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.7a, 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.7b, 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.8, we show the cost of displacement for the model without

Figure 1.A.7. Tenure and wealth

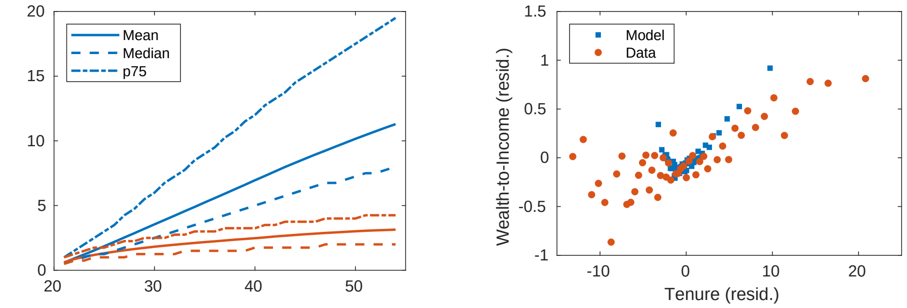

(a) Tenure

(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 lower red profiles are for the model without heterogeneity in job stability, and the upper blue profiles correspond to the baseline model. Panel (b) shows the relationship between wealth-to-income ratios and tenure after nonparametrically controlling for age.

cross-sectional heterogeneity in the separation rate. Contrary to the results from the baseline model in Figure 1.12, 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 compared to the baseline model. Job loss in the economy without heterogeneity is largely inconsequential, in contrast to the model with heterogeneity, as discussed in the main part of the paper.

Figure 1.A.8. Cost of displacement without heterogeneity in job stability

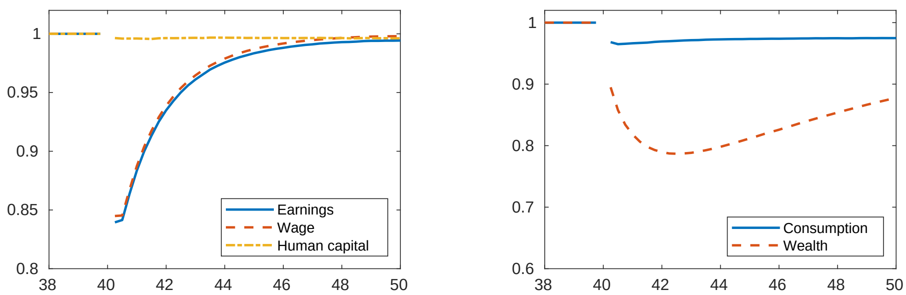

(a) Earnings

(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.10 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, 2010a). Second, as emphasized in Guvenen, Karahan, Ozkan, and Song (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.[29] 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, Karahan, Ozkan, and Song (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

29. Heathcote, Perri, and Violante (2010a) 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.

version of the job ladder model with a human capital process. The same also applies to our model. Figure 1.A.9 shows the distribution of one-year earnings growth rates in the model with a superimposed normal distribution that has the same standard deviation. Earnings changes in our model are left-skewed and strongly clustered around 0.

Figure 1.A.9. One-year earnings changes

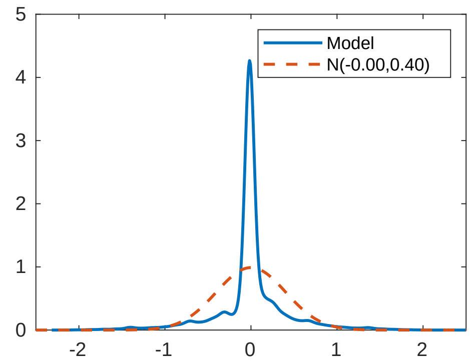

Notes: Figure shows the estimated kernel density of the model-based one-year earnings changes superimposed on Gaussian densities with the same standard deviation. Earnings growth net of average age effect shown. Kernel density estimation with a bandwidth of 0.05 used.

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.10 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.10a, 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 1.6). 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

Figure 1.A.10. Decomposing earnings dynamics over the life cycle

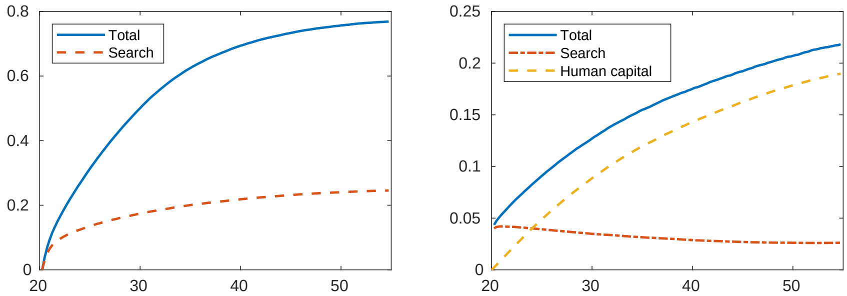

(a) Mean

(b) Variance

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

as many jobs offer lower wages, and therefore, these workers are more likely to stay with their current employer. As a result, initial wage differences decrease. This mechanism also highlights the general challenge when trying to account for wage dispersion relying on employer differences alone. As our decomposition shows, differences in human capital accumulation are the driver of rising earnings inequality over the life cycle. The covariance between human capital and wages (not shown) is small but positive and contributes little to earnings dispersion in the model. At age 40, the contribution of the covariance accounts for about 10 percent of the search component. This decomposition with a small contribution from search frictions (frictional wage dispersion) is consistent with results in Hornstein, 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.11 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.11 show selected conditional distribution functions for wealth by income quintile from model and SCF data.[30] 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).[31] For data on the joint distribution, we rely on the SCF data. We follow previous research (Diaz-Gimenez, Glover, and Rios-Rull, 2011; Kuhn, Schularick, and Steins, forthcoming) and rely on PSID data from 1984 to 1999 to trace out individual-level wealth dynamics. The repeated cross sections of the SCF data prevent such an analysis.[32]

Looking at the joint distribution in Figures 1.A.11a to 1.A.11c, we find that the model aligns closely with the SCF data for households ages 40 to 50. We focus on a single age group to alleviate concerns regarding the age structure in the model relative to the age structure in the data. The fit of this untargeted dimension is very good. The most notable difference between the model and data is that too many households from the fifth income quintile are at the bottom of the wealth distribution. The model hence generates too many income-rich but wealth-poor households, but

30. 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.

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

32. Pfeffer, Schoeni, Kennickell, and Andreski (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.11. Joint distribution of income and wealth and wealth dynamics

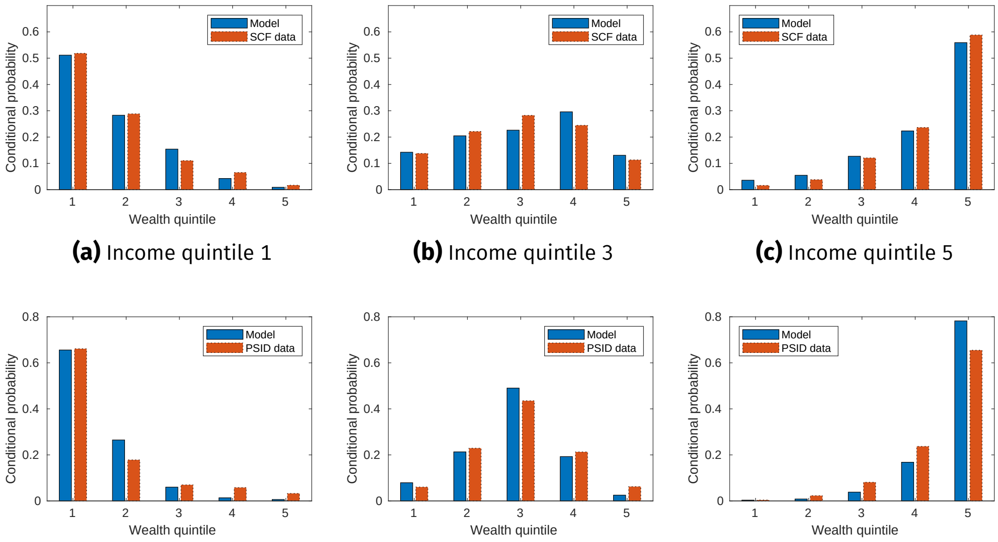

(a) Income quintile 1 **(b)** Income quintile 3 **(c)** Income quintile 5

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

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

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.11d to 1.A.11f, we find that the model matches closely the observed wealth dynamics in PSID data. The PSID surveys wealth only every five years during this time period, so we focus on five-year transition probabilities. We observe a high persistence of households' position along the wealth distribution. More than 80 percent of households from the first wealth quintile remain within the two bottom wealth quintiles over a five-year horizon. Similarly at the top, only about 30 percent of households from the top quintile end up in a lower quintile five years later. Comparing the model and data, we find that the model produces slightly too little wealth mobility at the bottom and the top. The middle of the distribution is closely matched. It is important to note that imputation of wealth information in the PSID likely leads to overstating the estimates of wealth mobility in the PSID data.

Table 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

1.A.12 Welfare effects of heterogeneity in job starting at age 20

Section 1.4.3 explores the welfare consequences of a bad start to the labor market. We consider 25-year-old workers in the main text to deal with the model assumption that all workers start working life ex ante identical as nonemployed workers with the lowest level of human capital. In this section, we use the employment distribution at age 25 to define stable and unstable jobs and endow 20-year-old labor market entrants with these states. We otherwise conduct the same welfare experiment and decomposition as in Section 1.4.3. Figure 1.A.12 shows the welfare effects at different levels of job stability and for three different wage levels. Wage levels are again taken using the wage distribution at age 25. We find again sizable welfare effects that are increasing in wage levels and job stability. The human capital component dominates the welfare effects at low-wage jobs, and the search component dominates at high-wage jobs. Comparing the decomposition at age 20 to age 25, we find that the human capital component gains in importance across all wage levels but that results lead to the same conclusions overall.

Figure 1.A.12. Welfare costs of job instability for 20-year-old workers in different job types

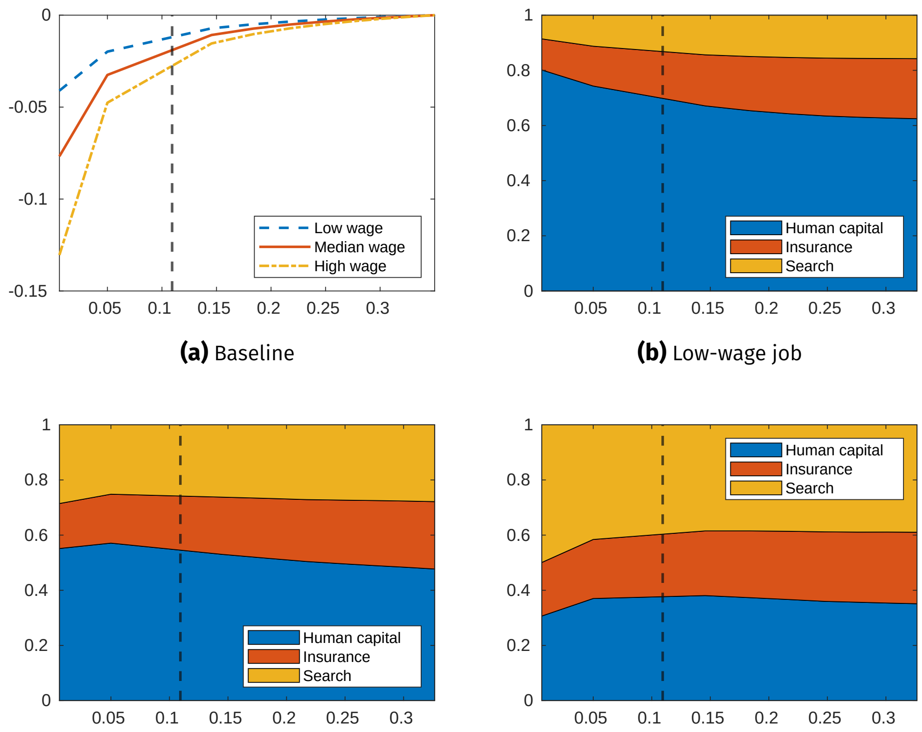

(a) Baseline

(b) Low-wage job

(c) Median-wage job

(d) High-wage job

Notes: Panel (a) shows the share of consumption that workers starting their career at a job characterized by wage w and separation rate λ (on horizontal axis) are willing to give up to avoid starting their career at the least-stable job with the same wage. Panels (b)-(d) show the relative contribution of the three effects for different wage levels. Medium wage correspond to the median wage at age 25, low wage corresponds to the 25th percentile of wages at age 25, and high wage corresponds to the 75th percentile of wages at age 25. The dashed vertical line shows the average separation rate at age 25.

1.A.13 Calibration for declining labor market dynamism

In Section 1.5, we describe how we calibrate declining labor market dynamism for the U.S. economy. Figure 1.A.13 shows the marginal distributions of separation rates for the baseline economy and in the economy with lower labor market dynamism ("Low mobility"). As explained in Section 1.5, the support of separation rates is compressed at the lower end so that unstable jobs become more stable. This compression shows up in Figure 1.A.13 as first-order stochastic dominance of separation rates in the baseline economy compared to the economy with lower mobility.

Figure 1.A.13. Marginal distribution of separation rates in the job-offer distribution

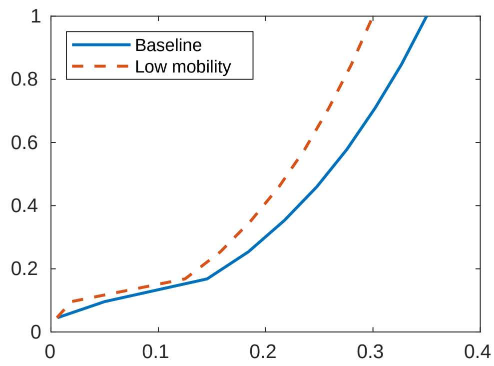

Notes: Marginal distribution of separation rates λ for the baseline economy and the economy with less labor market dynamism (Low mobility). The horizontal axis shows support of separation rates λ, with $\bar{\lambda}$ being higher in the baseline economy, as described in Section 1.5.

References

Aguiar, Mark, and Erik Hurst. 2013. "Deconstructing life cycle expenditure." *Journal of Political Economy* 121 (3): 437–92. [5, 28]

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *Quarterly Journal of Economics* 109 (3): 659–84. [7, 62]

Barnette, Justin. 2020. "Wealth After Job Displacement." [36]

Bayer, Christian, and Moritz Kuhn. 2018. "Which Ladder to Climb? Wages of workers by job, plant, and education." [61]

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, 28, 59]

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. [9, 10]

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

Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business & Economic Statistics* 21 (3): 339–53. [31]

Cajner, Tomaz, Ilhan Güner, and Toshihiko Mukoyama. 2020. "Gross Worker Flows over the Life Cycle." [8]

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1): 1–55. [6]

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull. 2003. "Accounting for the US earnings and wealth inequality." *Journal of political economy* 111 (4): 818–57. [22]

Couch, Kenneth A, and Dana W Placzek. 2010. "Earnings losses of displaced workers revisited." *American Economic Review* 100 (1): 572–89. [35]

Daly, Moira, Dmytro Hryshko, and Iourii Manovskii. 2016. "Improving the measurement of earnings dynamics." Working paper. National Bureau of Economic Research. [59]

Davis, Steven J. 2008. "The Decline of Job Loss and Why It Matters." *American Economic Review* 98 (2): 263–67. [41]

Davis, Steven J, and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, 1. [35]

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal choice and inequality." *Journal of political economy* 102 (3): 437–67. [28]

Diaz-Gimenez, Javier, Andy Glover, and Jose-Victor Rios-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. [62]

Fallick, Bruce, and Charles A Fleischman. 2004. "Employer-to-employer flows in the US labor market: The complete picture of gross worker flows." [4, 7, 41]

Friedman, Milton. 1957. *Theory of the consumption function.* Princeton university press. [33]

Fujita, Shigeru. 2018. "Declining labor turnover and turbulence." *Journal of Monetary Economics* 99: 1–19. [4, 7, 41, 42]

Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay. 2020. “Measuring Employer-to-Employer Reallocation.” Working paper. National Bureau of Economic Research. [41, 42]

Gourinchas, Pierre-Olivier, and Jonathan A Parker. 2002. “Consumption Over the Life Cycle.” *Econometrica* 70 (1): 47–89. [31]

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*, (710): [5, 6, 8, 12, 28, 59]

Hagedorn, Marcus, and Iourii Manovskii. 2010. “Search frictions and wage dispersion.” *Manuscript, University of Zurich and University of Pennsylvania*, [61]

Hall, Robert E. 1982. “The Importance of Lifetime Jobs in the US Economy.” *American Economic Review* 72 (4): 716–24. [3, 12, 13]

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L Violante. 2010a. “Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006.” *Review of Economic dynamics* 13 (1): 15–51. [28, 59]

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010b. “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006.” *Review of Economic Dynamics* 13 (1): 15–51. [13]

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. [23, 29, 38, 61]

Hubmer, Joachim. 2018. “The job ladder and its implications for earnings risk.” *Review of Economic Dynamics* 29: 172–94. [8, 22, 27, 28, 59]

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, 62]

Iacono, Roberto, and Marco Ranaldi. 2020. “The wage curve across the wealth distribution.” *Economics Letters* 196: 109580. [11]

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993a. “Earnings losses of displaced workers.” *American economic review*, 685–709. [35]

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993b. “Earnings Losses of Displaced Workers.” *American Economic Review* 83 (4): 685–709. [5]

Jarosch, Gregor. 2015. “Searching for Job Security and the Consequences of Job Loss.” [6, 8, 15, 35]

Jung, Philip, and Moritz Kuhn. 2018. “Earnings Losses and Labor Mobility Over the Life Cycle.” *Journal of the European Economic Association*, (May): [6, 8, 14, 15, 19, 22–24, 27, 28, 35, 36]

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. [9]

Krebs, Tom, Moritz Kuhn, and Mark LJ Wright. 2015. “Human capital risk, contract enforcement, and the macroeconomy.” *American Economic Review* 105 (11): 3223–72. [17]

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2017. “Gross worker flows over the business cycle.” *American Economic Review* 107 (11): 3447–76. [8]

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Sahin. 2010. “Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations.” *Review of Economic Studies* 77 (4): 1477. [8]

Kudlyak, Marianna, and Fabian Lange. 2014. “Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories.” *IZA DP*, (8663): [23]

Kuhn, Moritz. 2013. “Recursive Equilibria in an Aiyagari-Style Economy with Permanent Income Shocks.” *International Economic Review* 54 (3): 807–35. [36]

Kuhn, Moritz, and Jose-Victor Rios-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. [9, 10]

Kuhn, Moritz, Moritz Schularick, and Ulrike Steins. Forthcoming. “Income and Wealth Inequality in America, 1949-2016.” *Journal of Political Economy*, [9, 62]

Larkin, Kieran P. 2019. “Job Risk, Separation Shocks and Household Asset Allocation.” In *2019 Meeting Papers*. 1058. Society for Economic Dynamics. [8, 15]

Lise, Jeremy. 2012. “On-the-Job Search and Precautionary Savings.” *Review of Economic Studies* 80 (3): 1086–113. [8, 27]

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. “Wage Risk and Employment Risk over the Life Cycle.” *American Economic Review* 100 (4): 1432–67. [9, 35]

Meghir, Costas, and Luigi Pistaferri. 2004. “Income variance dynamics and heterogeneity.” *Econometrica* 72 (1): 1–32. [28, 59]

Michelacci, Claudio, and Hernan Ruffo. 2015. “Optimal Life Cycle Unemployment Insurance.” *American Economic Review* 105 (2): 816–59. [8, 29, 57]

Modigliani, Franco, and Richard Brumberg. 1954. “Utility analysis and the consumption function: An interpretation of cross-section data.” *Franco Modigliani* 1 (1): 388–436. [30]

Molloy, Raven, Christopher Smith, and Abigail K Wozniak. 2020. “Changing Stability in U.S. Employment Relationships: A Tale of Two Tails.” Working Paper 26694. National Bureau of Economic Research. [4, 7, 41, 42]

Molloy, Raven, Riccardo Trezzi, Christopher L Smith, and Abigail Wozniak. 2016. “Understanding declining fluidity in the US labor market.” *Brookings Papers on Economic Activity* 2016 (1): 183–259. [3, 7]

Morchio, Iacopo. 2020. “Work Histories and Lifetime Unemployment.” *International Economic Review* 61 (1): 321–50. [9, 12]

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. [22]

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. [62]

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

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. [28]

Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review* 95 (1): 25–49. [22]

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. [36]

Topel, Robert H, and Michael P Ward. 1992. “Job Mobility and the Careers of Young Men.” *Quarterly Journal of Economics* 107 (2): 439–79. [5, 28, 59, 60]

Chapter 2

The simple life? Heterogeneity in income risk and household portfolios*

Joint with Christian Bayer, Thomas Hintermaier and Moritz Kuhn

2.1 Introduction

The Matthew effect is typically summarized by “the rich get richer and the poor get poorer”. In this paper, we rely on German social security data that allow us to follow workers over the entire working life to provide evidence for such a Matthew effect for income, where low-risk individuals have higher average income growth and accumulate more wealth. Among workers, there is a group, the “unlucky” few, who not only face low income growth on average but also highly volatile incomes. Most income fluctuations, unemployment risk, and bad health outcomes are concentrated in this group. Looking at their wealth, one sees that these workers accumulate little even in comparison to their lower incomes. On top, they earn meager returns on what they manage to save, because they invest in poor return (i.e. liquid and safe) assets. On the other hand, the vast majority of workers face a much calmer income process with little turbulence over the life cycle, they have steady income growth, and can thus use high-return assets as savings devices.

* We thank Tatjana Mika and Michael Stegmann for their help regarding questions on the data. We thank conference participants at the mini conference on heterogeneous agents at Zurich and Kurt Mitman, Xavier Ragot for helpful comments and remarks. Financial support by ADEMU - 649396 grant is gratefully acknowledged. Kuhn thanks the DFG for financial support (DFG No. 433368816). Ploj gratefully acknowledges financial support from the DAAD, the Bonn Graduate School of Economics, and the DFG research training group 2281 “The Macroeconomics of Inequality.” The usual disclaimer applies.

We document these facts by combining largely unexplored data from the German retirement records on employment histories of males in West Germany with information on households' balance sheets. As in other data sources (Guvenen, Karahan, Ozkan, and Song, 2015), we document that income fluctuations have fat tails. Yet, we also document that this is driven to a significant extent by substantial and persistent heterogeneity in income risk across workers. While most workers face little income fluctuations, a smaller fraction has very volatile incomes. Most unemployment is highly concentrated within cohorts of workers and low-pay, no-pay cycles are prevalent among workers who become unemployed. We find that within one cohort 20 percent of workers account for around 80 percent of all transitions into unemployment. In the same vein, we also find a high concentration of changes in health status (healthy to sick) within cohorts, although health risks appear to be less concentrated than the incidences of unemployment.

Matching the income risks to household balance sheet data, we document that those who face the highest risks in the labor market earn the lowest returns in the asset market and also accumulate substantially less wealth even relative to their lower incomes. We rationalize this outcome as a result of persistent risk types. While households live most of the time in tranquil times with small income fluctuations, they can occasionally transit into a state of labor market turbulence. On average this happens rarely, once in 20 years, and mostly after a large negative income shock. This turbulent labor market situation persists for a significant period of time as we estimate the average duration of the turbulence to be around five years. The household optimally reacts to the negative income shock that preceded the labor market turbulence by drawing from its savings to smooth the negative income shock. In the aftermath of the shock, the household structures its portfolio in a way that it provides more insurance against negative income draws and this happens at the expense of worse old-age insurance. Because the portfolio becomes safer and more liquid, it earns lower returns.

We show in a calibrated model that such dynamics explain well the systematic differences in household portfolios and their returns both in the cross-section as well as over the life cycle. In the model, households face income shocks and stochastic transitions between risk types. They self-insure against these shocks by investing either into a low-return liquid asset that they can access at no cost every period or into an illiquid asset with higher but also risky returns. We model the illiquidity as a saving commitment, where the household can change the per-period flow into or out of its illiquid account by paying an adjustment cost. Such a saving contract, therefore, resembles a typical mortgage contract or traditional life insurance that pays out at retirement.

In both model and data, differences across risk groups in terms of wealth, portfolios, and returns increase as households age. For example, at age 40 high-risk households have a 50% lower wealth to income ratio, their portfolios contain a 20% higher share of liquid assets and they earn 20% lower returns on their assets compared to low-risk households. The presence of idiosyncratic second-order income risk increases aggregate capital accumulation and reduces welfare for low-risk households. The average household who enters the labor market as the low-risk type would consequently be willing to forgo 4% of its lifetime consumption to avoid living in the economy with second-order risk.

Related literature. This paper relates primarily to two strands of literature. First, it adds to the literature that investigates the heterogeneity in labor market risk. Going back at least to Hall (1982), economists have documented that workers have very different working lives. Recent empirical work has corroborated this finding and concluded that a representative earnings process will struggle to account for the regularities of earnings dynamics in the data (Guvenen, Karahan, Ozkan, and Song, 2015; Arellano, Blundell, and Bonhomme, 2017; Jung and Kuhn, 2018; Karahan, Ozkan, and Song, 2019; Kuhn and Ploj, 2020; Morchio, 2020). Findings of the literature show that labor income risk varies, among others, across age (Karahan and Ozkan, 2013; Blundell, Graber, and Mogstad, 2015; Jung and Kuhn, 2018; Kuhn and Ploj, 2020), educational and skill groups (Mukoyama and Şahin, 2006; Blundell, Graber, and Mogstad, 2015) and with the business cycle (Storesletten, Telmer, and Yaron, 2004; Bayer, Luetticke, Pham-Dao, and Tjaden, 2019). The increased availability of large detailed administrative data has also allowed for more granular and less parametric analysis of the heterogeneity in labor market risk with an increased focus on the higher moments of the earnings growth rate distribution (Guvenen, Karahan, Ozkan, and Song, 2015; Manuel, Blundell, and Stéphane, 2017). Our paper follows the spirit of these papers and presents additional evidence that workers systematically vary in terms of *lifetime* earnings risk, and that these differences can be to a large degree attributed to differences in lifetime unemployment and health risk. In this view, our work therefore also complements that of Schmillen and Möller (2012) and Morchio (2020) who report a high concentration of lifetime unemployment using German and U.S. data.

Second, we also contribute to the literature that studies the implications of heterogeneity in labor market risk for household savings behavior and portfolio heterogeneity. Findings of this literature predominantly show that income risk reduces the demand for risky assets, although the estimated effects are often small and insignificant (Guiso, Jappelli, and Terlizzese, 1996; Heaton and Lucas, 2000, 2001; Angerer and Lam, 2009; Palia, Qi, and Wu, 2014). Recently, Fagereng, Guiso, and Pistaferri (2018) find much larger effects than previously documented after improving on the identification of wage risk variability. Corroborating these findings, Chang, Hong, Karabarbounis, Wang, and Zhang (2020) also show that the empirically estimated response of the risky share to wage risk can be explained with a standard portfolio choice model that is augmented with idiosyncratic, time-varying income volatility.[1] We complement these findings on the empirical side and provide additional evidence for the effect of risk heterogeneity on wealth accumulation, portfolio returns, and asset participation. Furthermore, we also introduce a portfolio choice model with a frictional saving commitment, which we show can well explain the observed life-cycle patterns of heterogeneity in portfolio liquidity and wealth accumulation across households with different labor market risk.

The remainder of the paper is structured as follows. In Section 2.2 we introduce the data from the German social security records and present evidence on heterogeneity in unemployment risk, health risk, wage risk, and earnings risk. Section 2.3 explores the consequences of labor market risk heterogeneity for household savings decisions and portfolio composition using the HFCS data. In Section 2.4, we study the consequences of labor market risk through the lens of incomplete markets models. Finally, Section 2.5 concludes.

2.2 Heterogeneity in wage and earnings risk

Our empirical analysis revisits the question of how unequal working lives are based on largely unexplored administrative data from the German retirement records. These administrative data are particularly well suited to address this question as they cover the entire employment history of workers starting at age 14 and they also provide a high resolution of earnings dynamics with monthly earnings and labor market observations. Compared to social security data as, for example, used in Guvenen, Karahan, Ozkan, and Song (2015), these data also report life events like unemployment or health shocks that we can relate to observed earnings dynamics.

2.2.1 Data

German retirement records comprise the whole employment history of individuals starting from age 14 on a monthly basis. In addition to this exceptional coverage comes the high accuracy of the data. The data are used to determine old-age retirement and disability benefits and have been taken from the administrative records of employer reports to the social security system and are sent to individuals for double-checking of their recorded employment histories.[2] The data contain typical demographic information including age and sex. For our analysis, we focus on male

1. Taking a macroeconomic perspective, Bayer, Luetticke, Pham-Dao, and Tjaden (2019) show that business cycle variations in income risk induce aggregate portfolio rebalancing with non-negligible implications for aggregate consumption and output.

2. In contrast to the widely used German social security data from the unemployment insurance records (Antoni, Ganzer, Vom Berge, et al., 2019), the data used in this paper are composed by biography time not by calendar time so that entire employment histories of individuals are observed in these data.

workers with all of their employment records in West Germany because employment information from times in the German Democratic Republic is not comparable. We restrict the analysis to men, as labor market experiences of those women for whom we observe completed labor market histories in the data were still largely determined by a strongly changing institutional and cultural environment that makes studying their working lives an important and interesting topic of research of independent interest.

The data cover all workers covered by social security legislation and earnings represent labor market income subject to social security taxes. From this group, we select workers in cohorts born between 1935 and 1953, i.e. men that complete their 60th year of life between 1995 and 2013, so that we observe complete labor market histories for workers up to age 60. The sample is representative for each birth cohort and comprises approximately 410 worker observations for every year of birth. Earnings in the data are reported as so-called *pension points*. These pension points express monthly earnings relative to macroeconomic annual average earnings, an index that is updated since 1957 based on NIPA gross earnings data. The fact that earnings are expressed relative to an annual average allows for an easy comparison of earnings dynamics over time because time trends of average earnings are removed.[3] Some earnings observations are capped at the social security contribution limit and we focus for most of our analysis on uncapped earnings observations.[4] The share of capped observations varies with age from below 1 percent at age 20 to roughly 18 percent at age 55. We provide further details in Appendix 2.A.1.2.

The data provide us both with wages (daily earnings) and monthly gross earnings. We construct the average annual wage as the average daily earnings for all days with non-zero earnings and the annual earnings as the sum over all daily earnings of a calendar year. To compute wage and earnings growth between age j and age $j+1$, we compute growth rates as logarithmic first differences (log-growth rates), where we require that workers work for at least 90 days in both years. For computing the growth rates we, therefore, require that workers are at least partially attached to the labor market in both years used for computing the growth rates.

2.2.2 Wage and earnings risk

In this section, we first document facts on life-cycle earnings dynamics in the German labor market by playing on the strength of the data, which allow us to observe the entire labor market history of workers. Second, we provide evidence for large heterogeneity in earnings risk. Third, we document that earnings risk and earnings

3. Note that the construction of the data does typically not provide a representative cross-section of all employed workers at each calendar date. This implies that the average number of points at different calendar dates might significantly differ from one. See Appendix 2.A.1.1 for a discussion.

4. We show the time series of the social security limit in Appendix 2.A.1.2. The cap varied over time between 1.5 and 2.1 times the average income (= points).

levels are negatively correlated in contrast to the idea of a risk-return trade-off for human capital.

2.2.2.1 Wage and earnings growth dispersion over a worker's life

Figure 2.1 shows different measures of dispersion for wage and earnings growth over workers' working life. Figure 2.1a plots the distribution of log earnings growth rates by age and Figure 2.1b shows the corresponding life-cycle profiles for wages. Both figures highlight the overall large dispersion in growth rates across workers with a dispersion of earnings growth rates that is larger than that of wage growth rates. We also observe salient age pattern. The dispersion of wage and earnings growth rates is highest after labor market entry and decreases quickly within 5 years. The dispersion declines as workers age until roughly age 40 and remains mostly constant thereafter up until shortly before retirement.

Figure 2.1. Earnings and wage growth rate distribution

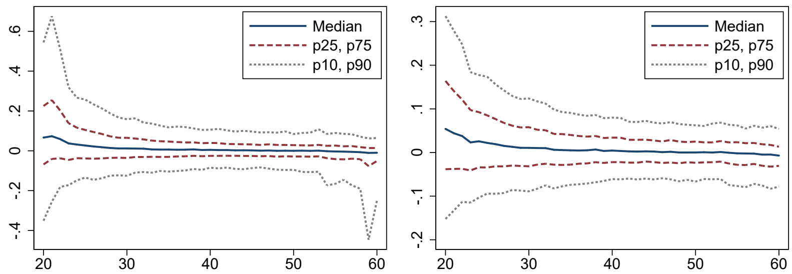

(a) Earnings growth by age **(b)** Wage growth by age

Notes: Percentiles of the annual earnings (the left panel) and wage (the right panel) growth rate distribution by age. We compute log growth rates for earnings observations that are below the social security contribution limit in both years and for workers that work at least 90 days in both years. Data for cohorts born between 1935 and 1953 used.

For earnings, we observe at labor market entry that the 90th and 10th percentile differ by about 100 percentage points. For wages, this dispersion is with over 50 percentage points still large but much lower. The difference highlights the importance of employment risk after labor market entry (Jung and Kuhn, 2018). The differences in wage and earnings growth dispersion decline with age and almost converge during the second half of the working life. The interquartile ranges for wage and earnings growth rates are with slightly over 25 percentage points still large but of similar magnitude at labor market entry. Generally, the decline in earnings variability is stronger than for wages and the decline is strongest between age 20 and 25. As workers age, the dispersion of wage and earnings growth rates does not only become more compressed but also more symmetric. At labor market entry, the distributions are strongly right-skewed and they converge to a roughly symmetric distribution for the second half of working life. The changing growth rate distribution suggests that there are underlying dynamics of the earnings process over the life cycle (Karahan and Ozkan, 2013; Kuhn and Ploj, 2020).

Looking at the median growth rate, we find a declining, slightly convex profile over the working life for both wages and earnings. The peak is at labor market entry and after that median growth rates decline to a slightly negative value at age 60. The declining convex growth rate profiles deliver the commonly observed concave life-cycle profiles of wages and earnings levels (for Germany, see Bayer and Kuhn (2019)).

Even when the distributions of growth rates become more symmetrical at around 40 years of age, they still remain significantly overdispersed. Figure 2.A.4 in Appendix 2.A.2 shows the growth rate distributions for the 40-year olds where we superimpose normal densities to highlight the strong deviation from a normal distribution, a fact emphasized by Guvenen, Karahan, Ozkan, and Song (2015) for the distribution of earnings growth in the United States. The most striking deviation from the normal distribution is the much higher kurtosis of the empirical distribution with much more mass at growth rates close to zero.

2.2.2.2 Persistent risk types

A standard way to model such an overdispersed income growth distribution is to assume mixing. Workers, or some workers, sometimes draw from a distribution with a high variance and sometimes from a distribution with a low variance (Guvenen, Karahan, Ozkan, and Song, 2015). Aside from this being a modeling device, we can also assign an economic interpretation to such mixing. Some workers undergo a more turbulent labor market than others. An interesting question when applying this interpretation is if the risk types are transitory, persistent, or permanent. For simplicity, we focus on the earnings growth distribution throughout the rest of the paper and approximate it by a mixing with only two risk types.[5]

Our data offer a unique opportunity to study the persistence of risk as they cover workers' entire labor market history. To understand the persistence and cross-sectional heterogeneity in earnings risk, we look at the distribution of earnings-growth volatility (within person) over the entire working life and (rolling) 5-year standard deviations of growth rates. In the first case, we compute worker-specific standard deviations of growth rates during the worker's prime-age working life from age 25 to 55. In the second case, we compute a worker-specific age-varying 5-year (rolling) standard deviation of the growth rate to allow for slow-moving changes in the earnings process over the life cycle. Specifically, we compute for each worker at each age j the standard deviation of growth rates using growth rates at ages

5. This will capture most features of the earnings growth distribution and simplifies the discussion. Allowing for more worker types will mechanically improve the fit to the data especially with respect to the tails of the distribution.

$j, j-1, \ldots, j-4$. Figures 2.2a and 2.2b provide estimates for the distributions of earnings-risk estimates across workers.

Looking first at the lifetime volatility in Figure 2.2a, we find a large dispersion of earnings growth volatilities. On average, the lifetime standard deviation of individual earnings growth rates is 0.17 but we also find that there is large heterogeneity with a substantial share of workers having standard deviations smaller than 0.05 but also larger than 0.3. The support of the distribution of volatilities is by construction positive and it shows a strong asymmetry with a long right tail. The 90-50 ratio of the distribution is 2.8 demonstrating that at least 10 percent of workers live working lives that are almost three times as volatile as that of the typical (median) worker.

Figure 2.2b considers the distribution of age-dependent earnings growth volatilities. We consider the distribution of 40-year-old workers roughly in the middle of working life. The distribution shows a similar shape to the lifetime volatilities. There is now more mass close to zero indicating that a substantial share of workers experiences wage growth rates at age 40 that feature almost no volatility. As for the lifetime volatilities, there is still a large fraction of workers with very volatile wage dynamics. The average volatility is slightly below 0.1 but the 90th percentile is at 0.23 resulting in a 90-50 ratio of 4.2 and, hence, even larger than for the lifetime distribution.

This higher dispersion of volatilities at a given age compared to the lifetime suggests some mean reversion in risk. At the same time, the large dispersion of lifetime volatility is hardly consistent with very transitory episodes of high earnings volatility. To explore the persistence of risk heterogeneity, we regress moving win-

Figure 2.2. Distribution of realized earnings growth rate volatility

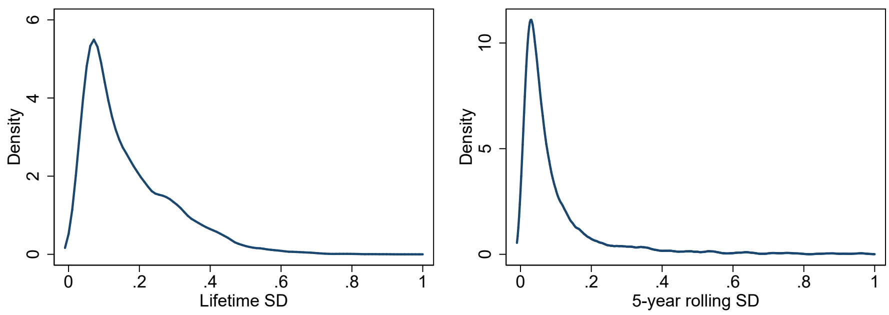

(a) Lifetime volatility **(b)** 5-year volatility at age 40

Notes: The left panel shows the distribution of the realized earnings growth rate volatility during the working life, computed as the standard deviation of annual earnings growth rate between ages 25 and 55. The right panel shows the distribution of the realized short-run earnings growth rate volatility at age 40, computed as the 5-year rolling standard deviation of annual earnings growth rate. Kernel density estimates shown. Log growth rates computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. Data for cohorts born between 1935 and 1953 used.

dows of future worker-specific volatilities on current worker-specific volatilities. If differences in realized volatility only stemmed from differences in the realization of a homogeneous stochastic process, e.g. because everyone draws from the same normal-mixture, then we should not be able to predict future volatility based on current volatility. To test this, we run the following predictive regression:

$$\text{SD(Growth rate)}_{i,j+5,t+5} = \alpha + \beta \times \text{SD(Growth rate)}_{i,j,t} + \gamma_j + \delta_t + \varepsilon_{i,j+5,t+5} \tag{2.1}$$

where $\text{SD(Growth rate)}_{i,j,t}$ is the 5-year rolling standard deviation of growth rate of earnings for worker i at age j at time t. As before, growth rate volatilities at age j are computed based on growth rates for periods $j, ..., j-4$ so that there is no overlap of samples for the right- and the left-hand side of the regression. We also control for age fixed effects γ_j and year fixed effects δ_t. The parameter of interest β indicates the persistence of individual-level earnings volatility, and hence, its predictability. Table 2.1 presents the estimated β coefficients from equation 2.1. We find a significant degree of persistence, see column (1). Annualizing the β-coefficient ($\beta^{\frac{1}{5}}$) yields a persistence of 0.75.

Column (2) shows that the persistence is not a pure artifact of rolling window estimates in the presence of homogeneous risk. It reports the results of estimating

Table 2.1. Predictability of future earnings volatility

	SD(Log growth rate)$_{t+5}$			
	(1) Data	(2) 1 type (simulation)	(3) 2 types (simulation)	(4) 2 types with transitions (simulation)
SD(Log growth rate)$_t$	0.2327***	0.0165***	0.6546***	0.2805***
	(33.94)	(5.37)	(177.09)	(76.81)
Observations	69530	105000	105000	105000
R^2	0.056	0.000	0.427	0.077

Notes: This table reports OLS regression estimates for equation 2.1. Regressions on empirical data include age and year fixed effects, while regressions on simulated data include age fixed effects. Workers between ages 30 and 50 considered. When dealing with empirical data, growth rates are computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. Data for cohorts born between 1935 and 1953 used. Simulation results presented for three cases. In the 1 type case, earnings of all workers are drawn for the same stochastic earning process. In the 2 types case, earnings of 80% of workers are drawn from the low-risk earnings process, and earnings of the other 20% are drawn from the high-risk earnings process. In the 2 types case with transitions, stochastic transitions between the 2 risk types are allowed. The low-risk type has a 4% probability of becoming a high-risk type, and the other transition probabilities are set in a way to preserve a constant share of both types over the life cycle. Robust standard errors used.

t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the persistence of the dispersion in earnings growth in a simulated sample with homogeneous earnings risk. Given this result, we next ask how we could match the empirical evidence with a simple mixing model with persistent risk types. For simplicity, we consider only two risk types. One type we match to the bottom 80% of the earnings volatility distribution, because we find these workers have similar risk (Appendix 2.A.3 for details), and the other, high-risk type we match to the top 20% of earnings volatility realizations. Column (3) in Table 2.1 reports estimates from a simulated sample that does an 80-20 split in terms of permanent risk types. Clearly, the empirical persistence in realized earnings growth rate standard deviations is too low to warrant the assumption of fixed risk types. Just as homogeneous types lead to a too low autocorrelation, permanent types would lead to a too strong autocorrelation in realized dispersion of earnings-growth compared to the data.

Column (4) in Table 2.1 finally shows the estimation results on simulated data where on average 4% of low-risk workers become high-risk every year and on average remain high-risk for 6.3 years, so that the share of high-risk workers remains constant over the life cycle. Such a simple refinement to the process fits the empirical persistence in volatility very well.

Figure 2.3 compares the empirical distribution of growth rate volatilities in the data with the ones from our simulations. The model with two persistent—but not permanent—risk types fits the shape of both the distribution of lifetime and 5-year volatility well in terms of shape. Potentially, one would need a further risk type of extremely high volatility to fully match the tail of earnings risk, which we abstain from for clarity of exposition.

Figure 2.3. Distribution of earnings growth rate volatility

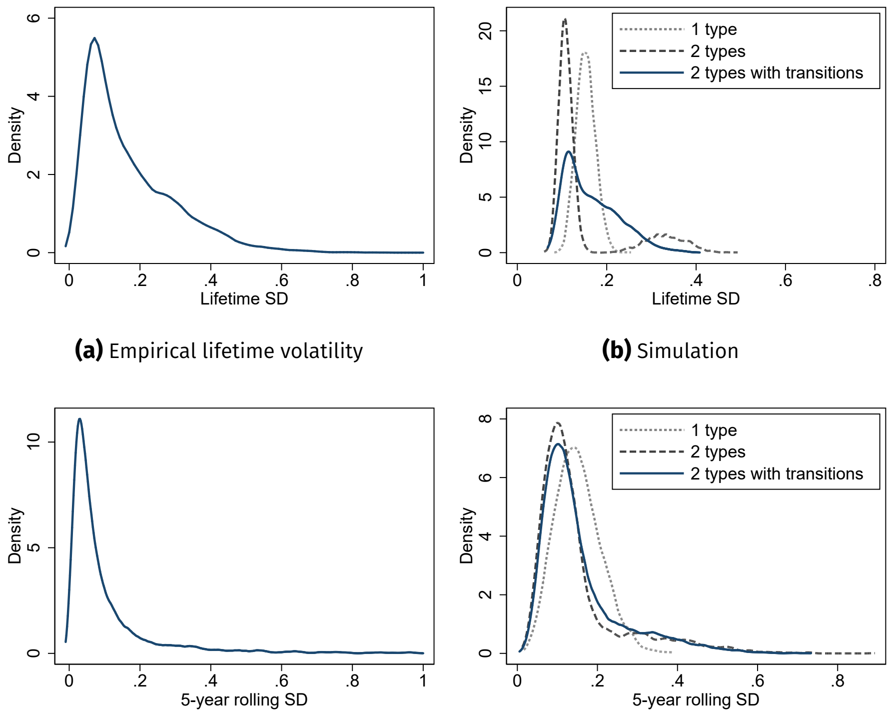

(a) Empirical lifetime volatility

(b) Simulation

(c) Empirical 5-year volatility at age 40

(d) Simulation

Notes: The left panels show the empirical distributions of the realized earnings growth rate volatility, while the right panels show the simulated distributions. Kernel density estimates shown. Log growth rates computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. Data for cohorts born between 1935 and 1953 used.

2.2.2.3 Non-existence of a risk-return trade-off

One interesting question is whether times of heightened earnings growth variance are compensated by higher average earnings offering a risk-return trade-off for human capital. Figure 2.4a shows that this is not the case. We classify workers based on their realized lifetime earnings volatility as described before. The figure looks at average earnings using pension points, i.e. earnings relative to the average worker. We find that all earnings profiles have the typical concave shape and that there is a striking negative correlation between earnings levels and risk. Workers in the high-risk group earn significantly less than the low-risk workers. This protracted lower earnings growth shows up also when looking at accumulated lifetime earnings in Figure 2.4b, where we compute lifetime earnings by accumulating pension points for each worker up to the respective age. A worker with average earnings in each year would hence have a linearly increasing profile with a slope of one. We find that the high-risk group increasingly lags behind the low-risk group in terms of accumulated earnings. Hence, workers with more volatile earnings histories also have lower average earnings and less earnings growth.[6]

Figure 2.4. Life-cycle earnings profiles by risk type

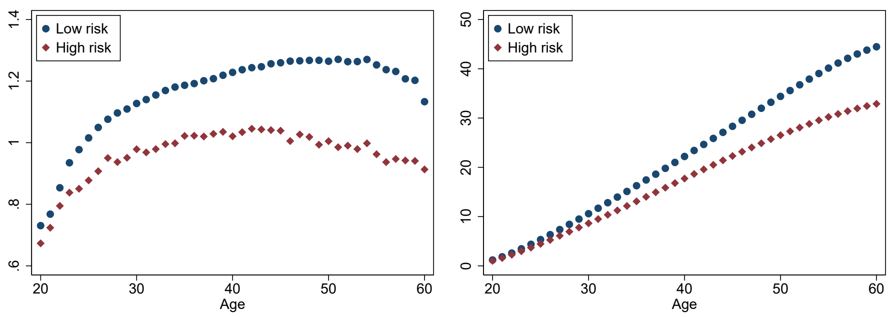

(a) Average earnings

(b) Cumulative earnings

Notes: Life-cycle profiles of annual and cumulative earnings. In the left panel, only observations with positive earnings used. Risk types are defined using the standard deviation of earnings growth rates between ages 25 and 55, where the bottom 80% of workers belong to the low-risk group and the top 20% to the high-risk group. Log growth rates for earnings are computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. At least 20 valid growth rate observations are required for assigning a risk type. Data for cohorts born between 1935 and 1953 used.

2.2.3 Life events and risk types

In the next step, we play upon a further strength of our data, the fact that we also observe important labor market events like unemployment and sickness spells. The

6. See Kuhn and Ploj (2020) for a microfoundation of such a negative correlation.

joint analysis of estimated earnings risk and such life events allows us to link risk type differences to observable realizations of labor market risk. We document that the high earnings volatility group not only has lower income growth, it also has higher unemployment and health risks as we show next. These two life events are typically associated with large earnings shocks (e.g., Low and Pistaferri, 2010; Guvenen, Karahan, Ozkan, and Song, 2015). The retirement records are particularly well suited for such an analysis as the social security system is also responsible for unemployment and long-term sickness insurance and therefore the data comprise high-quality and high-frequency information on these events. We assign employment and health status following the coding in the data. For employment, we only consider employment under social security legislation. For unemployment, we rely on a benefit-based definition.[7] Sickness spells in the data are either long-term sickness periods (typically longer than six weeks) during which workers are unable to work or situations when workers have to be hospitalized.[8]

2.2.3.1 Unemployment

The incidence of unemployment is highly unequally distributed across workers. Just as earnings growth risk is concentrated, so is unemployment risk. Most people do not experience any unemployment during their entire working lives while others circle repeatedly in and out of low-paid employment. To construct a measure of concentration of unemployment risk, we follow workers from cohorts born between 1935 and 1953 over their entire working life and count for each worker all transitions between employment and unemployment between age 20 and 60.[9] At age 60, we have a cross-section of workers with different stocks of accumulated employment to unemployment transitions. We compute the Lorenz curve of these stocks as our measure of the concentration of unemployment risk. Figure 2.5a shows the resulting Lorenz curve together with the Lorenz curve of a simulation assuming uniform unemployment risk. In Figure 2.5b we add to the extensive margin of unemployment risk also the intensive margin by considering the distribution of accumulated unemployment duration for the same cohorts.

We see that transitions into unemployment out of employment are highly concentrated among a small group of workers. While roughly 50 percent of all workers do not transit from employment into unemployment during their entire working life, 10 percent of all workers account for roughly 50 percent of all transitions and about 20 percent of workers account for almost 80 percent of all transitions. In the simula-

7. Hartung, Jung, and Kuhn (2016) show based on data from unemployment records that there is an overlap of 80 percent of benefit recipients and registered unemployed in the 2000s before the overhaul of the unemployment insurance system in Germany by the *Hartz* reforms.

8. Workers receive sick leave benefits (Lohnfortzahlung im Krankheitfall) during the first six weeks of a sickness spell.

9. We report figures for individual cohorts in Appendix 2.A.7.

Figure 2.5. Concentration of unemployment risk

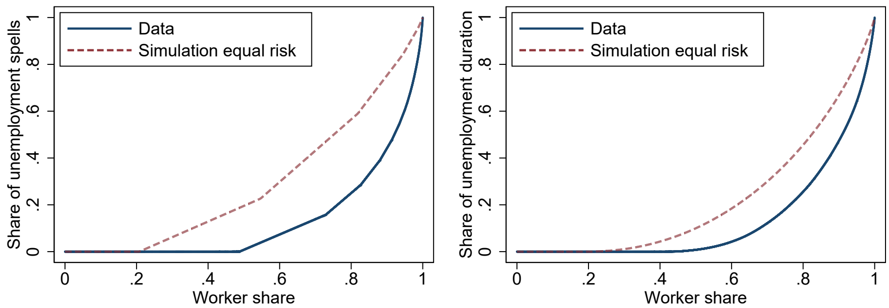

(a) Spells **(b)** Duration

Notes: Lorenz curve of accumulated unemployment spells (left panel) and unemployment duration (right panel) between ages 20 and 60 for cohorts of workers born between 1935 and 1953. The simulation assumes all workers have the same unemployment risk.

tion with uniform risk, we find only a fifth of workers without unemployment spells and that unemployment is much more equally spread with 50 percent of workers accounting for roughly 80 percent of unemployment transitions. Morchio (2020) reports a similar concentration of labor market transitions for the United States. When we look at the distribution of unemployment duration in Figure 2.5b, we find a similar pattern of inequality in unemployment risk.

Figure 2.6. Lifetime unemployment by risk type

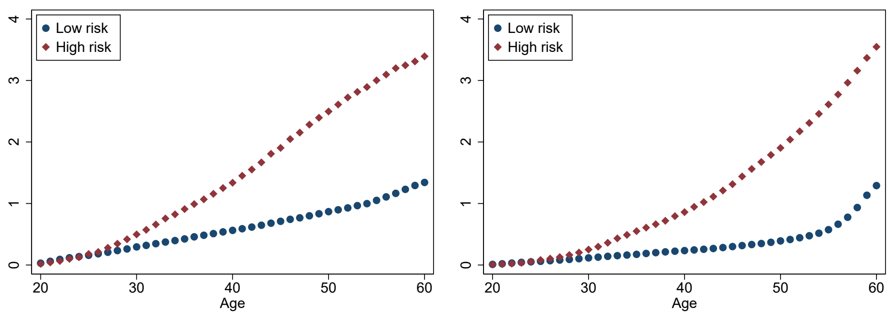

(a) Unemployment spells **(b)** Unemployment duration

Notes: The left panel shows the life-cycle profile of the number of independent unemployment spells. The right panel shows the life-cycle profile of the total number of years spent as unemployed. Risk types are defined using the standard deviation of earnings growth rates between ages 25 and 55, where the bottom 80% of workers belong to the low-risk group and the top 20% to the high-risk group. Log growth rates for earnings are computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. At least 20 valid growth rate observations are required for assigning a risk type. Data for cohorts born between 1935 and 1953 used.

Figure 2.6a relates observable unemployment risk to worker risk types from the previous subsection, where we categorize workers on the basis of their realized earnings growth rate volatility between age 25 and 55. For both risk types, we document the average accumulated number of unemployment spells by age. The resulting life-cycle patterns of accumulated unemployment spells are striking. We find a strong correlation between unemployment risk and risk types. Workers in the high-risk group experience on average around three unemployment spells during their working life, while the low-risk majority experiences only slightly more than one unemployment spell on average.

2.2.3.2 Health shocks

Health shocks are one of the major sources of earnings risk (Low and Pistaferri, 2010). To measure health risk, we proceed as in the case of unemployment risk by playing on the strengths of the social security data that provide direct information on transitions to long-term sickness periods capturing disability according to the social security system.[10] Health risk is defined as transiting to such a sickness period.

Figure 2.7. Concentration of sickness risk

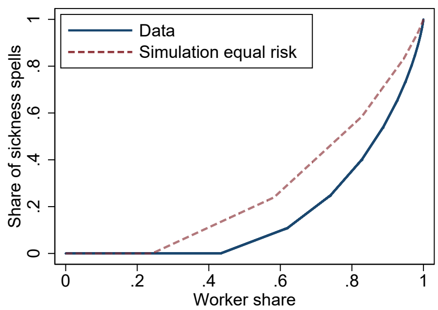

(a) Spells

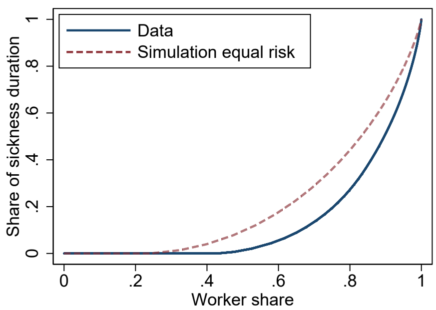

(b) Duration

Notes: Lorenz curve of accumulated sickness spells (left panel) and sickness duration (right panel) between ages 20 and 60 of cohorts of workers born between 1935 and 1953. Simulation assumes all workers have the same health risk.

Figure 2.7a shows the Lorenz curve of transitions to sickness periods together with a Lorenz curve for a simulation where all workers have the same probability of getting sick. Figure 2.7b adds to the extensive margin of health risk the intensive margin of health risk by considering the Lorenz curve of accumulated sickness duration for the same cohorts.

10. The German disability insurance system distinguishes between not being able to work in the old occupation and not being able to work at all. Disability in our data only refers to the latter definition. Occupational disability is not covered by the social security program.

Like in the case of unemployment risk, we find that health risk is highly concentrated, although slightly less than unemployment risk. The 20 percent of workers with the most transitions account for 60 percent of all transitions from healthy to sick and around 40 percent of workers do not go through a single sickness period up to age 60 (Figure 2.7a). Looking at the sickness duration in Figure 2.7b, we see that durations are slightly more concentrated than incidences with 20 percent of workers accounting for 70 percent of the total duration in sickness. We conclude that severe health shocks are rare, severe, and unequally distributed across workers.

Finally, Figure 2.8a relates observable health risk to the worker risk types from the previous section. For both risk types, we document the average accumulated number of sickness spells by age. As for unemployment before, we find again a strong positive correlation between observable health risk and risk types. High earnings risk workers have on average more than 50% more sickness spells and they have them much earlier in their lives. At age 50 the average low-risk worker has had slightly more than one long-term sickness episode while the average high-risk worker has had twice as many. In terms of health risk, the divergence happens mostly between ages 40 and 50. High-risk workers accumulate as many long-term sickness spells at age 50 as their low-risk counterparts accumulate when they turn 60.

To summarize, we document large heterogeneity in earnings risk and that risk is highly concentrated. But we also provide evidence for transitions between risk types over the working life. Finally, we document a strong negative correlation between earnings risk and earnings growth and levels, a finding inconsistent with a risk-

Figure 2.8. Lifetime sickness by risk type

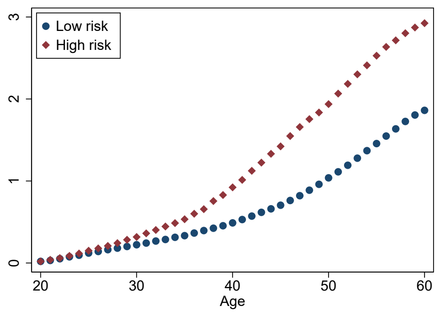

(a) Sickness spells

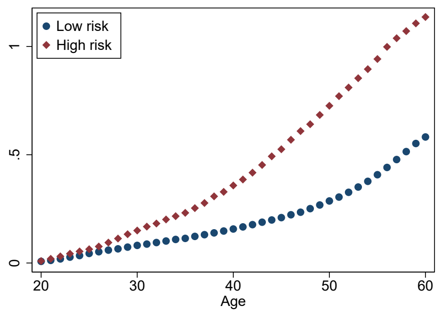

(b) Sickness duration

Notes: The left panel shows the life-cycle profile of the number of independent sickness spells. The right panel shows the life-cycle profile of the total number of years spent sick. Risk types are defined using the standard deviation of earnings growth rates between ages 25 and 55, where the bottom 80% of workers belong to the low-risk group and the top 20% to the high-risk group. Log growth rates for earnings are computed only for earnings observations that are below the social security contribution limit in both years and only for workers that work at least 90 days in both years. At least 20 valid growth rate observations are required for assigning a risk type. Data for cohorts born between 1935 and 1953 used.

return trade-off for human capital. In the next step, we explore if and how this heterogeneity in earnings risk correlates with the savings behavior by looking at wealth accumulation and portfolio allocation.

2.3 Risk heterogeneity and household portfolios

Given the large and persistent differences in earnings stability and other fundamental risks across workers, it is natural to expect that this has repercussions for household savings decisions. In this section, we use data from the 2010 and 2014 waves of the *Household Finance and Consumption Survey* (HFCS) to document that heterogeneity in labor market risk is indeed systematically associated with differences in wealth accumulation, realized returns, and portfolio allocation. The HFCS data is a euro-area equivalent of the U.S. Survey of Consumer Finances (SCF) and collects detailed information about the financial situation of households (ECB, 2016). In addition to data on the financial situation of households, it also provides information on the labor market experiences of household members. For our analysis, we focus on households from Germany and restrict the sample to households with male employed household heads between ages 20 and 60 at the survey date. We drop all households whose head is self-employed.

For financial variables, we focus on wealth, liquid and illiquid assets, and capital returns. We follow the construction of these variables as common in the literature (Kuhn and Rios-Rull, 2016; Bhutta, Bricker, Chang, Dettling, Goodman, et al., 2020; Kuhn, Schularick, and Steins, 2020). We construct household wealth as the net worth of the household by taking the difference between the value of all household assets and the value of all household debt. Household income is income from all sources before taxes, but including transfers. We consider two definitions of illiquid wealth. In a narrow definition, we include housing, real estate, private pension, and life insurance wealth. In a broader definition, we additionally include mutual funds, managed accounts, and stocks in line with the broader definition applied in parts of the literature (c.f., Kaplan and Violante, 2014; Bayer, Luetticke, Pham-Dao, and Tjaden, 2019). We compute the capital returns as the sum of income from financial assets, rental income, capital gains, and net imputed rents for owner-occupied housing and express this capital income relative to household wealth. While income from financial assets and rental income are directly observed in the HFCS data, capital gains and imputed rents need to be estimated based on external sources. We impute these components by applying average historical returns from Jordà, Schularick, and Taylor (2017) to individual asset categories and impute net rents as net rental income minus interest payments on housing and real estate loans using their data for Germany for the years 1963 to 2014.

2.3.1 Measuring labor market risk in the HFCS data

The key challenge for connecting the current analysis to the empirical analysis from the previous section is to identify risk types in the HFCS data as the HFCS data do not have high-frequency panel data on employment histories. To overcome this limitation, we use employment duration and earnings levels as proxies for exposure to labor market risk, as these are variables that significantly correlate with labor market risk. Evidence presented in Section 2.2 shows that workers with high labor market risk have lower earnings and a higher risk of unemployment, while numerous empirical studies also document a strong negative correlation between employment duration and job loss (see among others Farber, 1999; Menzio, Telyukova, and Visschers, 2016; Jung and Kuhn, 2018).

Building on this evidence, we infer the risk type of workers in the HFCS data using the information on employment tenure, earnings, and age. Due to differences in the dataset construction, the information on employment duration is not directly comparable between both datasets. In the social security data we observe employment duration, whereas the HFCS data give information on employer tenure. The former refers to the number of years since the end of the last nonemployment and the latter to the number of years that the worker has been working for his/her current employer. Consequently, job-to-job transitions affect tenure but not employment durations.[11] Furthermore, also earnings levels are not directly comparable between both datasets. In order to make variables in both datasets as comparable as possible, we, therefore, discretize tenure and earnings into age-specific quartiles.

In the next step, we specify the following regression model

$$\begin{aligned}\text{High risk type}_{i,t} &= \alpha + \sum_{j=2}^{8}\beta_j a_{i,t}^j + \sum_{k=2}^{4}\gamma_k d_{i,t}^k + \sum_{l=2}^{4}\delta_l y_{i,t}^l + \sum_{j=2}^{8}\sum_{k=2}^{4}\epsilon_{j,k} a_{i,t}^j d_{i,t}^k \\ &+ \sum_{j=2}^{8}\sum_{l=2}^{4}\eta_{j,l} a_{i,t}^j y_{i,t}^l + \sum_{k=2}^{4}\sum_{l=2}^{4}\nu_{k,l} d_{i,t}^k y_{i,t}^l + \varepsilon_{i,t} \qquad (2.2)\end{aligned}$$

where the probability of being a high-risk type is modeled as a function of dummies for age bins, age-specific quartiles for employment duration and earnings, and the interactions of these variables. In particular, dummy variables High risk type$_{i,t}$ and $a_{i,t}^j$ indicate for a worker i in year t whether he is characterized as the high risk type, and if he belongs to the age bin j, while $d_{i,t}^k$ and $y_{i,t}^l$ indicate if he belongs to the k-th and l-th age-specific quartile of employment duration and earnings.

Given the binary nature of the dependent variable, we estimate the model (2.2) on the social security data using a logit regression, and impute high-risk types in

11. Job-to-job transitions are substantially lower in the German labor market compared to the U.S. labor market (Jung and Kuhn, 2014).

the HFCS data through multiple imputations using $m = 5$ draws from the error term of the model. In this way, we account for the uncertainty related to the imputation of risk types. Our choice for the number of implicates is guided by the design of the HFCS dataset which already includes 5 implicates. Since our imputation model is highly parametrized and therefore difficult to present in a concise manner, Table 2.2 reports regression results for a simplified version of the imputation model, which includes linear effects for employment duration and earnings. As expected, we find that the probability of being characterized as the high-risk type is negatively related both with employment duration and with earnings. At the same time, the limited number of variables that can be included in the imputation model also limits its accuracy. As reported in the last row of Table 2.2, our fully specified imputation model given by equation 2.2 has a hit rate of 72.7%.[12]

Table 2.2. Employment tenure, earnings and risk type: results from retirement accounts data

High-risk type	(1) OLS	(2) Logit
Employment duration	-0.0114***	-0.0920***
	(-97.26)	(-89.03)
Earnings	-0.0795***	-0.3442***
	(-25.99)	(-17.21)
Observations	178591	178591
R^2	0.113	0.121
Hit rate		0.727

Notes: This table reports estimates for OLS and logit regressions of an indicator for the high-risk type on employment duration, earnings, and age dummies. Earnings expressed in pension points. Workers between ages 20 and 60 considered. Pseudo R^2 reported for the logit model. Hit rate reported for the imputation model specified in equation 2.2. Robust standard errors used.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2.3.2 Labor market risk and household portfolios

On the basis of this imputation, we explore the relationship between labor market risk and financial variables in the HFCS data. First, we present the differences in portfolio outcomes (wealth-to-income ratio, share of liquid assets, and returns on wealth) across the imputed types. In Section 2.4 we also provide additional details

12. Hit rate computed as the share of observations with the correctly imputed risk type: Hit rate $= \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(\hat{\mathrm{RT}}_i = \mathrm{RT}_i)$

by splitting up the results by age. We find that high-risk types have lower wealth-to-income ratios, hold a larger share of their wealth in liquid form and earn lower returns, see the top panel of Table 2.3.

Second, we run more flexible regressions where we regress portfolio outcomes (wealth-to-income ratios, the share of liquid assets, and returns on wealth) linearly on earnings and tenure, the variables we use to classify households in terms of risk. In addition, we also include fixed effects (dummies) for age, education, and year. All results from these regressions are reported in the bottom panel of Table 2.3. Given the noisiness of the imputation procedure, these regression results provide additional supporting evidence for the uncovered substantial differences in portfolio outcomes.

Table 2.3. Wealth, portfolio composition and returns on wealth

	(1) Net wealth (in 1000s)	(2) WTI	(3) Illiquid share (in %)	(4) Illiquid share (wide) (in %)	(5) Return (in %)
	Difference between imputed types				
Overall average	148.73	2.36	51.42	54.39	3.53
Average low-risk type	156.12	2.43	52.73	55.85	3.66
Average high-risk type	125.58	2.14	47.33	49.83	3.13
Difference low-high	30.54	0.28	5.40	6.02	0.53
	(2.22)	(1.22)	(1.27)	(1.42)	(1.70)
	Linear regression				
Tenure	2.517***	0.042***	0.469***	0.495***	0.040***
	(3.80)	(4.34)	(3.45)	(3.68)	(3.66)
Earnings (in 10 thousands)	16.212**	-0.002	0.588+	0.912**	0.099***
	(3.26)	(-0.12)	(1.94)	(3.00)	(4.55)
Observations	2389	2389	2389	2389	2389
R^2	0.218	0.071	0.130	0.145	0.112

Notes: The top panel shows sample means of net wealth, wealth-to-income ratio, portfolio share of illiquid wealth, and portfolio returns for both risk types. The bottom panel reports OLS regression estimates, where left-hand side variables are regressed on tenure and earnings of the household head. Additional controls: age FE, education FE, and year FE. Illiquid wealth includes net housing wealth, real estate wealth, pension, and life insurance wealth. The wide definition of illiquid wealth additionally includes holdings of mutual funds, stocks, and managed accounts. The sample includes households with male household heads between ages 20 and 60. Net wealth, wealth-to-income ratios, portfolio shares, and returns winsorized at the 2.5 and 97.5 percentiles. Data from the HFCS for Germany. Robust standard errors adjusted for the survey design with multiple imputations. Standard errors for the difference in means between risk types obtained using bootstrap.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In the first step, we look in column (1) at labor market risk and household wealth. We find that households with lower labor market risk have on average more wealth. On average the high-risk type has 31 thousand Euros less net wealth, a fifth less than the low-risk type. Nineteen thousand euros higher earnings or twelve more years of tenure correspond to the same size increase when using the linear regression results. This suggests that the wealth to income ratio is higher for low-risk households. Indeed this is what we find in column (2). We still find a significant negative effect of the risk type on wealth-to-income ratios. The low-risk type has a 14% higher wealth to income ratio. In the linear regression tenure is the driving factor. In fact, Kuhn and Ploj (2020) report the same systematic positive relationship between tenure and wealth-to-income ratios for the United States using 25 years of pooled data from the Survey of Consumer Finances.

Columns (3) and (4) report regression results with the illiquid portfolio share as the left-hand side variable of the regression. In column (3), we use the narrow definition of illiquidity and in column (4) the broader definition of illiquidity. In both cases, we find a significant negative effect of risk on illiquidity. Households with higher labor market risk have an almost 14% larger share of liquid assets. On top, in column (5) we see that low-risk households earn, with an extra 50bp, a more than 16% higher return on their wealth. This largely mirrors their higher share of illiquid assets, which have on average higher returns than the more liquid ones (Kaplan and Violante, 2014; Jordà, Knoll, Kuvshinov, Schularick, and Taylor, 2019a), such that differences in portfolio allocation lead workers with lower labor market risk to achieve a higher return on wealth.

In Table 2.4, we regress the return on wealth on the share of illiquid assets in the household portfolio using our broad and narrow definition of illiquidity.[13] In both cases, we find a highly statistically significant effect of the illiquidity share on the return on wealth. The point estimates imply that a portfolio that is fully invested in liquid assets earns a 4.7, respectively, 5.3 percent lower return compared to a portfolio that is fully invested in illiquid assets. This large difference implies that also empirically observed differences of 5 percentage point differences in portfolio shares will translate in return differences of more than 25 bp. The estimated tenure effect in columns (3) and (4) implies an increase in the illiquid share of 5 percentage points from 10 years of additional tenure and, hence, a roughly 25 bp higher return.

Going further into details, we find in line with the literature on household finance (see Campbell, 2006; Guiso and Sodini, 2013) that a large part of the differences in the illiquidity share is accounted for by the extensive margin of asset ownership for the most important illiquid asset classes, i.e. housing, pension, and investment plans (see Table 2.5). High-risk households have an 18% lower proba-

13. In a setup with two assets, a liquid and an illiquid asset, the return on wealth can be computed as the weighted return on both assets. Consider a liquid asset with return r_L and an illiquid asset with return r_I. Then if s denotes the share of illiquid wealth in a portfolio, the portfolio return can be expressed as $r_P = (1-s)r_L + sr_L = r_L + s(r_I - r_L)$. A regression of portfolio return on the share of illiquid wealth, therefore, identifies the average return on liquid wealth (r_L) and the return spread between the illiquid and the liquid wealth ($r_I - r_L$).

Table 2.4. Return and portfolio share of illiquid wealth

	(1)	(2)
	Return (in %)	
Illiquid share	4.688***	
	(20.52)	
Illiquid share (wide)		5.278***
		(25.98)
Constant	1.123***	0.663***
	(7.64)	(5.22)
Observations	2389	2389
R^2	0.359	0.446

Notes: This table reports OLS regression estimates of portfolio returns on the share of illiquid wealth. The sample includes male households heads between ages 20 and 60. Return and illiquid shares winsorized at the 2.5 and 97.5 percentiles. Robust standard errors adjusted for the survey design with multiple imputations. Data from the HFCS for Germany.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

bility to be homeowners. For taking the higher risk of being a homeowner with a mortgage, the probability decreases even by 26%. High-risk households also have a 17% lower probability to contribute to a voluntary pension plan and are 27% less likely to invest in equities. Regression results reported in the bottom panel support these findings. Both tenure and earnings levels are strongly and highly significantly related with increased ownership of illiquid assets.

In Appendix 2.A.5, we provide further robustness checks for these findings. Because the analysis in this section uses ratios of variables, the results can be heavily influenced by outliers and extreme values. To reduce their influence on regression results, we, therefore, winsorize data for household wealth, wealth-to-income ratios, portfolio returns, and portfolio shares at the 2.5% and the 97.5% level. In Table 2.A.2, we show that our conclusions are not influenced by such treatment of the tails of the distribution. To provide an alternative analysis that is robust to outliers and extreme values, we model the conditional median instead of the conditional mean. The results of median regressions are in line with the findings presented above. Additionally, Table 2.A.3 shows that car wealth and car loans do not drive our results and that excluding cars from household balance sheets does not alter the results. Lastly, results in Table 2.A.4 show that the results remain very similar also when we do not drop households with female household heads from the sample.

To summarize, our empirical analysis finds that employment duration and earnings levels correlate negatively with labor market risk. We exploit this relationship

Table 2.5. Asset participation

	(1) Homeowner	(2) Homeowner with a mortgage	(3) Voluntary pension plan	(4) Equity investments
	Difference between imputed types			
Overall average	0.47	0.29	0.57	0.24
Average low-risk type	0.49	0.31	0.60	0.26
Average high-risk type	0.40	0.22	0.50	0.18
Difference low-high	0.09	0.08	0.10	0.07
	(2.02)	(2.05)	(1.92)	(2.47)
	Linear regression			
Tenure	0.008^{***}	0.006^{***}	0.007^{***}	0.005^{***}
	(4.67)	(4.42)	(3.99)	(4.10)
Earnings (in 10 thousands)	0.016^{***}	0.010^{**}	0.010^{+}	0.021^{***}
	(4.10)	(2.86)	(1.93)	(4.70)
Observations	2389	2389	2389	2389
R^2	0.195	0.093	0.049	0.154

Notes: The top panel shows sample shares for homeownership without and with a positive mortgage balance, voluntary pension plan participation, and investments into equities for both risk types. The bottom panel reports OLS regression estimates, where left-hand side variables are regressed on tenure and earnings of the household head. Additional controls: age FE, education FE, and year FE. The sample includes households with male household heads between ages 20 and 60. Robust standard errors adjusted for the survey design with multiple imputations. Standard errors for the difference in means between risk types obtained using bootstrap. Data from the HFCS for Germany.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

and use tenure and earnings level as proxies for labor market risk and explore the relationship between differences in earnings risk and wealth accumulation and allocation across asset classes and differences in returns on wealth. We find that lower labor market risk is associated with higher wealth accumulation, higher portfolio returns, and a wealth allocation that is tilted towards more illiquid assets. Hence, we provide evidence for a double dividend of low labor market risk. First, we have documented in the last section that lower labor market risk leads to higher returns on human capital in the form of higher wages and earnings. Second, we document in this section that lower labor market risk leads to higher returns on financial wealth as low labor market risk allows households to invest a larger share of their portfolio into illiquid high return assets.

2.4 Life-cycle choices of heterogeneous risk types

The fact that high-risk households accumulate less wealth is surprising at a first glance given that they would have a higher precautionary motive to accumulate wealth. However, the fact that high-risk episodes oftentimes come in conjunction with low earnings can possibly explain this surprising fact. To check whether the differences in risk are in line with forward-looking behavior of households, we build a life-cycle model of agents who are heterogeneous with respect to their labor earnings risk in this section. Risk heterogeneity in this model matters for portfolio choices, and thereby affects wealth accumulation over the life cycle. We calibrate the model, inform it by our empirical results, and use it to obtain quantitative results for the consequences of risk heterogeneity.

2.4.1 The model

Our model is an incomplete markets life-cycle model with portfolio choice. Agents in the model are risk-averse and maximize their expected life-time utility, applying a discount factor β. The life cycle consists of J periods, indexed by $j = 1, \ldots, J$. We distinguish two phases of the life cycle, a working phase $j = 1, \ldots, J_{ret} - 1$, where agents work and receive stochastic labor income, and a retirement phase $j = J_{ret}, \ldots, J$, where they receive retirement benefits. Agents in the model make investment decisions to self-insure against idiosyncratic income fluctuations and to smooth consumption over the life cycle.

In every period of their life cycle agents can invest in two assets. These assets differ in terms of their return, risk, and liquidity. The first asset is a safe asset, denoted by a, with a risk-free return π, which is perfectly liquid in the sense that it can be traded without any frictions. The other asset is an illiquid asset, denoted by k. The valued stock of the illiquid asset evolves stochastically, thereby driving its risky return which fluctuates around its mean return r. The illiquid asset is subject to a trading friction. Agents can invest in this asset only by committing to a saving plan d, which can be changed at a utility cost κ. Borrowing is not allowed in either asset. The return on the liquid asset is lower than the mean return on the illiquid asset, $\pi < r$, thus inducing a trade-off between return and liquidity.

The higher average rate of return makes the illiquid asset an attractive vehicle for accumulating wealth over the life cycle. However, holdings of the liquid asset are crucial for smoothing consumption in the presence of earnings risk. The extent and structure of earnings risk will therefore affect how agents trade off the benefits from each of the assets available. The resulting portfolio composition will thus be dependent on risk types of agents.

Agents face earnings processes that are distinguished by risk types, whose calibration is explained in detail below. In the following formal description of the life-cycle decision problem in its recursive form $y(y_{p,j}, y_{\tau,j}, j)$ denotes earnings, allowed to be dependent on a persistent state $y_{p,j}$, a transitory state $y_{\tau,j}$, and on age j. This formulation accommodates all processes calibrated for the quantitative analysis.

2.4.1.1 Recursive formulation

We now formulate the agent's decision problem recursively. The state of an agent is described by the age j, the liquid asset position a_j, the illiquid asset position k_j, the current saving plan d_j, and the income states $y_{p,j}$ and $y_{\tau,j}$.

The decision problem is framed as featuring a discrete choice between the two branches of non-adjustment and adjustment of the saving plan. On each of these two branches, the decision problem then specializes in determining the appropriate financial decisions, and thereby the corresponding level of consumption.

Conditional on the discrete choice of not adjusting the saving plan, the recursive problem is the following:

$$\begin{aligned} V_{N,j}(a_j,k_j,d_j,y_{p,j},y_{\tau,j}) &= \max_{(c_j,a_{j+1})} u(c_j) + \tilde{\beta}_j \mathbb{E}_{p,\tau} V_{j+1}(a_{j+1},k_{j+1},d_{j+1},y_{p,j+1},y_{\tau,j+1}) \\ \text{s.t.} \quad & c_j + a_{j+1} = (1+\pi)a_j + rk_j + y(y_{p,j},y_{\tau,j},j) - d_{j+1} \\ & k_{j+1} = (k_j + d_{j+1})\xi_{j+1}, \quad d_{j+1} = d_j \\ & a_{j+1} \geq 0 \end{aligned} \tag{2.3}$$

Non-adjustment of the saving plan, leaving the state of the saving plan unchanged, $d_{j+1} = d_j$, is a feasible choice only if the implied illiquid asset position is not negative, $k_{j+1} \geq 0$. In this formulation $u(c_j)$ is utility from current-period consumption and $\tilde{\beta}_j$ is the effective discount factor, i.e. the product of the discount factor β and the probability of surviving from age j to age $j+1$.

Conditional on the discrete choice of adjusting the saving plan, the recursive problem is as follows:

$$\begin{aligned} V_{A,j}(a_j,k_j,d_j,y_{p,j},y_{\tau,j}) &= \max_{(c_j,a_{j+1},d_{j+1})} u(c_j) + \tilde{\beta}_j \mathbb{E}_{p,\tau} V_{j+1}(a_{j+1},k_{j+1},d_{j+1},y_{p,j+1},y_{\tau,j+1}) \\ \text{s.t.} \quad & c_j + a_{j+1} + d_{j+1} = (1+\pi)a_j + rk_j + y(y_{p,j},y_{\tau,j},j) \\ & k_{j+1} = (k_j + d_{j+1})\xi_{j+1} \\ & a_{j+1} \geq 0, \quad k_{j+1} \geq 0 \end{aligned} \tag{2.4}$$

The choice d_{j+1} of a flow for the saving plan has an immediate impact on the budget constraint relevant for financing consumption from liquid resources. In case of future non-adjustment of the saving plan, the same impact will remain present, as long as the amount committed to this per-period flow is unchanged.

The state-dependent value of the problem, which takes into account both branches, is obtained by considering the following maximization.

$$\begin{aligned} V_j(a_j,k_j,d_j,y_{p,j},y_{\tau,j}) \;=\; \mathbb{E}_\zeta \max\{V_{N,j}(a_j,k_j,d_j,y_{p,j},y_{\tau,j}) &\quad + \zeta_{N,j}, \\ V_{A,j}(a_j,k_j,d_j,y_{p,j},y_{\tau,j}) - \kappa &+ \zeta_{A,j}\} \end{aligned} \tag{2.5}$$

For the discrete choice of adjustment, the utility cost parameter κ is taken into consideration. We allow for random components in the valuation of discrete choice options. Using the assumption that these random components are distributed according to an extreme-value (type I) distribution, the relevant expectation is obtained by the familiar log-sum formula for such a discrete choice problem.

2.4.2 Calibration

We calibrate the model to German data. We set the model period to correspond to a year and assume a constant relative risk aversion utility function with a risk aversion parameter $\gamma = 2$. Workers work for 40 years and then spend (subject to uncertain survival) up to 25 years in retirement.

In the working phase agents face a stochastic process for gross log-income y (see eq. 2.6 below), which is composed of a deterministic age profile $\bar{y}_{it}$ and a stochastic part $\tilde{y}_{it}$, where i denotes the worker and t is time. The stochastic part of income consists of a persistent AR(1) component p_{it} and a transitory component η_{it}. The innovations for the persistent component and the transitory component are drawn from normal distributions with variances σ^2_ε and σ^2_η. Differences in labor market risk in the model are reflected by heterogeneity in the distribution of income shocks, as captured by these variances.[14]

$$\begin{aligned} y_{it} &= \bar{y}_{it} + \tilde{y}_{it}, & \tilde{y}_{it} &= p_{it} + \eta_{it}, & \eta_{it} &\sim N(\mu_\eta, \sigma^2_{\eta,it}) \\ & & p_{it} &= \rho p_{it-1} + \varepsilon_{it}, & \varepsilon_{it} &\sim N(\mu_{\varepsilon,it}, \sigma^2_{\varepsilon,it}) \end{aligned} \tag{2.6}$$

We calibrate the characteristics of the stochastic income process to the empirically observed earnings patterns in the German retirement accounts data. To keep the representation of heterogeneity in earnings risk tractable in the model, we model it in the same way as in Section 2.2 and split workers into two risk types. We assign the bottom eight deciles of workers based on their realized lifetime earnings volatility to the low-risk group and the top two deciles to the high-risk group. For each risk type, we estimate the deterministic life-cycle earnings profile and the corresponding stochastic earnings process. Figure 2.9 shows the resulting life-cycle profiles of average earnings, and Table 2.6 presents the estimation results for the earnings process. We provide additional details on estimation in Appendix 2.A.6.

14. Means of the distribution of innovations are always adjusted so that the variance of innovations does not affect the average level of earnings $\mathbb{E}(\exp(\tilde{y}_j)) = 1$.

Figure 2.9. Estimated life-cycle earnings profiles

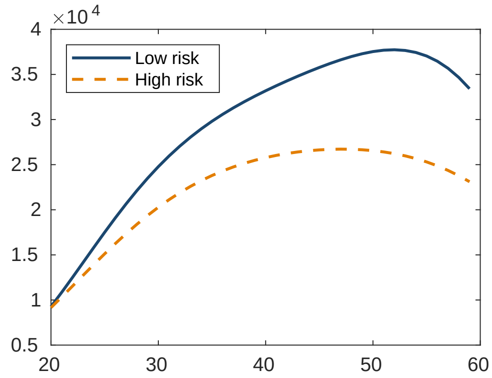

Notes: Life-cycle earnings profiles for the two risk types used in the model calibration. Earnings expressed in 2009 euros.

Table 2.6. Estimated parameter values for the earnings processes

Parameter	Pooled sample	Low-risk type	High-risk type
ρ	0.940 (0.013)	0.948 (0.018)	0.938 (0.021)
σ_ε^2	0.006 (0.001)	0.003 (0.000)	0.018 (0.003)
σ_η^2	0.014 (0.001)	0.006 (0.000)	0.059 (0.003)

Notes: Parameters of the earnings process estimated using a minimum-distance estimator as described in Appendix 2.A.6. Bootstrapped standard errors in parenthesis.

Reflecting the results presented in Section 2.2, the two risk types differ substantially in terms of earnings levels and the riskiness of their earnings processes. Throughout the life cycle, earnings of the low-risk workers are higher than earnings of workers belonging to the high-risk type. The high-risk type faces an earnings process with higher variances for both types of shocks. The variance of shocks to the persistent component is six times larger for the high-risk type than for the low-risk type, and the ratio of variances of transitory shocks is approximately ten.

When numerically solving the model, we approximate the estimated earnings processes by a discrete Markov chain using the method proposed by Tauchen (1986). As discussed in Section 2.2, the evidence presented in Table 2.1 indicates that risk types are persistent but not permanent. We model this by stochastic transitions. The probability of switching between risk types depends on the magnitude of the realized change in the persistent component of earnings. We assume that for workers of the low-risk type the probability of switching to the high-risk type is highest when they experience a very negative persistent earnings shock. The opposite holds for workers of the high-risk type, who are most likely to switch to the low-risk type when they experience a very positive persistent earnings shock. We calibrate the probability of switching between risk types to reproduce the empirically observed persistence in earnings volatility shown in Table 2.1 (first column). The resulting transition probabilities are shown in Figure 2.10.

Figure 2.10. Probability of switching between risk types

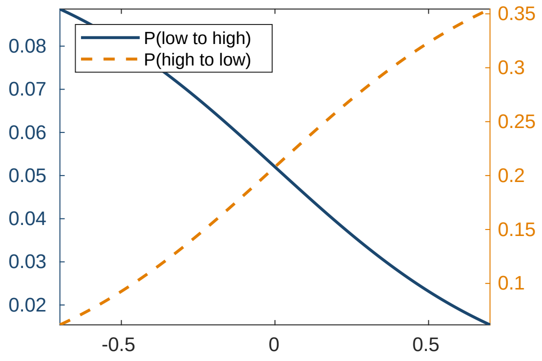

Notes: The probability of switching between risk types as a function of the realized change in the persistent component of earnings (in log points).

German retirement accounts provide only information on earnings before taxes. To convert gross earnings to net earnings we use the parameterization of the tax and social security system as documented and applied for Germany by Hintermaier and Koeniger (2018). We apply their version of a tax function, which is based on information in the OECD tax database on tax exemptions, tax rates and social security contributions at different levels of earnings. We follow their calculation of retirement benefits, which considers a pay-as-you-go component of the pension systems. This calculation uses information on the adjustment factor for pre-retirement earnings (the valorization rate), set to 1%, the number of earning years considered for the determination of retirement benefits, set to 35 years, and the applicable net-replacement rates at different levels of net earnings. The approximation of pension benefits is based on the last pre-retirement income state, considering the average working-age earnings according to the distribution of paths of earnings that lead to the relevant terminal income state.

The returns on both assets are calibrated using the evidence from the HFCS data. We use the results in Table 2.4 and adjust them to real returns. Inflation averaged 1% in the HFCS sample used. Therefore, we set the return on the liquid asset to -0.3% and the return on the illiquid asset to 5.0%. We use the historical housing returns data for Germany from Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019b) to calibrate the volatility of the return on the illiquid asset. We find that the standard deviation of the detrended housing return between 1948 and 2014 equals 2.5%.

The remaining preference and utility cost parameters are estimated within the model using a simulated method of moments approach. We use the empirically ob-

Table 2.7. Overview of model parameters

Parameter	Value		Description
γ	2.000		Relative risk aversion
π	-0.003		Liquid asset return
r	0.050		Illiquid asset return
σ_r	0.025		Volatility of the illiquid return
	With transitions	Without transitions	
β	0.961	0.956	Discount factor
κ	2.822	2.715	Adjustment cost
σ_ε	0.905	0.961	Randomness of adj. cost

served patterns of wealth accumulation over the life cycle to pin down the value of the discount factor β and the life-cycle evolution of the portfolio composition to estimate the parameters κ and σ_ζ that govern the portfolio adjustment frictions. Following the approach in Section 2.3, we construct age profiles for the wealth-to-income ratio, the portfolio share of illiquid assets, and the realized return on wealth from the HFCS data for both risk types. Using these empirical targets, we estimate the parameters for two setups: in our baseline setup we allow switching between risk types, while in the alternative case workers have fixed risk types. The resulting estimated parameters together with the other calibrated parameters are shown in Table 2.7. In the case of our baseline calibration which allows for transitions between risk types, the annual discount factor β is estimated at 0.96. We find that the model requires a non-negligible degree of adjustment frictions to match the empirical data. The adjustment cost κ is estimated at 2.82 and the standard deviation of the adjustment cost is estimated at 0.91.

2.4.3 Results

The empirical findings presented in Section 2.3 indicate that the observed patterns of wealth accumulations and portfolio composition are strongly associated with differences in labor market risk across workers. Our structural model allows us to study if model-implied consequences of earnings-risk heterogeneity are consistent with the observed empirical findings. We find that our model generates results that are consistent with the empirical evidence, but also that this finding crucially depends on the way we model workers' transitions between risk types.

We proceed by considering two scenarios. In the first scenario, our baseline, worker risk types are persistent but not fully fixed. Workers enter the labor market either as a low or a high-risk type, but they face a probability to switch between risk types throughout the working life. The switching between types happens according to the probabilities shown in Figure 2.10. In the alternative case, risk types are fixed and workers belong to the same risk type throughout their working lives. We treat the model implied data the same way as we treated the social security data. Concretely, we assign an average risk type at the end of working life in the simulated data based on the same rules we used for the actual data (see Section 2.2) and allocate workers to the two risk groups using the realized volatility of their income growth rates. The bottom 80% of workers are designated as low risk and the top 20% as high risk. Based on this sample split we can then compare the empirical profiles for wealth accumulation and portfolio composition with the model profiles.

First, we examine the scenario where workers can switch between risk types. The probability of transitioning between risk types depends on the magnitude of the earnings shock workers experience. For workers of the low-risk type, the probability to switch to the high-risk type is particularly large when they experience a very negative earnings shock. Alternatively, for workers that belong to the high-risk type, their probability to switch to the low-risk type is the largest when they experience a very positive earnings shock. Figure 2.11 shows the resulting life-cycle profiles for the wealth-to-income ratio, liquid share and realized returns.

At first abstracting from differences between risk groups, we find that the estimated model is able to replicate the stylized facts on wealth accumulation and portfolio choice over the life cycle. In line with the empirical evidence, workers increase their wealth holdings as they age and the share of the illiquid wealth increases since workers allocate more wealth towards the illiquid asset as they age. Consequently also the return on wealth increases over the life cycle. In terms of quantitative results, we also find that our model matches the data well, although it implies a somewhat more convex profile of wealth accumulation than the data suggests. However, it is important to note that we are able to match the liquid share at the beginning of the working life only by endowing a quarter of workers with an initial endowment of the illiquid asset. In the absence of such an initial endowment, portfolio liquidity would be too high at the start of the working life. Since wealth is held mostly for self-insurance against idiosyncratic income fluctuations at the beginning of working life (Gourinchas and Parker, 2002; Cagetti, 2003), workers start saving in the liquid asset. Only after a sufficient stock of the liquid asset has been accumulated, they start to invest more into the illiquid asset. Consequently, the model implies that absent any initial endowment of the illiquid asset, portfolios would be composed entirely of liquid wealth at the beginning of working life.

Now turning our attention to the differences between risk groups, we find that the model can match the empirically observed patterns of differences between workers to a great extent. First, we find that high-risk workers hold more liquid portfolios, replicating the empirical finding. Higher volatility of earnings induces workers in the high-risk group to save relatively more in the liquid asset, which can be easily used

Figure 2.11. Results when workers switch between risk types

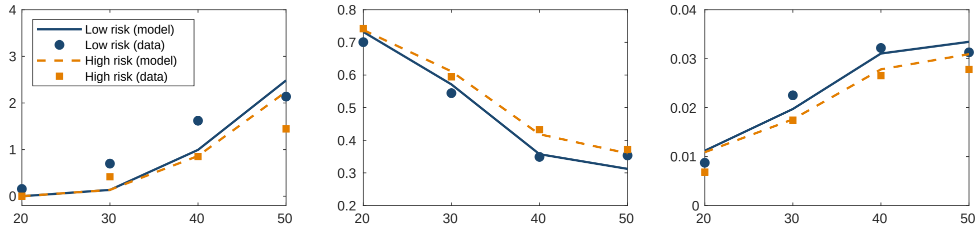

(a) Wealth-to-income **(b)** Liquid share **(c)** Realized return

Notes: Workers allocated to risk groups on the basis of realized volatility of their income growth rates. The bottom 80% of workers are classified as low-risk and the top 20% as high-risk. Wealth-to-income ratios normalized by the value of the high-risk group in the initial point. Liquid share in the initial point matched by endowing a quarter of workers with an initial stock of the illiquid asset.

to smooth income fluctuations. Asset illiquidity matters a lot for workers with high earnings fluctuations and therefore a higher stock of the liquid asset is required to optimally smooth consumption over time. Second, we also find that our model can replicate the empirical finding that low-risk workers develop higher wealth-to-income ratios than high-risk workers, although model-based differences between risk types are smaller than the empirical ones.

As an alternative, we consider a scenario where switching between risk types is not possible. Workers can belong either to the low-risk type or to the high-risk type, and the risk structure a given worker faces does not change during the working life. Figure 2.12 shows the life-cycle profiles for the wealth-to-income ratio, liquid share and realized returns for this situation. Similarly to the earlier scenario, where transitions between risk types were possible, the model is able to match the direction of the empirically observed differences in portfolio liquidity and realized returns. Most importantly, though, the results on wealth accumulation are now inconsistent with the empirical evidence. We find that with fixed worker risk types the model fails to replicate the empirical finding that low-risk workers have relatively more wealth. This version of our model would imply that workers in the high-risk group hold more wealth relative to their income throughout the working life. With a fixed assignment of workers to risk types, earnings are fully mean-reverting and therefore higher volatility of shocks induces workers to save more in order to smooth the larger deviations around the expected mean throughout the working life. Without the probability to transition between risk types, there is also no additional precautionary motive to self-insure against the idiosyncratic second-order risk for the initially low-risk type.[15]

15. In Figure 2.12 workers are split into risk groups based on the realized earnings volatility. In the situation where workers do not switch between risk groups throughout their working life, splitting

Figure 2.12. Results when workers have fixed risk types

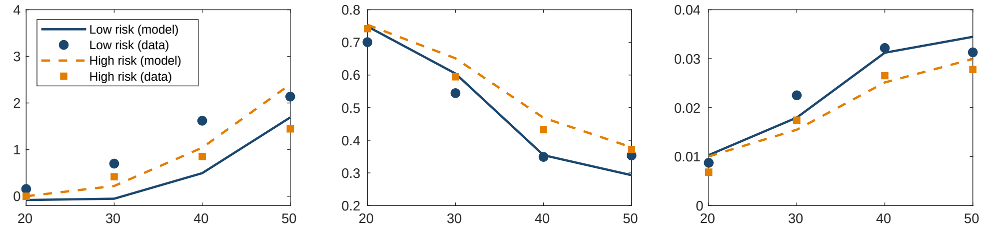

(a) Wealth-to-income **(b)** Liquid share **(c)** Realized return

Notes: Workers allocated to risk groups on the basis of realized volatility of their income growth rates. The bottom 80% of workers are classified as low-risk and the top 20% as high-risk. Wealth-to-income ratios normalized by the value of the high-risk group in the initial point. Liquid share in the initial point matched by endowing a quarter of workers with an initial stock of the illiquid asset.

2.4.4 Aggregate and welfare effects of risk heterogeneity

Given the significant differences in portfolios between high and low-risk households, we ask next what are the welfare consequences of the existence of high-risk episodes. Table 2.8 shows the results of the following experiments. First, we consider a setup in which households are born with fixed types and ask how different the aggregate capital accumulation is in this setting. Second, we look at welfare and ask how much higher/lower welfare is on average in terms of consumption equivalents for the high/low-risk type at labor market entry compared to the model with stochastic types.[16]

We find that stochastic risk increases the overall accumulation of capital because households additionally insure against the second-order risk. The average capital stock in the economy with fixed risk types is 2.2% lower than in the economy with idiosyncratic second-order risk. This leads to a high willingness to pay for removing the second-order risk for the initially low-risk households because they would need to accumulate much less for motives of self-insurance. At labor market entry, these households would be willing to give up 4.2% of baseline lifetime consumption to avoid living in an economy with second-order risk. On the other hand, the initially high-risk households benefit from the possibility of lower income risk in the future and would therefore have to be compensated by 13.3% of baseline lifetime consumption in order for them to have the same welfare in both economies.

---

into risk groups can also be performed on the basis of the initial assignment to risk types. In Appendix 2.A.8 we report the results for this case, showing that the effect from the types not being perfectly observable is negligible.

16. When we compare economies with different risk setups, we always use parameter values that were estimated for the case with stochastic risk types (see column 2 of Table 2.7).

Table 2.8. Capital accumulation and welfare (in % deviation from baseline)

	Fixed risk types	Transitory risk types
	Capital accumulation	
Overall	-2.2	0.3
	Welfare (as CEV)	
Low-risk type	4.2	-0.5
High-risk type	-13.3	2.3

Notes: Consumption-equivalent welfare gains expressed in percent of baseline consumption. Welfare evaluated at the initial model period for agents with no wealth. Situation with persistent risk types as the baseline.

Finally, we also evaluate a setup where risk types are entirely transitory, a setup akin to the Gaussian mixtures studies in Guvenen, Karahan, Ozkan, and Song (2015). When risk types are completely transitory, capital accumulation is only marginally higher than in the baseline economy with persistent risk types. Contrary to the earlier case with fixed risk types, welfare now increases for the initially high-risk type households. For them, the expected time spent as the high-risk type is reduced when risk types are entirely transitory, which results in a welfare gain of 2.3% of baseline lifetime consumption.

2.5 Conclusions

We documented large differences in earnings risk across workers. Most workers lead a relatively quiet life. At any point in time, four out of five workers face only mild earnings risks from year to year, with a standard deviation of transitory and persistent earnings shocks that is about 50% smaller than the average shock in the overall population. The remaining 20% of workers however face very large earnings risks, with 150% higher standard deviations than the low-risk types. When facing higher risks, households structure their portfolios in a more liquid way forgoing returns in order to be more easily able to tap into their savings to smooth consumption. This leads high-risk households to face not only higher earnings risks but also earn lower returns. On top, high-risk households often have also meager earnings growth on average. Compared to these high-risk households, the low-risk ones therefore lead a simple life in many dimensions.

However, we show that while risk types are persistent, they are not permanent. This is important because it introduces second order risk for households. We have shown that households react systematically to these second order risks and to changes in their risk type. On average they save more to insure against the second order risk. As a result, low-risk households would be willing to forgo a substantial amount of consumption to eliminate this second order risk.

2.A Appendix

2.A.1 Data details

2.A.1.1 Representativeness

Figure 2.A.1 shows the average age of all employed[17] by calendar year in the data. It also shows the support of the age distribution by indicating the observed minimum and maximum age. The data is constructed so that only workers who are at the time when the data is drawn are between age 30 and 67. Our first data is drawn in 2002 so that the oldest cohort is born in 1935 and the last data is drawn in 2013 so that the youngest cohort is born in 1983. This also implies that we only observe employment histories up to 2013.

Figure 2.A.1. Mean, minimum, and maximum age by year

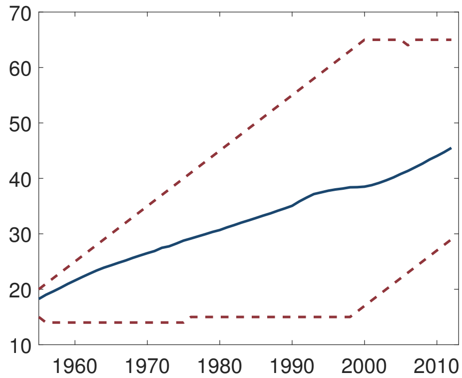

Notes: Blue solid line shows average age of all employed by calender year in the data. The red dashed lines show minimum and maximum age of employed in the data. Age is indicated on the vertical axis in years. Calender time in years is on the horizontal axis.

To analyze the representativeness of the data, we compare the data from the social security records to independent data from the 2006 Structure of Earnings Survey ("Verdienststrukturerhebung"). The SES data is a linked employer-employee cross-sectional data that covers all German workers in dependent employment. Earnings data in the SES is not capped at the social security contribution limit. To make the data comparable, we, therefore, use annual gross earnings and drop all observations above the social security contribution limit for 2006. We transform all pension points in the worker's account for a given calendar year as described in Section 2.A.1.2 and divide by the number of days employed. Figure 2.A.2 compares the log annual earnings profiles from the two data sets. We see that the two data sets are very similar in slope and shape.

17. We count employed workers with SES 13 as employed.

Figure 2.A.2. Age-earnings profile for 2006

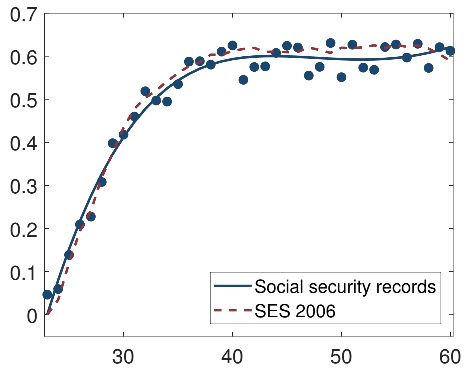

Notes: The red dashed line shows the age-earnings profile from the 2006 Structure of Earnings Survey ("Verdienststrukturerhebung"). The blue points show the raw data for the age-earnings profile from the social security records and the blue solid line shows the smoothed profile (fourth-order polynomial). All profiles have been normalized to zero at age 23. For the social security data, we use earnings from the smoothed profile.

2.A.1.2 Transform to Euro earnings

The German old-age pension scheme ("Deutsche Rentenversicherung (DRV)") uses the so-called "Durchschnittsentgelt" (average earnings) to express worker contributions. The individual accounts carry so-called "Rentenpunkte" that are the ratio of workers' earnings and average earnings. If a worker gets one point credited to her account, then she earned the "Durchschnittsentgelt" of the respective year. Importantly, the "Durchschnittsentgelt" does not correspond to the average income of all workers in the old-age pension scheme but is an earnings index that is constructed and updated since 1957. The growth rate comes from NIPA data on gross earnings per employee of all employees ("Bruttolöhne und Gehälter je Arbeitnehmer"). The initial 1957 value was based on full-time employees in 1957. Since earnings data are reported in pension points, we use the time series of "Durchschnittsentgelt" to get nominal values for earnings.

It is important to note, though, that Social security contributions in Germany are capped at an upper bound which changes over time. Figure 2.A.3a shows the historical evolution of the social security contribution limit in pension points. Due to this cap, earnings for a certain fraction of workers are top-coded in the data. In line with the observed life-cycle increase in average earnings, almost no workers are affected in the early twenties, however, at age 50 around 15% of workers are affected by censoring.

Figure 2.A.3. Social security contribution limit and the share of capped observations

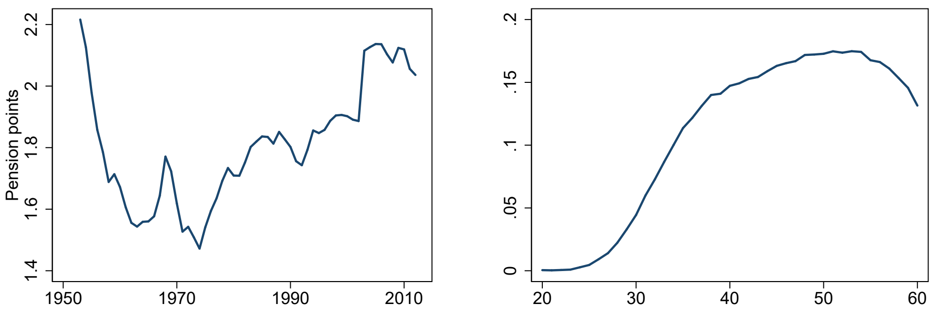

(a) Social security contribution limit

(b) Share of earnings observations

Notes: The left panel shows the social security contribution limit expressed in the number of pension points. The right panel shows the share of earnings observations that are affected by the social security contribution limit at a given age (on the horizontal axis).

2.A.2 Distribution of growth rates at age 40

Figure 2.A.4 shows the distribution of earnings and wage growth rates at age 40. The blue solid line shows the empirical distribution and the red dashed line the fitted normal distribution. In line with findings of Guvenen, Karahan, Ozkan, and Song (2015), we find that the empirical distribution exhibits a strong deviation from normality with much higher kurtosis.

Figure 2.A.4. Earnings and wage growth rate distribution at age 40

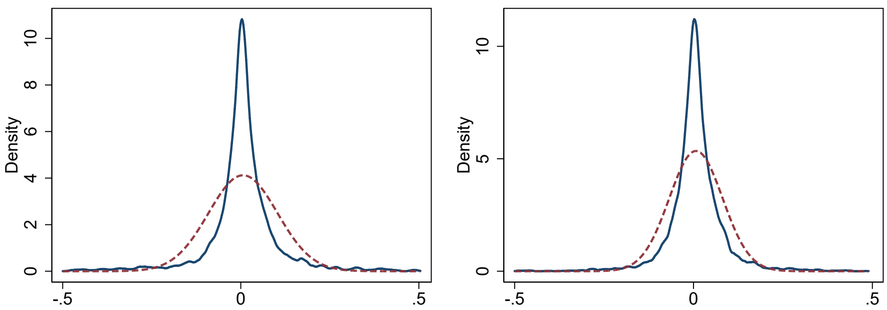

(a) Earnings

(b) Wages

Notes: The left panel shows the distribution of the annual earnings growth rate at age 40. The right panel shows the distribution of the annual wage growth rate at age 40. The blue solid line shows the kernel density estimate and the red dashed line the fitted normal distribution. For both figures, we compute log growth rates for earnings observations that are below the social security contribution limit in both years and for workers that work at least 90 days in both years. Data for cohorts born between 1935 and 1953 used. We restrict to growth rates in the range $[-0.5, 0.5]$.

2.A.3 Risk types

In Section 2.2 we group workers into risk types and explore life-cycle earnings dynamics across risk types. We compute worker-specific standard deviations of earnings growth rates during workers' prime-age working life between age 25 to 55 and assign workers belonging to the bottom 80% to the low-risk type and workers belonging to the top 20% to the high-risk type. Since no clear-cut separation between risk types can be observed in the distribution of the realized lifetime earnings growth rate volatility (see Figure 2.A.5), the results could therefore be driven by our choice of the cutoff point. However, we find that our main conclusions remain unchanged also if we use a 70-30 split instead of an 80-20 split (see Figure 2.A.6). Furthermore, our sample split is also supported by a k-means clustering algorithm since such an algorithm produces a 73-27 split. Lastly, in Figure 2.A.7 we show that persistent differences between the two risk types can also be observed when only information between age 25 and 35 is used to assign workers to risk types.

Figure 2.A.5. Lifetime earnings growth rate volatility

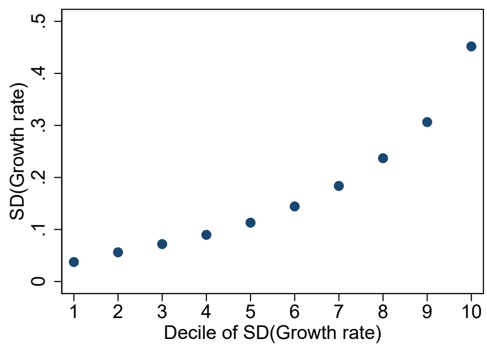

Notes: This figure shows the standard deviation of earnings growth rates during workers' prime-age working life from age 25 to 55.

Figure 2.A.6. Life-cycle profiles for worker risk types with a 70-30 sample split

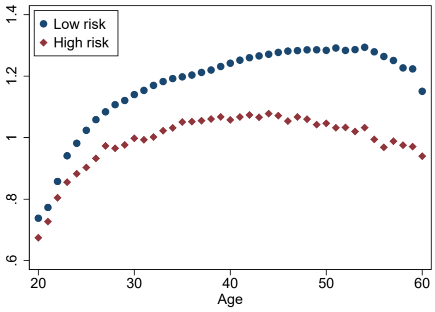

(a) Average earnings

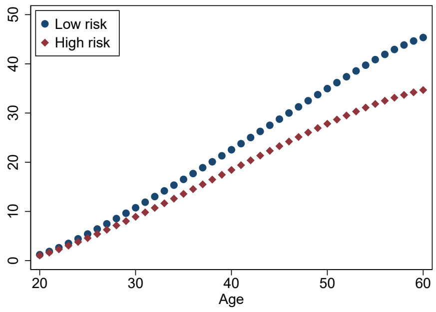

(b) Cumulative earnings

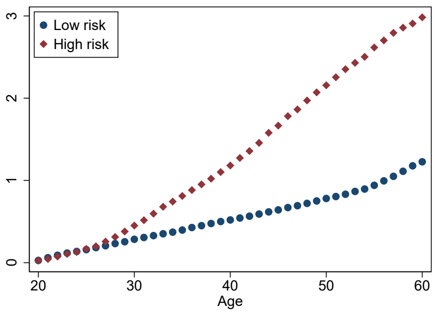

(c) Unemployment spells

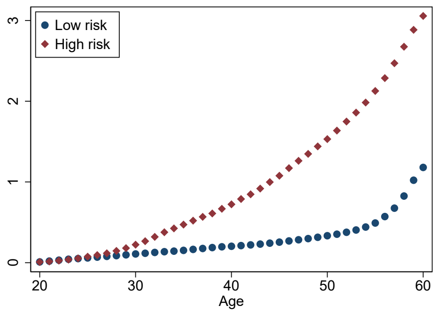

(d) Unemployment duration

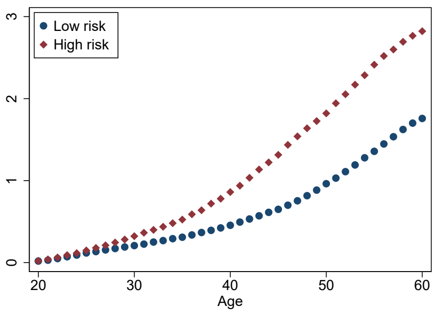

(e) Sickness spells

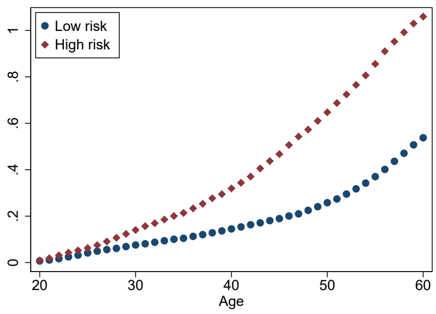

(f) Sickness duration

Notes: Risk types defined using the standard deviation of earnings growth rates between ages 25 and 55, where the bottom 70% of workers belong to the low-risk group and the top 30% to the high-risk group.

Figure 2.A.7. Life-cycle profiles for worker risk types when only information between ages 25 and 35 is used to assign types

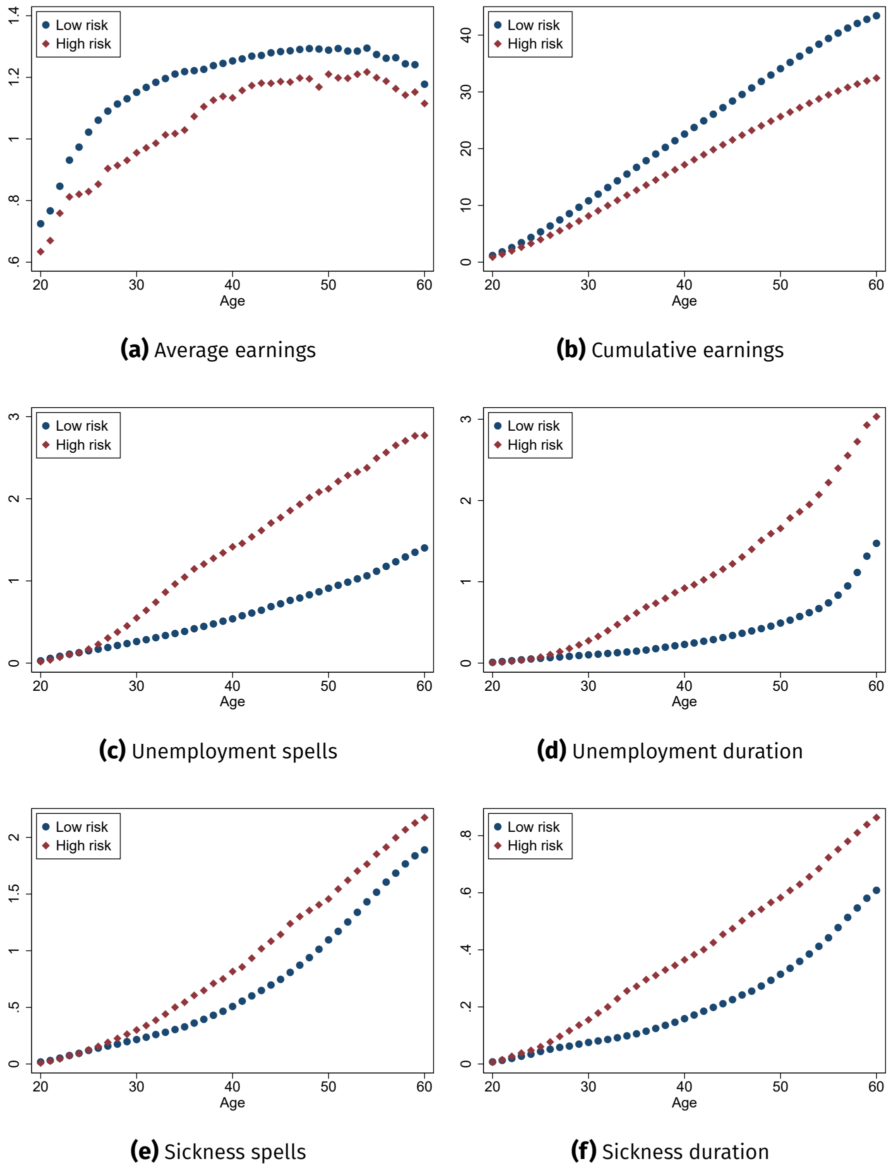

(a) Average earnings

(b) Cumulative earnings

(c) Unemployment spells

(d) Unemployment duration

(e) Sickness spells

(f) Sickness duration

Notes: Risk types defined using the standard deviation of earnings growth rates between ages 25 and 35, where the bottom 80% of workers belong to the low-risk group and the top 20% to the high-risk group.

2.A.4 Risk heterogeneity simulations

Section 2.2.2 presents simulation-based evidence to support our claim that the empirically observed patterns of earnings volatility are not consistent with a uniform risk process. In these simulations, we assume that earnings of workers are generated by a stochastic earnings process with permanent and transitory shocks. Such specification is one of the most prevalent specifications in applied work (see e.g. Meghir and Pistaferri, 2004; Heathcote, Perri, and Violante, 2010).

More specifically, we assume the following stochastic process for log earnings e

$$\begin{aligned} e_{i,j} &= p_{i,j} + \eta_{i,j} \\ p_{i,j} &= p_{i,j-1} + \varepsilon_{i,j} \end{aligned}$$

where $p_{i,j}$ denotes the permanent component, $\varepsilon_{i,j}$ the permanent innovation and $\eta_{i,j}$ the transitory innovation. Innovations are assumed to be drawn from a normal distribution with a mean of zero and variances σ^2_η and σ^2_ε.

To estimate the parameters of the earnings process we use an identification scheme based on growth rates. Using the social security data, we estimate the earnings process first on the pooled data and then also separately for each of the risk types. Table 2.A.1 presents the parameter estimates.

Table 2.A.1. Earnings processes used in the simulations

St. deviation	Pooled sample	Low-risk type	High-risk type
Permanent innovation	0.095	0.072	0.202
Transitory innovation	0.085	0.057	0.184

2.A.5 Labor market risk and household portfolios: additional results from the HFCS

In this section we provide robustness analysis for findings presented in Section 2.3.2. Table 2.A.2 presents results of median regressions which show that our main results are not driven by winsorization. Table 2.A.3 shows that car wealth and car loans do not drive our results and that excluding cars from household balance sheets does not alter the results. Results in Table 2.A.4 show that the results remain very similar also when we do not drop households with female household heads from the sample.

Table 2.A.2. Wealth, portfolio composition and returns on wealth: results from quantile regressions

	(1) Net wealth (in 1000s)	(2) WTI	(3) Illiquid share (wide) (in %)	(4) Return (in %)
Tenure	2.337***	0.061***	0.457**	0.033***
	(4.36)	(7.18)	(3.20)	(4.65)
Earnings (in 10 thousands)	23.921***	0.069***	0.905*	0.101***
	(9.08)	(5.69)	(2.35)	(6.07)
Observations	2389	2389	2389	2389
Pseudo R^2	0.139	0.075	0.116	0.066

Notes: This table reports median regression estimates of net wealth, wealth-to-income ratio, illiquid share, and portfolio return on tenure and earnings of the household head. Included controls: age FE, education FE, and year FE. The sample includes households with male household heads between ages 20 and 60. Data from the HFCS for Germany. Robust standard errors adjusted for the survey design with multiple imputations.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.A.3. Wealth, portfolio composition and returns on wealth: excluding cars and other vehicles

	(1) Net wealth (in 1000s)	(2) WTI	(3) Illiquid share (wide) (in %)	(4) Return (in %)
Tenure	2.343***	0.040***	0.537***	0.044***
	(3.59)	(4.19)	(3.77)	(3.32)
Earnings (in 10 thousands)	15.742**	-0.000	1.023***	0.107***
	(3.22)	(-0.00)	(4.21)	(4.14)
Observations	2389	2389	2383	2383
R^2	0.213	0.072	0.116	0.068

Notes: This table reports OLS regression estimates of net wealth, wealth-to-income ratios, portfolio returns, and the illiquid share on tenure and earnings of the household head. Net wealth excludes the net value of household vehicles. Included controls: age FE, education FE, and year FE. The sample includes households with male household heads between ages 20 and 60. Net wealth, wealth-to-income ratio, illiquid share and return winsorized at the 2.5 and 97.5 percentiles. Data from the HFCS for Germany. Robust standard errors adjusted for the survey design with multiple imputations.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.A.4. Wealth, portfolio composition and returns on wealth: all households

	(1) Net wealth (in 1000s)	(2) WTI	(3) Illiquid share (wide) (in %)	(4) Return (in %)
Tenure	2.254***	0.039***	0.538***	0.037***
	(4.13)	(4.76)	(4.79)	(4.18)
Earnings (in 10 thousands)	16.322***	-0.019	1.162***	0.101***
	(4.08)	(-1.01)	(3.81)	(5.12)
Observations	3459	3459	3459	3459
R^2	0.212	0.077	0.126	0.091

Notes: This table reports OLS regression estimates of net wealth, wealth-to-income ratio, illiquid share, and portfolio return on tenure and earnings of the household head. Included controls: age FE, education FE, and year FE. The sample includes households with household heads between ages 20 and 60. Net wealth, wealth-to-income ratio, illiquid share and return winsorized at the 2.5 and 97.5 percentiles. Data from the HFCS for Germany. Robust standard errors adjusted for the survey design with multiple imputations.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2.A.6 Model calibration: risk types and earnings process estimation

We calibrate the properties of the earnings process in the model to the empirical patterns from the German retirement accounts data. To capture the empirically observed heterogeneity in earnings risk, workers are split into two risk types based on the standard deviation of earnings growth rates between age 25 and 55. We assign the bottom eight deciles of workers based on their realized lifetime earnings volatility to the low-risk type and the top two deciles to the high-risk type. For both risk types, we estimate the deterministic life-cycle earnings profiles and the corresponding stochastic earnings processes for their residual earnings. When computing the earnings profiles, we transform the reported pension points into nominal earnings levels and express them in real 2009 euros using the consumer price index. Constructed in such a way, these earnings profiles also take into account the aggregate real growth in earnings that occurred during workers' working lives.

To capture the heterogeneity in earnings volatility, a stochastic earnings process is estimated on residualized earnings separately for each risk group. We residualize earnings by regressing log annual earnings on risk group-specific age fixed effects and on year fixed effects. Limited and incomplete information on education prevents us from controlling for education fixed effects, likewise, we are also not able to control for other observable demographic characteristics. For the resulting residual gross earnings y^* of a worker i of age j we assume that they follow the stochastic process in (2.A.1) and are composed of a persistent, a transitory, and a worker-fixed component:

$$\begin{aligned} y^*_{ij} &= p_{ij} + \eta_{ij} + \mu_i \\ p_{ij} &= \rho p_{ij-1} + \varepsilon_{ij} \\ \eta_{ij} &\sim N(0,\sigma^2_\eta), \quad \varepsilon_{ij} \sim N(0,\sigma^2_\varepsilon), \quad \mu_i \sim N(0,\sigma^2_\mu) \end{aligned} \tag{2.A.1}$$

where p_{ij} is the persistent component, η_{ij} is the transitory shock, and μ_i is a worker fixed effect. We estimate the parameters of the earnings process using a minimum distance estimator with the identity weighting matrix, fitting the empirical autocovariance function of residual log-earnings to the autocovariance function implied by the model. We use the variances and autocovariances up to the 5th order for ages 35-50. Standard errors are computed using a bootstrap procedure with 1000 repetitions.

2.A.7 Concentration of unemployment and sickness risk across cohorts

Figure 2.A.8 shows how the concentration of unemployment and sickness risk varies across cohorts used in our sample. We find that the degree of inequality remains broadly the same across cohorts, although there is some indication that the degree of concentration was higher for the youngest cohorts in our sample.

Figure 2.A.8. Concentration of unemployment and sickness risk across cohorts

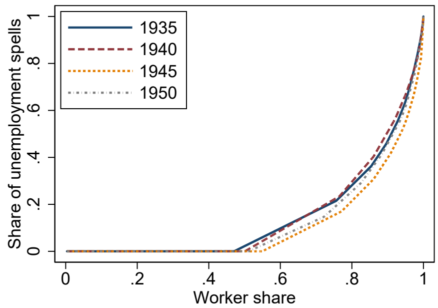

(a) Unemployment spells

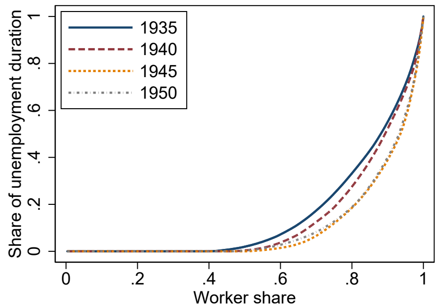

(b) Unemployment duration

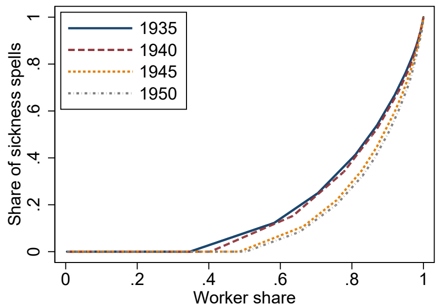

(c) Sickness spells

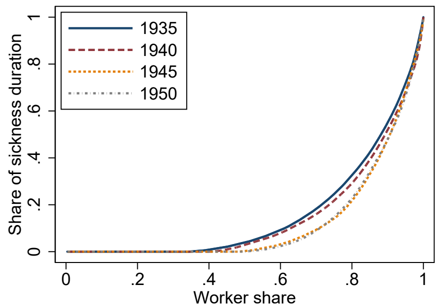

(d) Sickness duration

Notes: Lorenz curves of accumulated unemployment spells (panel a), unemployment duration (panel b), sickness spells (panel c) and sickness duration (panel d) between ages 20 and 60 for cohorts of workers born in 1935, 1940, 1945 and 1950.

2.A.8 Model simulation with fixed *observable* types

Figure 2.A.9 shows the life-cycle profiles for the scenario when risk types are fixed and observable. In this case with observable types, there is no need to infer the risk type of workers from the realized distribution of earnings growth rates. We find that the results do not change significantly, since our allocation rule based on realized earnings volatility correctly assigns workers to their initially assigned risk type in 99% of cases.

Figure 2.A.9. Results when workers have fixed risk types: risk groups based on the true assignment

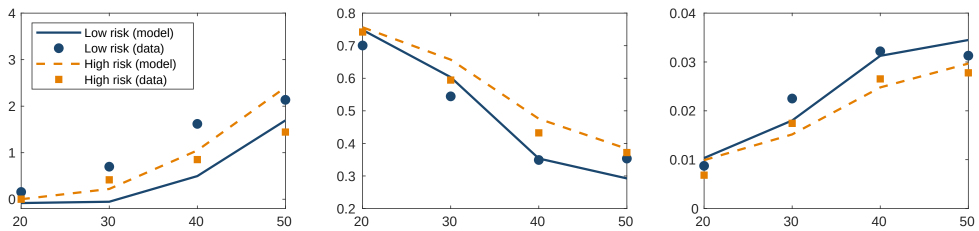

(a) Wealth-to-income **(b)** Liquid share **(c)** Realized return

Notes: Workers allocated to risk groups using their true type. Wealth-to-income ratios normalized by the value of the high-risk group in the initial point. Liquid share in the initial point matched by endowing 25% of workers with an initial stock of the illiquid asset.

References

Angerer, Xiaohong, and Pok-Sang Lam. 2009. "Income Risk and Portfolio Choice: An Empirical Study." *Journal of Finance* 64 (2): 1037–55. [73]

Antoni, Manfred, Andreas Ganzer, Philipp Vom Berge, et al. 2019. "Stichprobe der Integrierten Arbeitsmarktbiografien Regionalfile (SIAB-R) 1975-2017." Working paper. Institut für Arbeitsmarkt-und Berufsforschung (IAB). [74]

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. "Earnings and consumption dynamics: a nonlinear panel data framework." *Econometrica* 85 (3): 693–734. [73]

Bayer, Christian, and Moritz Kuhn. 2019. "Which ladder to climb? Decomposing life cycle wage dynamics." Working paper. IZA Discussion Paper. [77]

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. [73, 74, 87]

Bhutta, Neil, Jesse Bricker, Andrew C Chang, Lisa J Dettling, Sarena Goodman, Alice Henriques Volz, Joanne W Hsu, Kevin B Moore, Sarah Reber, Richard Windle, et al. 2020. "Changes in US Family Finances from 2016 to 2019: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 106 (5): 1–42. [87]

Blundell, Richard, Michael Graber, and Magne Mogstad. 2015. "Labor income dynamics and the insurance from taxes, transfers, and the family." *Journal of Public Economics* 127: 58–73. [73]

Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business & Economic Statistics* 21 (3): 339–53. [100]

Campbell, John Y. 2006. "Household Finance." *Journal of Finance* 61 (4): 1553–604. [91]

Chang, Yongsung, Jay H. Hong, Marios Karabarbounis, Yicheng Wang, and Tao Zhang. 2020. "Income Volatility and Portfolio Choices." Working paper. FRB Richmond. [73]

ECB. 2016. "The Household Finance and Consumption Survey: results from the second wave." Working paper. ECB Statistics Paper. [87]

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2018. "Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk." *Review of Economic Studies* 85 (1): 437–74. [73]

Farber, Henry S. 1999. "Mobility and stability: The dynamics of job change in labor markets." *Handbook of labor economics* 3: 2439–83. [88]

Gourinchas, Pierre-Olivier, and Jonathan A Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (1): 47–89. [100]

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1996. "Income Risk, Borrowing Constraints, and Portfolio Choice." *American Economic Review* 86 (1): 158–72. [73]

Guiso, Luigi, and Paolo Sodini. 2013. "Chapter 21 - Household Finance: An Emerging Field." In. Edited by George M. Constantinides, Milton Harris, and Rene M. Stulz. Vol. 2, Handbook of the Economics of Finance. Elsevier, 1397–532. [91]

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2015. "What do data on millions of US workers reveal about life-cycle earnings risk?" Working paper. National Bureau of Economic Research. [72–74, 77, 83, 103, 107]

Hall, Robert E. 1982. "The Importance of Lifetime Jobs in the U.S. Economy." *American Economic Review* 72 (4): 716–24. [73]

Hartung, Benjamin, Philip Jung, and Moritz Kuhn. 2016. "What hides behind the German labor market miracle? A macroeconomic analysis." Working paper. mimeo. [83]

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. [111]

Heaton, John, and Deborah Lucas. 2000. "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk." *Journal of Finance* 55 (3): 1163–98. [73]

Heaton, John, and Deborah Lucas. 2001. "Portfolio Choice in the Presence of Background Risk." *Economic Journal* 110 (460): 1–26. [73]

Hintermaier, Thomas, and Winfried Koeniger. 2018. "Differences in Euro-Area Household Finances and their Relevance for Monetary-Policy Transmission." CESifo Working Paper Series 7088. CESifo. [98]

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. 2019a. "The rate of return on everything, 1870–2015." *Quarterly Journal of Economics* 134 (3): 1225–98. [91]

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. 2019b. "The Rate of Return on Everything, 1870–2015*." *Quarterly Journal of Economics* 134 (3): 1225–98. [98]

Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2017. "Macrofinancial history and the new business cycle facts." *NBER macroeconomics annual* 31 (1): 213–63. [87]

Jung, Philip, and Moritz Kuhn. 2014. "Labour market institutions and worker flows: comparing Germany and the US." *Economic Journal* 124 (581): 1317–42. [88]

Jung, Philip, and Moritz Kuhn. 2018. "Earnings Losses and Labor Mobility Over the Life Cycle." *Journal of the European Economic Association*, (May): [73, 76, 88]

Kaplan, Greg, and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4): 1199–239. [87, 91]

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. [73, 77]

Karahan, Fatih, Serdar Ozkan, and Jae Song. 2019. "Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital." *FRB of New York Staff Report*, (908): [73]

Kuhn, Moritz, and Gašper Ploj. 2020. "Job stability, earnings dynamics, and life cycle savings." Working paper. mimeo. [73, 77, 82, 91]

Kuhn, Moritz, and Jose-Victor Rios-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. [87]

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins. 2020. "Income and wealth inequality in America, 1949–2016." *Journal of Political Economy* 128 (9): 3469–519. [87]

Low, Hamish, and Luigi Pistaferri. 2010. "Disability risk, disability insurance and life cycle behavior." Working paper. National Bureau of Economic Research. [83, 85]

Manuel, Arellano, Richard Blundell, and Bonhomme Stéphane. 2017. "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica* 85 (3): 693–734. [73]

Meghir, Costas, and Luigi Pistaferri. 2004. "Income variance dynamics and heterogeneity." *Econometrica* 72 (1): 1–32. [111]

Menzio, Guido, Irina A Telyukova, and Ludo Visschers. 2016. "Directed search over the life cycle." *Review of Economic Dynamics* 19: 38–62. [88]

Morchio, Iacopo. 2020. "Work Histories and Lifetime Unemployment." *International Economic Review* 61 (1): 321–50. [73, 84]

Mukoyama, Toshihiko, and Ayşegül Şahin. 2006. "Costs of business cycles for unskilled workers." *Journal of Monetary Economics* 53 (8): 2179–93. [73]

Palia, Darius, Yaxuan Qi, and Yangru Wu. 2014. “Heterogeneous Background Risks and Portfolio Choice: Evidence from Micro-level Data.” *Journal of Money, Credit and Banking* 46 (8): 1687–720. [73]

Schmillen, Achim, and Joachim Möller. 2012. “Distribution and determinants of lifetime unemployment.” *Labour Economics* 19 (1): 33–47. [73]

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. “Cyclical Dynamics in Idiosyncratic Labor Market Risk.” *Journal of Political Economy* 112 (3): 695–717. [73]

Tauchen, George. 1986. “Finite state markov-chain approximations to univariate and vector autoregressions.” *Economics Letters* 20 (2): 177–81. [97]

Chapter 3

Home equity, mortgage credit and firm creation: evidence from the Great Recession*

3.1 Introduction

New and young firms are a key element of a dynamic job-creating economy, as they account for a disproportionately large chunk of aggregate job creation (Haltiwanger, Jarmin, and Miranda, 2012). During the Great Recession, the U.S. experienced an unprecedented decline in firm creation (see Figure 3.1), which is often at least partially attributed to developments in the housing market. The existing evidence on the funding structure of firms shows that most new firms heavily rely on funds financed through personal balance sheets of entrepreneurs (Robb and Robinson, 2014; FED, 2016), however much less is known about the importance of home equity financing for firm creation. Such scarcity of evidence is particularly puzzling given the fact that residential property accounts for the majority of all personal wealth in the U.S. (Corradin and Popov, 2015). At the same time, the existing empirical evidence with its focus on house price changes also does not provide a clear answer to the question of how much the declining availability of home equity financing in the aftermath of the subprime mortgage crisis contributed to the observed decline in firm creation during the Great Recession.

In this paper, I study how important home equity financing is for firm creation and how much of the decline in firm and job creation during the Great Recession can be attributed to the declining availability of home equity financing. Furthermore, I

* I would like to thank Moritz Kuhn, Petr Sedláček, Marek Ignaszak, Alina Bartscher, and the participants at the BGSE Brown Bag Seminar and the BGSE Macro Reading Group for helpful comments and suggestions. I also thank Karl Schulze for sharing the data. Financial support from the DAAD, the Bonn Graduate School of Economics, and the DFG research training group 2281 "The Macroeconomics of Inequality" is gratefully acknowledged.

show how the funding structure of new firms responded to the tightening credit availability in the aftermath of the 2007 financial crisis, and how these changes helped to dampen the negative impact of the declining availability of home equity financing on firm creation.

Figure 3.1. Firm entry, exit and jobs created by entering firms during the Great Recession

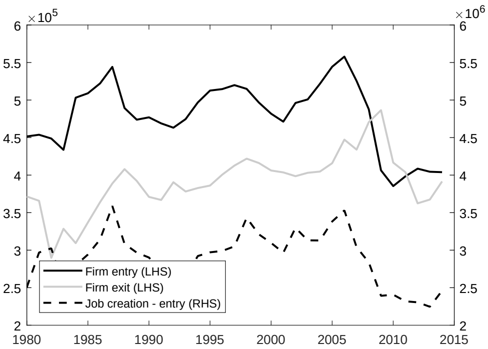

Notes: This figure plots firm entry and exit, as well as the number of jobs created by newly created firms (on the right-hand side axis) using the Business Dynamics Statistics data from the U.S. Bureau of Census.

First, I use the data from the Survey of Business Owners to provide new evidence on the importance of home equity funding for startup capital. I find that home equity is the third most used source of outside capital, behind only bank loans and credit card debt. Around 13% of new firms with employees in the SBO sample use home equity to finance startup capital, compared to 17% for credit cards and 18% for bank loans. I show that home equity is an important source of startup capital for new firms with a low number of employees and above-average capital needs and that there exists substantial heterogeneity in reliance on home equity financing across sectors.

Next, I investigate the effects of the declining availability of home equity financing on firm and job creation during the Great Recession. In the aftermath of the subprime mortgage crisis, lenders tightened their mortgage credit standards (Amromin, De Nardi, and Schulze, 2017). This constituted a significant shock to credit access and limited entrepreneurs' use of home equity to finance startup capital. Using mortgage credit availability as a proxy for the availability of home equity financing, I estimate how changes in mortgage credit availability affected the creation of new firms during the Great Recession. Using county-level data on firm creation from the Statistics of U.S. Businesses and granular loan application mortgage data from the HMDA, I find convincing evidence that the contraction of mortgage credit availability negatively affected firm creation, although the measured elasticity of firm creation with respect to mortgage credit availability appears to be relatively small. Based on a back-of-the-envelope calculation, the decline in mortgage credit availability contributed to the loss of 51,000 new firms and 297,000 jobs during the Great Recession, which is approximately 25% of the actual observed decline.

I show that the estimated effect is larger when I allow for geographical clustering of counties into local search markets for mortgage credit, and that the results cannot be explained by changes in local demand induced by changes in credit availability. To further strengthen the evidence that the estimated effect is driven by the credit channel, I identify three dimensions of heterogeneity which should affect the magnitude of the estimated effect. First, I use survey results on the use of home equity for startup capital to verify if the estimated effects differ across subsamples in such a way that could be explained by the heterogeneity in the use of home equity funding. When I split sectors according to the intensity with which they use home equity funding, I find some evidence that firm creation is more sensitive in sectors which are more reliant on home equity compared to the pooled sample estimate. When I look for heterogeneity in the estimates across size groups of new firms, I find that smaller firms are more affected by changes in mortgage credit availability than larger firms - a pattern which matches the survey results. Lastly, I split the sample according to the Saiz (2010) housing supply elasticity and find that firm creation responds much more strongly in areas with higher housing supply elasticity and where homes provide more stable collateral, corroborating Robb and Robinson (2014) who find that new firms in high-elasticity areas rely more heavily on bank debt due to more stable house prices. Finally, I use data from the Survey of Consumer Finances to provide an explanation for the relatively small magnitude of the estimated elasticity of firm creation with respect to mortgage credit availability. Evidence from the SCF suggests that although there was a significant reduction in the availability of mortgage credit, entrepreneurs were able to partially compensate for this reduction by increasing the use of other costlier and less stable types of debt.

This paper is related to several strands of existing literature. First, I contribute to the literature on the funding structure of new firms. Robb and Robinson (2014) provide detailed information on the capital structure of new firms without providing any information on the importance of home equity financing, and Schott (2015) outlines some limited evidence that home equity can be an important source of startup financing. I complement their findings by providing a more detailed analysis of the importance of home equity financing for firm creation, and show how the funding structure of new firms evolved in the aftermath of the 2007 financial crisis. Second, this paper is related to the literature which studies how changes in credit conditions affect employment and firm creation. Mondragon (2015) and Di Maggio and Kermani (2017) study how changes in aggregate mortgage credit supply affect the real economy, but do not focus on firm creation. Closely related to my research question are Jensen, Leth-Petersen, and Nanda (2014), who exploit a unique mortgage reform in Denmark to study how greater access to mortgage credit impacted entrepreneurship. Chodorow-Reich (2014) and Greenstone, Mas, and Nguyen (2020) study how other non-mortgage types of credit affect employment and business dynamics. Additional evidence on the effects of changes in credit supply on firm creation and employment is also provided by Black and Strahan

(2002) and Kerr and Nanda (2009, 2010), who study the effects of intrastate and interstate banking deregulations in the U.S. on firm dynamics. A separate, though not unrelated, literature looks at how house price changes affect entrepreneurial activity, firm creation and employment. The majority of this empirical literature points towards the conclusion that house prices have a significant effect on entrepreneurial activity and employment[1], however some authors find little evidence for such an effect[2]. While some of these studies use changes in house prices to estimate the effect of changes in the availability of housing-backed financing for new firms, I complement their findings by using a more precise measure for the availability of home equity financing.

The remainder of this paper is organized as follows. I begin in Section 3.2 by presenting the survey results on the use of home equity among new firms. Section 3.3 describes the empirical design, and Section 3.4 presents and discusses the results. Finally, Section 3.5 concludes.

3.2 Home equity financing of new firms

Existing evidence on the use of home equity financing as a source of startup capital is limited.[3] Robb and Robinson (2014) use data from the Kauffman Firm Survey to study capital structure choices that entrepreneurs make in their firms' initial year of operation. They find that newly founded firms rely heavily on formal debt financing, such as owner-backed loans, business bank loans, and business credit lines. Although they include home equity loans in their calculation of outside loans, they provide no further details on the prevalence of home equity financing. Schott (2015), on the other hand, presents some limited information on the importance of home equity financing for new firms. I complement their findings by providing additional information on how entrepreneurs use different types of debt financing, such as home equity loans, bank loans, credit card loans, etc., and show that home equity loans represent an important source of outside debt financing.

In order to shed some light on this question, I use the 2007 Survey of Business Owners, which is available as an anonymized public use micro-sample from the U.S. Census Bureau. The Survey of Business Owners provides detailed information on the characteristics of U.S. business owners. It is the only comprehensive, regularly

1. See e.g. Corradin and Popov (2015), Adelino, Schoar, and Severino (2015), Schmalz, Sraer, and Thesmar (2017), Fairlie and Krashinsky (2012), Bahaj, Foulis, Pinter, and Surico (2019), and Schott (2015).

2. See e.g. Hurst and Lusardi (2004), and Kerr, Kerr, and Nanda (2015).

3. When focusing on the already established firms, the importance of personal assets and property as a source of collateral in the business loan market is well established. See e.g. Berger and Udell (1995), Avery, Bostic, and Samolyk (1998), Ono and Uesugi (2009) and Bahaj, Foulis, and Pinter (2017).

collected source of information on selected economic and demographic characteristics for businesses and business owners by gender, ethnicity, race, etc. Included in the survey are all nonfarm businesses (employer and nonemployer) filing Internal Revenue Service tax forms as individual proprietorships, partnerships, or any type of corporation, and with receipts of $1,000 or more. The survey is conducted every five years, however, only the 2007 edition is available as a Public Use Micro Sample (PUMS). Although the data provide various details on the background of business owners, the main interest of my analysis is their choice for the source of startup capital. Available data provide evidence on the use of a specific type of funding source (extensive margin), but it does not provide any information about the distribution of total startup capital among various types (intensive margin).

Table 3.1 shows the prevalence of different types of startup capital funding over time. Approximately 65% of all firms established in 2007 required some startup funding. With 53.5%, the most common sources of startup capital are personal savings, followed by credit cards (11%), other personal assets (5.5%), bank loans (4.3%), and home equity (4.2%). A big portion of companies (35%) required no startup capital. Over time the funding structure of startup capital has undergone significant changes: while the importance of personal savings has stayed relatively the same until the Great Recession, the use of credit cards and home equity increased significantly. In the 1980s, 6.9% and 4.4% of business owners used credit cards and home equity for funding, whereas in 2005 the same percentages stood at 14.6% and 7.1%, respectively. The importance of bank loans declined significantly, from 13.3% in the 1980s to 7.5% in 2005. The onset of the Great Recession is marked by a significant decrease in the use of all sources of funding, e.g. the use of home equity declined from 7.3% to 4.2% between 2005 and 2007. On the other hand, the proportion of companies that required no startup capital increased from 22.7% to 35%.

When interpreting these statistics, it is important to note that such aggregate statistics do not reveal the true importance of various funding sources for labor market outcomes such as job creation. The vast majority of firms in the U.S. (78% in 2007) are non-employers and therefore have no impact on job creation. Furthermore, their economic impact is limited, since non-employer firms account for a very small fraction of total receipts (3.2% in 2007). Therefore, it is essential to focus on firms with positive employment when macroeconomic importance is in the focus.

Initial firm size is a substantial source of heterogeneity in the funding structure. As shown in Figure 3.2, firms with an initially low number of employees tend to use home equity as a source of funding in much bigger numbers than larger firms. More specifically, among companies that have around 10 - 20 employees, around 15% of firms use home equity as the source of initial capital, whereas for companies that have more than 100 employees this percentage drops to zero. Non-employer firms are also less likely to use home equity financing for startup capital, with less than 4 percent of them using home equity financing.

Table 3.1. Use of different sources of startup capital across time

	1980-1989	1990-1999	2000-2002	2003	2004	2005	2006	2007	Total
Personal savings	0.632	0.639	0.635	0.637	0.641	0.626	0.603	0.535	0.618
Other personal assets	0.075	0.080	0.079	0.079	0.080	0.078	0.073	0.055	0.075
Home equity	0.044	0.051	0.059	0.073	0.073	0.071	0.066	0.042	0.056
Credit cards	0.069	0.111	0.133	0.148	0.143	0.146	0.143	0.110	0.118
Gov't loan	0.007	0.007	0.006	0.005	0.006	0.005	0.004	0.003	0.005
Gov't guaranteed bank loan	0.007	0.007	0.006	0.007	0.007	0.006	0.005	0.003	0.006
Bank loan	0.133	0.108	0.090	0.089	0.084	0.075	0.067	0.043	0.091
Loan from family/friends	0.028	0.025	0.022	0.024	0.022	0.023	0.020	0.014	0.023
Venture capital	0.003	0.003	0.003	0.003	0.003	0.003	0.004	0.002	0.003
Grant	0.002	0.002	0.003	0.002	0.002	0.002	0.002	0.002	0.002
Other	0.016	0.017	0.018	0.017	0.020	0.018	0.019	0.016	0.017
Don't know	0.034	0.026	0.020	0.019	0.018	0.017	0.016	0.017	0.022
None needed	0.185	0.196	0.212	0.211	0.206	0.227	0.252	0.350	0.228
Not reported	0.012	0.009	0.008	0.008	0.007	0.007	0.006	0.007	0.008

Notes: This table shows the fraction of all newly created firms that used a particular source of funding for startup capital in a given time period. Observations are weighted using the provided survey weights. Data from the U.S. Bureau of Census, Survey of Business Owner, 2007 Public Use Microdata Sample.

Table 3.2 shows the prevalence of different sources of startup capital by employment size bins. These are chosen to match the size bins available from the Business Dynamics Statistics database, which will facilitate the interpretation of the empirical results in the later section. Looking at the table, it is clear that there are substantial differences in the funding structure of non-employer (first column) and employer firms (last column). Employer firms require more startup capital: 36.2% of non-employer firms require no capital, whereas for employer firms the same percentage is only 7.8%. As a consequence, all funding sources are more common with employer firms: 68.5% of the owners of employer firms used personal savings compared to 52.9% for non-employer firms. Institutional sources of startup capital are substantially more prevalent: in the case of employer firms 18.3% of them use bank loans, 16.8% credit cards, and 13% home equity, whereas the same percentages for non-employer firms are 3.6%, 10.8% and 3.8% respectively.

For several funding sources, an inverse U-shaped relationship between employment size and their prevalence can be observed. The use of personal savings, other personal assets, home equity, credit cards, and loans from family/friends initially increases with employment size up to a given level and then decreases. In contrast, formal institutional sources of outside capital, such as government loans, bank loans, and government-guaranteed loans display a positive relationship between employment size and their prevalence. While only 3.6% of newly created non-employer firms use bank loans as a source of startup capital, this percentage rises to almost 40% for firms with more than 20 employees. This implies that formal sources of outside capital are substantially more frequently used by firms with higher employment.

Figure 3.2. Use of home equity financing by employment size, 2007

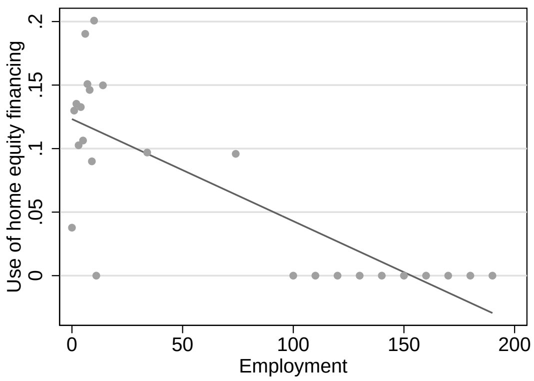

Notes: This figure plots the fraction of newly created firms in 2007 that used home equity as a source of startup capital against the (noise-infused) establishment employment. Observations are weighted using the provided survey weights. Data from the U.S. Bureau of Census, Survey of Business Owner, 2007 Public Use Microdata Sample.

Table 3.2. Use of different sources of startup capital by employment size

	0	1-4	5-9	10-19	20+	**Total**	**Employers**
Personal savings	0.529	0.699	0.606	0.606	0.540	0.535	0.685
Other personal assets	0.052	0.121	0.113	0.152	0.086	0.055	0.120
Home equity	0.038	0.129	0.142	0.155	0.087	0.042	0.130
Credit cards	0.108	0.173	0.151	0.110	0.123	0.110	0.168
Gov't loan	0.002	0.013	0.019	0.029	0.036	0.003	0.015
Gov't guaranteed bank loan	0.002	0.020	0.024	0.023	0.037	0.003	0.021
Bank loan	0.036	0.168	0.245	0.269	0.391	0.043	0.183
Loan from family/friends	0.013	0.050	0.061	0.047	0.045	0.014	0.050
Venture capital	0.002	0.008	0.002	0.019	0.014	0.002	0.008
Grant	0.002	0.002	0.002	0.002	0.006	0.002	0.002
Other	0.014	0.043	0.055	0.045	0.054	0.016	0.044
Don't know	0.016	0.021	0.044	0.039	0.082	0.017	0.025
None needed	0.362	0.079	0.066	0.082	0.058	0.350	0.078
Not reported	0.007	0.006	0.010	0.002	0.006	0.007	0.006

Notes: This table shows the fraction of all newly created firms in 2007 that used a particular source of funding for their startup capital. Columns represent different size groups, defined using the (noise-infused) establishment employment. Observations are weighted using the provided survey weights. Data from the U.S. Bureau of Census, Survey of Business Owner, 2007 Public Use Microdata Sample.

Lastly, Table 3.3 shows that a substantial sectoral heterogeneity in the use of home equity for startup capital can be observed. The difference between the sectors is almost tenfold, with only 1.5% of firms in the sector Educational Services and almost 13% of firms in Accommodation and Food Services using home equity.

3.3 Assessing the importance of mortgage credit for firm creation during the Great Recession

During the period between the beginning of 2007 and the end of 2009, households experienced a tightening of credit standards for home mortgages and other credit products. Amromin, De Nardi, and Schulze (2017) report that the median FICO credit score at mortgage origination increased from 720 to 780 during this period, while at the same time the distribution of credit scores for the U.S. population changed only slightly. As a result, a sizable dip in success rates of credit applications occurred, as shown in Figure 3.3.[4] The drop in success rates was especially pronounced for borrowers with the lowest FICO scores, who tend to be young and/or lower-income, whereas for the group with a super-prime score the drop was very limited. For sub-prime borrowers, the loan application success rate dropped

4. I thank Karl Schulze for generously sharing the Equifax time series data with me.

Table 3.3. Use of home equity financing for startup capital by sector

NAICS Code	Description	Use of home equity
72	Accommodation and Food Services	0.128
31	Manufacturing	0.072
53	Real Estate and Rental and Leasing	0.070
42	Wholesale Trade	0.067
48	Transportation and Warehousing	0.067
44	Retail Trade	0.060
52	Finance and Insurance	0.051
81	Other Services (except Public Administration)	0.049
21	Mining, Quarrying, and Oil and Gas Extraction	0.044
11	Agriculture, Forestry, Fishing and Hunting	0.043
23	Construction	0.036
71	Arts, Entertainment, and Recreation	0.033
51	Information	0.028
55	Management of Companies and Enterprises	0.027
62	Health Care and Social Assistance	0.026
56	Administrative, Support, Waste Management, Remediation Services	0.024
54	Professional, Scientific, and Technical Services	0.022
22	Utilities	0.021
61	Educational Services	0.015
99	Nonclassifiable Establishments	0.003
	Total	0.042

Notes: This table shows the sector specific fraction of all newly created firms that used home equity funding for their startup capital. Observations are weighted using the provided survey weights. Data from the U.S. Bureau of Census, Survey of Business Owner, 2007 Public Use Microdata Sample.

Figure 3.3. Success rate in obtaining credit

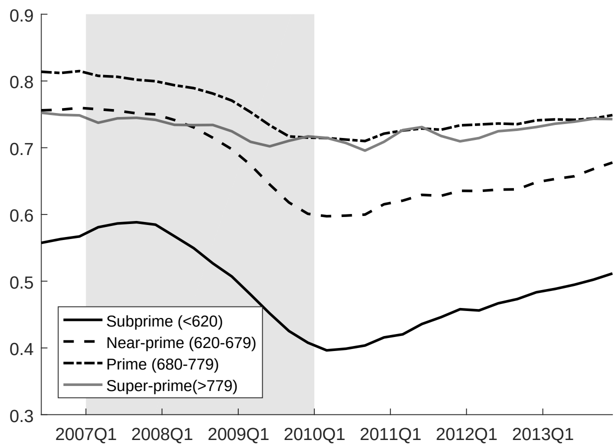

Notes: This figure plots success rates for loan applications by credit score, which includes mortgages, as well as car loans and credit cards. Figure adapted from "Household Inequality and the Consumption Response to Aggregate Real Shocks", by Amromin, De Nardi, and Schulze (2017) with the permission of the authors. Original data from the Equifax Credit Panel. The gray area indicates the period between the beginning of 2007 and the end of 2009.

by approximately 60% to 40%, and for prime borrowers it approximately declined from 80% to 70%. Overall, this evidence suggests that an important segment of the population experienced declining credit availability during the period 2007 - 2009.

The coming section assesses how such decline in the availability of mortgage credit affected aggregate firm and job creation during the Great Recession. As shown in the previous section, a significant fraction of employer firms depends on home equity for financing the startup capital. Therefore, a positive empirical relationship between mortgage credit availability and aggregate firm creation can be expected. A similar relationship is expected to also hold for the number of jobs created by these newly created firms, although changes in credit availability could also affect entrepreneurs differentially. If entrepreneurs who typically create smaller firms are substantially more affected by the decline in credit availability than entrepreneurs who create larger firms, the aggregate effect on job creation would be attenuated. However, aggregate data from the BDS does not seem to support this view, since the average number of jobs created by newly created firms only marginally increased from 5.8 in 2007 to 5.9 in 2009.

3.3.1 Econometric model

I use the following model to estimate the impact of the decline in mortgage credit availability on firm creation

$$\begin{aligned}\text{Firm Creation}_{ikt} &= \beta_0 + \beta_1 \text{Credit Availability}_{it} \\ &\quad + \boldsymbol{\beta}_2' \mathbb{1}_i + \boldsymbol{\beta}_3' \mathbb{1}_t + \boldsymbol{\beta}_4' \mathbb{1}_k + \boldsymbol{\beta}_5' \mathbb{1}_{kt} + \boldsymbol{\beta}_6' \mathbb{1}_{st} + \epsilon_{ikt},\end{aligned} \tag{3.1}$$

where firm creation in county i, sector k and time t depends on the availability of credit, as well as on other county-specific characteristics (county FE), time shocks that affect the whole economy (time FE), sector-specific characteristics (sector FE), sector-time specific shocks (sector-time FE), and state-time specific shocks (state-time FE), which allow for different economic dynamics across sectors and states. Because the primary interest lies in explaining the relationship between both variables during the Great Recession, I reformulate the model using time differences. This results in the cross-sectional specification given by equation 3.2:

$$\begin{aligned}\Delta^{2009,2007}\text{Firm Creation}_{ik} \;=\;& \gamma_0 + \gamma_1\Delta^{2009,2007}\text{Credit Availability}_i \\ &+ \boldsymbol{\gamma}_2'\mathbb{1}_k + \boldsymbol{\gamma}_3'\mathbb{1}_s + \boldsymbol{\gamma}_4'\boldsymbol{X}_i + \epsilon_{ik}\end{aligned} \tag{3.2}$$

Since time differences eliminate all time-invariant characteristics that enter equation 3.1, I am left only with $\mathbb{1}_k$, which is a set of indicator variables for sectors defined using the NAICS classification, and $\mathbb{1}_s$, which are indicators for U.S. states. These fixed effects capture what was previously captured by sector-time and state-time fixed effects. This set of variables allows me to control for heterogeneous sector and state-specific responses of changes in firm creation to the reduction of mortgage credit availability. To additionally control for likely differences in fundamental county characteristics, I also include an extensive set of pre-recession demographic and industry composition variables ($\boldsymbol{X}_i$). Where indicated in the results, I allow for an arbitrary correlation in the error term at the county level by clustering errors accordingly.

To measure changes in firm creation, I use the symmetric growth rate which has been commonly used in the literature that uses establishment-level employment microdata (Davis, Haltiwanger, and Schuh, 1996; Chodorow-Reich, 2014; Greenstone, Mas, and Nguyen, 2020).[5] Growth of the number of created firms $\Delta^{2009,2007}\text{Firm Creation}_{ik}$ for county i and NAICS sector k between 2007 and 2009 is calculated as

$$\Delta^{2009,2007}\text{Firm Creation}_{ik} \;=\; \frac{\text{Entry}_{ik,2009} - \text{Entry}_{ik,2007}}{0.5\left(\text{Entry}_{ik,2007} + \text{Entry}_{ik,2009}\right)}, \tag{3.3}$$

where $\text{Entry}_{ik,t}$ denotes the number of created firms in county i and NAICS sector k in period t.

5. The symmetric growth rate definition in equation 3.3 is a second-order approximation of the log difference growth rate around 0 and it is bounded in the range [–2,2] (Davis, Haltiwanger, and Schuh, 1996).

3.3.2 Estimating changes in credit availability

A necessary step in carrying out this empirical strategy is to construct a plausibly exogenous measure of credit availability. A naive regression of the change in firm creation on the change in total mortgage origination would most likely fail to uncover the desired effect because of possible reverse causality problems or due to the omitted variable bias. As a result, the estimated coefficient would be biased due to the failure of the orthogonality assumption $Cov(Credit_{ik}, \epsilon_{ik}) = 0$. To correctly estimate the effect of mortgage credit availability on firm creation, I, therefore, need a proxy for mortgage credit availability which is cleaned of any county-specific demand component. The approach I take to achieve this is similar to Greenstone, Mas, and Nguyen (2020), who designed a procedure to estimate credit supply shocks using business loan data from the Community Reinvestment Act. Their approach exploits the structure of the CRA dataset, which provides loan amounts for each bank and Metropolitan Statistical Area combination. This combination of bank and geographical dimensions then allows them to use an econometric procedure to estimate credit supply shocks which are orthogonal to the error term. The idea behind their approach is that different banks expand credit at different rates across geographical regions and that there is substantial heterogeneity in bank market shares across these regions. As a result, a plausibly exogenous credit supply variable can be constructed by exploiting the heterogeneity in regions' exposure to banks in the sample.

Applied to my case, this empirical strategy proceeds in the following way:

(1) For each bank in the sample, I aggregate all the originated loans up to the county level. This gives me a two-dimensional dataset where I have one observation for each bank-county combination. Having performed such aggregation for 2007 and 2009, I then calculate the growth rate of mortgage originations. I exclude refinance loans from the calculation because this segment of the mortgage market was heavily influenced by government intervention in the aftermath of the subprime crisis.[6]

(2) Next, I estimate a statistical model which decomposes the overall change in mortgage origination into two parts: the first one driven by county factors (which can be interpreted as demand factors), and the second one driven by bank factors (which can be interpreted as national-wide bank-specific supply factors):

$$\Delta^{2009,2007}\text{Mortgage Originations}_{ij} = c_i + b_j + \epsilon_{ij} \tag{3.4}$$

6. The U.S. government introduced two programs that influenced the refinancing market. The Home Affordable Refinance Program was designed to help underwater and near-underwater homeowners refinance their mortgages, whereas the Home Affordable Modification Program was designed to help financially struggling homeowners avoid foreclosure by modifying loans.

The dependent variable in equation 3.4 is the percentage change in the dollar volume of approved mortgage loans by bank j in county i between the two years. County fixed effects are denoted by c_i, and bank fixed effects by b_j.

(3) Lastly, I construct a county-level measure of mortgage credit availability using the estimated bank fixed effects b_j and banks' market shares in 2007 $m_{ij,2007}$. Market shares are calculated using the total volume of approved non-refinance mortgages by bank j in county i in the base year 2007.

$$\Delta^{2009,2007}\text{Credit Availability}_i = \sum_{j=1}^{J} m_{ij,2007} \times \hat{b}_j \tag{3.5}$$

The above procedure then produces a credit availability instrument that fulfills the orthogonality assumption, which can be now written as $Cov(\sum_{j=1}^{J} m_{ij,2007} \times \hat{b}_j, \epsilon_{ik}) = 0$. Estimated bank effects are orthogonal to the error term by construction, while on the other hand, I reduce the possibility of nonzero correlation of the error term with the market shares by using the market shares from the base period (2007).

3.3.3 Data

To carry out the empirical analysis, I combine several sources of publicly available data. Data on mortgage originations come from the Home Mortgage Disclosure Act (HMDA) database, while the data on firm creation come from two sources: Statistics of U.S. Businesses and Business Dynamics Statistics. Data on the use of home equity financing come from the Survey of Business Owners. Below is the full description of the used data.

The Home Mortgage Disclosure Act (HMDA) database is a loan-level database maintained by the Federal Financial Institutions Examination Council (FFIEC) and is constructed using the disclosure reports submitted by mortgage lenders. For every reported mortgage loan application, HMDA provides details on characteristics of the applicant (income, race, location, etc.) and the loan (type, amount, response, etc.). Although only banks and other mortgage institutions above a certain asset threshold have to report their mortgage activity, it has been estimated that HMDA data cover around 90% of all mortgage activity (Dell'Ariccia, Igan, and Laeven, 2012; Mondragon, 2015) and they can therefore be considered as a good approximation to the whole mortgage market.

The Statistics of U.S. Businesses database of the U.S. Census Bureau (USCB) provides data on enterprise creation and destruction. It contains disaggregated information on the number of newly created or destroyed (incorporated) establishments and associated job changes for every county-year-NAICS sector combination. Although SUSB provides data at a very granular level of geographic aggregation, it does not give a detailed disaggregation of firm creation on some key firm characteristics of interest like firm size.

The Business Dynamics Statistics database gives annual measures of business dynamics for the whole U.S. economy and aggregated by selected firm characteristics. It includes measures of firm openings and closings, job creation and destruction by firm size, age, and industrial sector, and several other statistics on business dynamics. The BDS is created from the Longitudinal Business Database (LBD), which is a confidential longitudinal database that allows researchers to track establishments and firms over time. The data in the BDS is provided at a more aggregated geographical level than the SUSB, as the smallest geographical unit is a Metropolitan Statistical Area. Nevertheless, the BDS data provide data on firm creation for detailed size classes of newly created firms, which is not available in the SUSB.

The Elasticity of the housing supply, which I use to measure counties' propensity to experience house price fluctuations, comes from Saiz (2010). He estimates land supply elasticities for Metropolitan Statistical Areas by processing satellite-generated data on elevation and presence of water bodies and by incorporating factors that account for endogenous restrictions on land use through zoning. This measure has been widely used in the literature (e.g. Mian, Rao, and Sufi, 2013; Mian and Sufi, 2014; Robb and Robinson, 2014, etc.) as an instrument for house price changes.

3.4 Results

Figure 3.4 shows the geographical distribution of changes in firm creation and the estimated changes in credit availability. In both cases, areas with the biggest decline are colored with stronger shades of red. The decline in firm creation is relatively evenly distributed across the whole U.S., with no strong pattern of geographical clustering. On the other hand, counties that were hit the most by the collapse in credit availability during the Great Recession are predominantly located either on the Eastern or on the Western coast, while most of the Central U.S. has experienced only a small reduction in credit availability. The central empirical question of the rest of this section is to establish how closely related these two developments are.

Table 3.4 shows the OLS estimates for the effect of changes in mortgage credit availability on changes in firm creation using the sample of county-industry observations with positive firm creation in years 2007 and 2009.[7] The coefficient estimates can be interpreted as elasticities since growth rates are used on both sides of the re-

7. Ideally, the same regression analysis would be repeated also for job creation resulting from firm entry. Unfortunately, the data from SUSB is unsuitable for this purpose since the Census Bureau infuses county and industry level employment values with a significant measure of noise (up to 5%) and values for some counties are withheld in entirety to avoid disclosing sensitive information. In unreported regression results, I find that the estimated coefficients using data of such quality are

Figure 3.4. Geographical distribution of changes in firm creation and credit availability

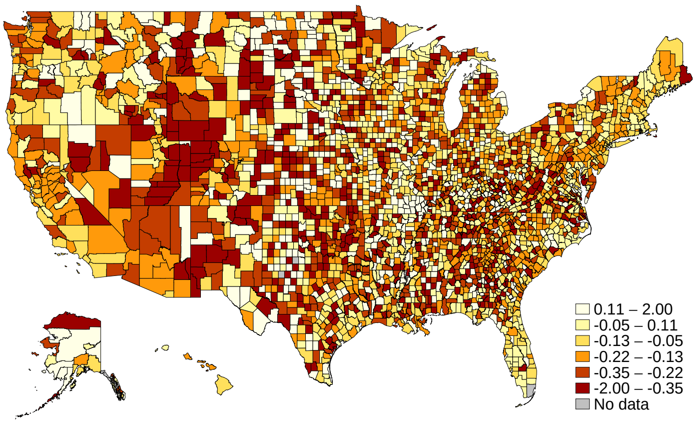

(a) Changes in firm creation

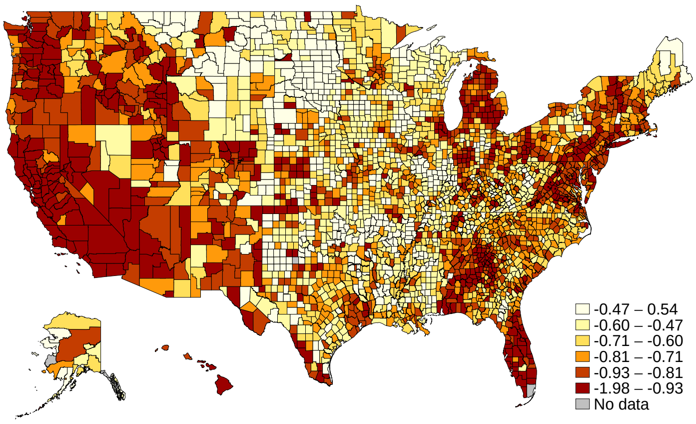

(b) Estimated changes in credit availability

Notes: The top panel plots changes in firm creation between 2007 and 2009 as defined by equation 3.3. Darker colored areas indicate a bigger decline in firm creation. Data from the Statistics of U.S. Businesses. The bottom panel plots the estimated changes in credit availability between 2007 and 2009 as defined by equation 3.5. Darker-colored areas indicate a bigger contraction in mortgage credit availability. Data from the Home Mortgage Disclosure Act dataset. Unit of observation is a county.

gression equation. The results show that the decline in mortgage credit availability negatively affected firm creation during the Great Recession, although the overall effect appears to be fairly small. In column (1) changes in firm creation are regressed on changes in credit availability without including any additional controls. The estimated coefficient of 0.057 implies that a one-percent decrease in credit availability is associated with a 0.057 percent decrease in the number of newly created firms. A clear concern with such a simplistic estimate is that the estimate is influenced by developments that are specific to certain industries and states. In columns (2)-(4) I then estimate models which allow for the possibility of sector and state-specific shocks to firm creation by including sector and/or state fixed effects. The estimated coefficient stays approximately the same when I exploit only the within-sector and within-state variation for the estimation of the elasticity. Additionally, as shown in column (5), the statistical significance of the estimate is only marginally impacted when I allow for arbitrary correlation of the error term within counties.

While my measure of credit availability is designed to be orthogonal to county-level demand and consequently also to county-level characteristics which determine local demand, a possible concern is that the estimated coefficients are still influenced by county-specific factors. I test for this concern in column (6) by including a number of pre-recession county-level control variables, including the percentage of the white population, median household income, percentage of owner-occupied housing, percentage with less than high school diploma, percentage with only a high school diploma, unemployment rate, poverty rate, and percentage of the urban population. Another concern might be that changes in credit availability interact with the existing industry structure in each county. I control for such concerns in column (7) by including the pre-recession share of total county employment that is in each of the 23 two-digit industries. Column (8) includes both demographic and industry-structure controls. In all three cases, the estimate of the coefficient is reduced by roughly a quarter, but it remains statistically significant at the ten-percent level.

Next, I consider if the estimated effect is driven by mean reversion in firm creation. On average, counties that experienced faster growth in firm creation during the expansionary period 2002-2005 suffered a greater decline in firm creation during the Great Recession. In column (10) I include pre-recession growth in firm creation to control for this concern, however, the results do not change substantially. Lastly, in column (9) I weight observations by county size, as it is sometimes common in the literature. The size of the coefficient remains similar but the statistical significance of the estimate is reduced. This seems to indicate that the relationship is more clearly identifiable in smaller counties.

Results in Table 3.4 suggest that the elasticity of firm creation with respect to credit availability lies in the range between 0.03 and 0.05. A possible reason why

---

unreliable, centered around 0, and statistically insignificant. However, I return to estimating the effect of credit availability on job creation in a later section using BDS data.

Table 3.4. Firm creation and credit availability: county-level shock

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Credit Availability	0.057***	0.042*	0.070***	0.051**	0.051**	0.038^{+}	0.033^{+}	0.037^{+}	0.035	0.046*
	(4.01)	(2.31)	(4.97)	(2.82)	(2.71)	(1.92)	(1.72)	(1.84)	(1.36)	(2.35)
State FE	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS FE	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	No	No	No	Yes	No	Yes	Yes	No
Industry composition	No	No	No	No	No	No	Yes	Yes	Yes	No
Growth 2002-05	No	No	No	No	No	No	No	No	No	Yes
Weights	No	No	No	No	No	No	No	No	Yes	No
Clustering	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32311	32311	32311	32311	32311	32283	32283	32283	32283	30314
Adjusted R^2	0.000	0.004	0.017	0.020	0.020	0.022	0.022	0.023	0.063	0.021

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equation 3.5. Standard errors are heteroscedasticity robust and where indicated clustered at the county level. Demographic controls include the pre-recession county-level percentage of the white population, median household income, percentage of owner-occupied housing, percentage with less than high school diploma, percentage with only a high school diploma, unemployment rate, poverty rate, and percentage of the urban population. Industry composition contains the pre-recession share of total county employment that is in each of the 23 two-digit industries. In column (9) observations are weighted by the total number of households in the county.

t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

these estimates are substantially lower than the survey evidence on the use of home equity financing might suggest could be the fact that the shock to credit availability is estimated at the incorrect geographical level. In Table 3.4 the credit availability shock is estimated at the county level, whereas the geographical area in which firm founders search for startup financing might in fact be much wider. As a result, the estimates could be biased towards finding a lower effect. To test this idea, I re-estimate the credit availability shock for Commuting Zones developed by Tolbert and Sizer (1996). Commuting Zones are clusters of U.S. counties that are characterized by strong within-cluster and weak between-cluster commuting ties and have been used in the studies of local labor markets (Autor and Dorn, 2013). As such, they can also be a relevant geographical level for defining local credit markets within which firm founders search for startup financing. I compute a measure of the Commuting Zone level credit availability shock by computing the mean of all county-based estimated credit availability shocks in a given Commuting Zone, as shown in equation 3.6. For a given county i that belongs to a commuting zone z_i, the CZ-level estimated credit availability shock is the average credit availability shock over all N^{z_i} counties that belong to the same Commuting Zone z_i as the given county i.

$$\Delta^{2009,2007}\text{Credit Availability}_i^{CZ} = \frac{1}{N^{z_i}} \sum_{n \in z_i} \Delta^{2009,2007}\text{Credit Availability}_n \tag{3.6}$$

Table 3.5 shows the results when the analysis of Table 3.4 is repeated with the CZ-level measure of credit availability. The estimated effect of credit availability is in the range between 0.05 and 0.09 and is substantially larger compared to earlier results. Additionally, the estimated coefficients have higher statistical significance, which seems to indicate that Commuting Zones are in fact a more appropriate geographical area for estimating the size of the shock to credit availability.

3.4.1 Economic magnitude of the effect

To put these estimates into perspective, I perform a back-of-the-envelope calculation to see how much of the overall decline in firm creation can be attributed to the decline in mortgage credit availability. For each individual county, I compute the model-predicted relative decline in firm creation by multiplying the estimated decline in credit availability (see Figure 3.4b) with the estimated elasticities from Tables 3.4 and 3.5. Next, I transform the relative decline in firm creation for each county into an absolute number of missing firms and sum up across all counties. In both Tables 3.4 and 3.5, I use the specification in column (8), meaning that I use elasticities of 0.037 (county-based shock) or 0.07 (CZ-based shock) in my computations.

This simplistic calculation shows that firm creation declined by 5.2-9.4% as a result of tighter credit availability. Since firm creation declined by 21% during the

Table 3.5. Firm creation and credit availability: commuting zone-level shock

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Credit Availability	0.081^{***}	0.079^{**}	0.097^{***}	0.089^{***}	0.089^{***}	0.073^{**}	0.065^{*}	0.070^{**}	0.053	0.080^{**}
	(4.85)	(3.25)	(5.85)	(3.67)	(3.49)	(2.72)	(2.51)	(2.63)	(1.55)	(3.01)
State FE	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS FE	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	No	No	No	Yes	No	Yes	Yes	No
Industry composition	No	No	No	No	No	No	Yes	Yes	Yes	No
Growth 2002-05	No	No	No	No	No	No	No	No	No	Yes
Weights	No	No	No	No	No	No	No	No	Yes	No
Clustering	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32311	32311	32311	32311	32311	32283	32283	32283	32283	30314
Adjusted R^2	0.001	0.004	0.017	0.021	0.021	0.022	0.023	0.023	0.063	0.021

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equations 3.5 and 3.6. Standard errors are heteroscedasticity robust and where indicated clustered at the county level. Demographic controls include the pre-recession county-level percentage of the white population, median household income, percentage of owner-occupied housing, percentage with less than high school diploma, percentage with only a high school diploma, unemployment rate, poverty rate, and percentage of the urban population. Industry composition contains the pre-recession share of total county employment that is in each of the 23 two-digit industries. In column (9) observations are weighted by the total number of households in the county.

t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

period 2007-2009, this implies that absent the decline in credit availability, the decline in firm creation would be 25-45% lower compared to the realized decline. In absolute terms, this amounts to 25,500-45,900 fewer firms per year or 51,000-90,800 fewer firms in the whole period. If I take into account that on average a newly created firm in 2007 created 5.8 jobs[8], I can calculate that during the two years between 2007 and 2009, the decline in mortgage credit availability accounts for 297,000-534,000 fewer created jobs. Compared with the total number of jobs created by entering firms (2.8 million in 2007), the magnitude of the effect of tighter credit availability seems to be relatively small but not economically insignificant. When interpreting these results, it is important to keep in mind that such a simplistic calculation does not take into account all the possible general equilibrium channels that can amplify the negative effect of the reduction in credit availability on firm creation and employment. For example, the decline in credit availability can, through general equilibrium and wealth effects, reduce household consumption and negatively influence aggregate demand, leading to less firm entry. Therefore, it is important to keep in mind that the primary aim of my estimation procedure is to establish the importance of variations in credit availability for firm creation, where the transmission channel in the center of my attention is the credit channel.

One might still wonder how reliable this back-of-the-envelope calculation is since it is done using the SUSB data which report establishment births and not firm births. To address this concern, I compare the evolution of firm and establishment births for 2007 and 2009 using the BDS data, where it is possible to discriminate between both enterprise definitions. As shown in Table 3.6, firm births account for roughly 2/3 of all establishments births. What is even more important, changes in establishment births can be almost completely explained by changes in firm births. The number of new establishments declined by 103,711 in the period 2007-2009, while the number of new firms declined by 102,315 in the same periods. The number of new establishments which are not new firms declined by only 1.396. This means that 99% of the decline in establishment births between 2007 and 2009 can be explained by the decline in firm births. Consequently, it seems that the SUSB data can be used to make valid statements about changes in firm creation.

3.4.2 Additional evidence supporting the credit channel hypothesis

In this section, I provide additional evidence which shows that changes in my measure of credit availability affect firm creation through the credit channel and not through alternative channels. A possible explanation for the positive relationship between the estimated changes in credit availability and changes in firm creation is based on the possibility that my measure of credit availability captures changes

8. As discussed earlier, the average number of jobs per newly created firm stayed relatively flat during 2007-2009.

Table 3.6. Comparison of firm and establishment-level dynamics

	2007	2009	Δ	%Δ
New firms (age 0)[a]	487,673	385,358	-102,315	-21.0%
New establishments (all)[b]	713,967	610,256	-103,711	-14.5%
New establishments ex. new firms	226,294	224,898	-1,396	-0.6%
Share of new firms (a/b)	68.3%	63.1%	98.7%	

Notes: This table compares aggregate business dynamics at the firm and the establishment level. In line with the rest of the empirical analysis, data on firm and establishment creation is shifted by one year to maximize the time coverage of the calendar year. This is done due to the fact that the reference day for data collection is March 12. Due to this timing convention, a firm is considered a new firm in the BDS in a given year if it reports positive employment in the Longitudinal Business Database for the first time on March 12 of that year. An example to clarify the time shift: data on firm creation for 2007 in this table includes firms that were created between March 12, 2007, and March 12, 2008. Data from the Business Dynamics Statistics.

in local demand due to its effect on local house prices. Mian and Sufi (2014) show that declining house prices played a significant role in the sharp decline in U.S. employment during the Great Recession. However, they also show that the effect is present only in non-tradable industries, which are more exposed to changes in local demand. Alternatively, my estimates could be influenced by how credit availability affects local construction activity. The contraction in credit availability would then be associated with contraction in the construction sector, which would imply that the estimated effect is not driven by the credit channel but by the construction demand channel. In Table 3.7 I test these hypotheses by comparing elasticities across different industry groups. In column (2) the construction sector is excluded and results remain unchanged. Also when I estimate the effect separately for the construction sector in column (3), it can be seen that the elasticity is roughly the same. Next, in column (4) I test the validity of the local demand hypothesis by excluding non-tradable industries.[9] The estimated elasticity increases in this situation, which goes against the prediction of the hypothesis since non-tradable industries should have a higher elasticity if credit availability is correlated with changes in local demand. Lastly, in column (5) I exclude non-tradable industries and the construction sector, which does not affect the estimate in a significant way. Table 3.A.2 in the Appendix shows the results when the analysis of Table 3.7 is repeated with the commuting zone based measure of credit availability. The findings are qualitatively unchanged.

To further strengthen the evidence that changes in mortgage credit availability in fact affect firm creation through the credit channel, I now use prior literature on capital structure decisions of new firms and the results from my analysis of the

9. The definition of non-tradable sectors is the same as in Mian and Sufi (2014).

Table 3.7. Industry heterogeneity in the effect of credit availability on firm creation

	(1) All industries	(2) Without construction	(3) Construction	(4) Without non-tradables	(5) Without non-tradables and construction
Δ Credit Availability	0.051**	0.050*	0.056	0.070**	0.072**
	(2.71)	(2.54)	(0.98)	(3.27)	(3.18)
State FE	Yes	Yes	Yes	Yes	Yes
NAICS FE	Yes	Yes	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes	Yes
Observations	32311	29609	2702	26150	23448
Adjusted R^2	0.020	0.021	0.030	0.014	0.014

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equation 3.5. Standard errors are heteroscedasticity robust and where indicated clustered at the county level.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

survey data on home equity financing to identify sub-samples which should exhibit bigger treatment effects from the decline in mortgage credit availability.

The first dimension that I explore is the geographical heterogeneity in the reliance on bank credit. Robb and Robinson (2014) have shown that entrepreneurs are more reliant on bank loans as a source of capital in areas where homes provide better collateral. They proxy the ability of entrepreneurs to use housing as collateral using the Saiz (2010) housing supply elasticity and argue that in areas with high elasticity of housing supply homes provide better loan collateral because the underlying home equity is less sensitive to changes in local demand for housing.[10] In high-elasticity areas housing supply expands in response to rising demand, limiting price changes, whereas the adjustment to changes in the demand in low-elasticity areas happens more through prices and less through quantities. Therefore, I expect the elasticity of firm creation with respect to mortgage credit availability to be higher in areas with higher housing supply elasticity. I test this hypothesis by splitting counties along the median of Saiz (2010) housing supply elasticity.

The second dimension that I explore is the heterogeneity in sectoral dependence on home equity for startup capital. As shown in Table 3.3, new firms in the most dependent sector are 8 times more likely to use home equity than new firms in the least dependent sectors. The most dependent sector is Accommodation and Food Services (NAICS Code 72) with 12.8% of all new firms using home equity for startup capital, while Education Services (NAICS Code 61) is the least dependent sector with only 1.5% of new firms using home equity as a source of initial capital. In line with the above results, I expect sectors that are more reliant on home equity to exhibit a higher sensitivity to the mortgage availability shock. The crucial assumption behind this test is that differences in the use of home equity across sectors reflect underlying structural differences, which make firms in some sectors more likely to use home equity compared to others. An example of such a mechanism would be sector-specific bank lending standards. However, such a hypothesis would require additional external validation, and as a result, the problem of self-selection cannot be completely ruled out.

The third dimension that I use is the heterogeneity in the likelihood with which new firms of different size classes use home equity for startup capital. As already discussed in Section 3.2 and shown in Figure 3.2, new firms with a smaller number of employees use home equity as a source of startup capital much more frequently than those with higher employment. Therefore, I expect new firms with fewer employees to be more responsive to changes in mortgage credit availability. The underlying reason for differential sensitivity lies in the fact that businesses of different sizes (and employment) require different amounts of startup capital to begin operations.

10. Adelino, Schoar, and Severino (2018) provide micro-level evidence that housing is perceived as less risky in areas with high housing supply elasticity.

Home equity financing is limited at the upper bound by the value of the owner's home, and as a result, firms above a certain level of startup capital and employment level should show little sensitivity to changes in mortgage availability. On the other hand, firms with low startup capital requirements (and low employment) should also exhibit a lower sensitivity, since home equity financing makes financial sense only above a certain minimum value due to the loan-approval costs involved. Table 3.A.1 in the Appendix provides additional evidence for this argument.

Table 3.8 contains results for the first and the second hypothesis. Results in column (2) show that areas with higher housing supply elasticity (HSE) respond more strongly to changes in credit availability. Counties with an above-median housing supply elasticity show more than twice as high elasticities (0.125) as the whole sample (0.057). The estimate for counties below the median is statistically insignificant and with an unexpected sign. Column (3) displays the results for my second hypothesis. I test this hypothesis by splitting sectors along the median of home equity dependence. Estimated elasticities are, contrary to expectations, bigger for sectors with low dependence on home equity (HE), which might be an indicator of the omitted variable bias. In column (4) I then interact the change in credit availability with both the housing supply elasticity and the dependence on home equity. Here the results are more supportive of my hypothesis since counties with an above-median housing supply elasticity and sectors with an above-median reliance on home equity show the highest estimated elasticity (0.148), whereas counties with an above-median housing supply elasticity and sectors with a below-median reliance on home equity exhibit a substantially lower elasticity (0.102). In columns (5) - (7) I include additional controls which generally increase the size of the effect for the groups with the highest effect, but the results are otherwise not substantially different. Table 3.A.3 in the Appendix shows the results when the analysis of Table 3.8 is repeated with the commuting zone level measure of credit availability. The findings are qualitatively unchanged.

To test if firm size (the third dimension of heterogeneity) affects the transmission of credit availability shocks, I estimate the baseline equation for different firm size groups. Because SUSB data do not contain any information on the size of new firms, I use data from the Business Dynamics Statistics. The BDS contains data on newly created firms by size groups, however, such disaggregation is not available at the county level. As a result, a Metropolitan Statistical Area represents a unit of observation in the following results. This new level of geographical aggregation requires also a correctly aggregated measure of credit availability shocks. I compute a measure of the MSA-level credit availability shock by computing the mean of all county-based estimated credit availability shocks in a given Metropolitan Statistical Area, as shown in equation 3.7. For a given Metropolitan Statistical Area i, the MSA-level estimated credit availability shock is the average credit availability shock over all N^i counties that belong to the same Metropolitan Statistical Area.

Table 3.8. Heterogeneous effects of credit availability on firm creation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Credit Availability	0.051**						
	(2.71)						
Δ CA × HSE high		0.125**			0.143***		
		(3.08)			(3.50)		
Δ CA × HSE low		-0.067			-0.011		
		(-1.29)			(-0.20)		
Δ CA × HE high			0.034			0.018	
			(1.51)			(0.75)	
Δ CA × HE low			0.070**			0.057*	
			(2.82)			(2.25)	
Δ CA × HE high × HSE high				0.148**			0.166***
				(3.12)			(3.45)
Δ CA × HE low × HSE high				0.102+			0.121*
				(1.84)			(2.19)
Δ CA × HE high × HSE low				-0.101			-0.046
				(-1.62)			(-0.70)
Δ CA × HE low × HSE low				-0.034			0.022
				(-0.55)			(0.33)
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	No	No	Yes	Yes	Yes
Industry composition	No	No	No	No	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32311	12204	32311	12204	12204	32283	12204
Adjusted R^2	0.020	0.105	0.064	0.105	0.108	0.066	0.108

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equation 3.5. HSE high (low) corresponds to a variable that indicates if a county has an above (below) median housing supply elasticity provided by Saiz (2010). HE high (low) corresponds to a variable that indicates if a NAICS sector has an above (below) median reliance on home equity financing as found in SBO. Standard errors are heteroscedasticity robust and where indicated clustered at the county level. Demographic controls include the pre-recession county-level percentage of the white population, median household income, percentage of owner-occupied housing, percentage with less than high school diploma, percentage with only a high school diploma, unemployment rate, poverty rate, and percentage of the urban population. Industry composition contains the pre-recession share of total county employment that is in each of the 23 two-digit industries.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

$$\Delta^{2009,2007}\text{Credit Availability}_i^{MSA} = \frac{1}{N_i}\sum_{a\in i}\Delta^{2009,2007}\text{Credit Availability}_a, \quad (3.7)$$

Estimates in Table 3.9 are in line with the expectations based on the survey results from the Survey of Business Owners. The group of new firms with a smaller number of employees respond significantly more strongly to changes in mortgage credit availability than the group with a higher number of employees (Panel A). The same also holds for the number of jobs created by new firms (Panel B). Point estimates for the effect on firm and job creation for the most responsive size group (5-9 employees) are with 0.18 more than three times the size of the effect based on the pooled sample. A concern might be that differences in estimated elasticities are driven by changes in the size composition of new firms. The available BDS data only allow me to compare the shares of each firm size group in the total population. These shares have stayed approximately the same between 2007 and 2009, meaning that there is (in aggregate) little evidence of movements between size group bins. Compared to the results in the previous section, BDS data design and small sample size do not allow me to control for sector and state-specific differences in firm and job creation, and therefore these results are of more indicative nature. Nevertheless, they give, together with the earlier results, a coherent picture that points to the conclusion that changes in mortgage credit availability do affect firm and job creation, and what is more important, that this happens through the credit channel.

3.4.3 Discussion

The results seem to indicate that the decline in mortgage credit availability had a statistically significant but economically limited effect on firm creation during the Great Recession. Estimated elasticities of firm creation with respect to credit availability are somewhat below those implied by the evidence from the Survey of Business Owners. Table 3.2 shows that 13% of employer firms[11] use home equity as a source of startup capital. If we make the extreme assumption that firms cannot substitute home equity financing for other sources of funding, this implies that the upper bound for the elasticity is 0.13. This is substantially higher than 0.037-0.07, which are the values estimated in the previous section, and as a result, such results require additional clarification. To shed some light on the smaller-than-expected elasticities, I now turn to the data from the Survey of Consumer Finances. The Federal Reserve Board conducted an SCF survey in 2007 and also an exceptional follow-up survey with 2007 participants in 2009, which gives me a unique possibility to inspect the evolution of business owners' portfolios during the Great Recession using a true panel dataset. I use the data from this panel survey to inspect the changes in the composition and the cost structure of business owners' funding. To place the

11. Both SUSB and BDS include only enterprises with non-zero employment.

Table 3.9. The effect of credit availability on firm and job creation by firm size classes: MSA level

	Panel A: Firm creation				Panel B: Job creation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1-4	5-9	10-19	20+	1-4	5-9	10-19	20+
Δ Credit Availability	0.132***	0.188*	-0.032	0.136	0.120***	0.184*	0.048	0.149
	(4.31)	(2.29)	(-0.26)	(1.27)	(3.38)	(2.17)	(0.45)	(1.00)
Observations	354	354	349	804	354	349	272	244
R^2	0.045	0.016	0.000	0.002	0.029	0.015	0.001	0.004

Notes: The dependent variable in Panel A is the change in firm creation, calculated using equation 3.3, and in Panel B the change in job creation. Credit availability is calculated using equations 3.5 and 3.7. Columns correspond to firm size groups (number of employees in a new firm). Unit of observation is Metropolitan Statistical Area. Standard errors are heteroscedasticity robust.

t statistics in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

dynamics during the Great Recession into the wider historical context, I also report the results from 2001, 2004, 2010, and 2013 (non-panel) SCF surveys.

In the following analysis, I limit myself to the sample of respondents who own and actively manage a business, which equals 12.7 (13.2) % of the whole sample in 2007 (2009).[12] Ideally, I would limit the sample only to entrepreneurs who own only newly founded companies, but sample size limitations prevent me from doing so. Nonetheless, I believe that the behavior of funding provided through the personal balance sheet of entrepreneurs should behave similarly for all firm age groups in response to a contraction in mortgage credit availability, as the one experienced during the Great Recession. Alternatively, one could argue that funding of new firms is more fragile with respect to the overall financial conditions, however, this should not cause significant problems as it only implies that the true effect is bigger than the measured one. The analysis of the overall sample would produce nonsensical results for new firms only if funding properties of older firms responded in the completely opposite direction than funding properties of new firms.

Out of this sample, 19.2% of entrepreneurs stated in 2007 that they facilitated funding for their business using their own balance sheet. This fraction slightly increased to 20.9% by 2009, which seems to imply that entrepreneurs were required to use their personal balance sheets to support their business during the financial crisis. The share of entrepreneurs who guaranteed a loan for their business declined from 45.3% to 39.3% between 2007 and 2009, while the share of those who collateralized a loan increased from 27.8% to 33.1%, leaving the share of those who did both roughly unchanged (26.9% and 27.6%). Answers to the question of how much in total they guaranteed and/or collateralized reveal some interesting facts. While the mean total value decreased substantially by approximately 28%, the median total value increased by 17%, implying that the financial crises affected the ability of entrepreneurs to borrow mainly through its effect on the far right tail of the borrowing capacity.

Looking at the wider historical dynamics, it is not possible to observe any trend in the use of personal funds to provide firm funding or in the way these funds are provided. On the other hand, the total amount of loans provided has a clear upward trend, with a cyclical decline during the Great Recession.

Such aggregate dynamics mask a substantial shift in the underlying distribution of sources of funding in the period between 2007 and 2009. Table 3.11 shows the relative shares of different types of loans used to provide firm funding. It is immediately evident that from 2007 to 2009 there was a substantial shift away from mortgage and housing-related debt towards lines of credit and credit and store debt. During this period home equity and home purchase loans completely disappeared

12. All of the reported statistics are calculated using the nonresponse-adjusted sampling weights provided in the SCF.

Table 3.10. Use of personal assets for firm funding

	2001	2004	2007	2009	2010	2013
Used personal assets to facilitate funding for own business?						
Yes	18.98	19.68	19.20	20.91	18.19	12.92
No	81.02	80.32	80.80	79.09	81.81	87.08
In what way?						
Collateralized a loan	28.05	19.51	27.78	33.14	26.05	28.26
Guaranteed a loan	48.89	55.51	45.33	39.26	45.34	50.19
Both	23.07	24.98	26.89	27.60	28.62	21.55
In total guaranteed/collateralized?						
Mean	338,065	414,721	1,010,749	732,695	713,523	742,588
Median	50,000	76,000	90,000	105,000	80,000	100,000

Notes: Sample includes SCF respondents who own and actively manage a business. Observation are weighted using the provided survey weights. Data from the Survey of Consumer Finances; data for 2007 and 2009 comes unmodified from the SCF 2007-09 Panel, while data for other years has been adjusted to be methodologically consistent with the 2007-2009 Panel definitions.

from reporting entrepreneurs' balance sheets, while mortgage debt declined by approximately 5 percentage points. On the other hand, credit card and store debt increased by 16.7 percentage points and lines of credit by 11 percentage points. These movements, together with the fact that the median value of provided funds between 2007 and 2009 increased, suggest that although there was a substantial decline in the availability of mortgage credit, entrepreneurs were able to substitute away from mortgage-related debt without much impact on their overall borrowing capacity.

A longer time span reveals some additional insights. The rise of the use of credit card and store debt during the Great Recession, when it reached around 17% in 2009, is exceptional in a historical perspective since in all other years it remained below 5%. The use of mortgage debt generally followed the dynamics of the general housing market: it almost doubled between 2001 and 2007, and afterwards declined by almost 1/3 until 2010. In 2013 it rebounded and reached a historical peak. The use of other housing-related products does not exhibit such a pattern. Lines of credit have progressively become by far the most widely used funding source: their share nearly doubled from 30% in 2001 to 58% in 2013. On the other hand, other installment loans (including various sorts of consumer loans) have declined in popularity and almost completely disappeared by 2013.

Although it seems that the Great Recession had little impact on the availability of firm funding, the changing composition of funding had important consequences for the financing costs. Table 3.12 reports the interest rates paid on different types of loans by SCF respondents who actively manage a business. The data for 2009 is,

Table 3.11. Types of loans used to provide firm funding

	2001	2004	2007	2009	2010	2013
Credit card or store debt	3.26	5.13	0.02	16.67	3.78	-
Mortgage debt	15.70	26.99	28.03	23.09	19.49	32.08
Home equity loan	9.21	8.88	9.32	-	10.02	6.94
Other home purchase loan	-	-	2.64	-	2.02	-
Loan for other real estate	13.75	4.47	5.55	10.45	0.74	2.57
Line of credit	30.32	37.49	29.99	41.48	53.75	58.26
Business loan	6.01	-	0.24	-	-	0.09
Other installment loan	20.19	8.13	24.17	8.29	10.19	0.03
Pension loan	-	8.91	0.03	-	-	-
Insurance loan	1.56	-	-	-	-	0.03
Margin loan	-	-	-	0.02	-	-
Total	100.00	100.00	100.00	100.00	100.00	100.00

Notes: This table shows the prevalence of types of loans used to provide firm financing through personal balance sheets. Observation are weighted using the provided survey weights. Data from the Survey of Consumer Finances; data for 2007 and 2009 comes unmodified from the SCF 2007-09 Panel, while data for other years has been adjusted to be methodologically consistent with the 2007-2009 Panel definitions.

unfortunately, unavailable, since the vast majority of interest rate information was not collected in the 2009 follow-up interview with 2007 participants. Nonetheless, it is immediately apparent that the shift away from mortgage debt to credit card debt importantly raised funding costs, since the interest rates on the latter are more than twice as high as the interest rates on the former. This was, on the other hand, partially offset by the relatively lower rates on lines of credit, however, the change in the composition still implies an increase in the overall cost of funding. The evolution of interest rates on lines of credit can also provide an explanation for their increasing popularity. Lines of credit have by 2013 become the cheapest source of funding which has, most likely together with their greater loan size flexibility compared to mortgage debt and home equity loans, contributed to the rise in popularity.

All in all, the presented findings suggest two key drivers behind the smaller-than-expected estimated effect of the contraction in mortgage credit availability on firm creation during the Great Recession. On the one hand, entrepreneurs were able to compensate for the declining availability of mortgage credit by adjusting their funding structure, which allowed them to mitigate the decline in the amount of available funds. On the other hand, this change implied rising funding costs since there was a substantial increase in the use of credit card and store debt, which carries higher interest rates. The combined finding that credit availability had only a limited effect on firm creation suggests that other channels, such as aggregate demand, were more important. Recent findings by Adelino, Ma, and Robinson (2017) and Decker, Mccollum, and Upton (2018), who show that the majority of net employment creation

Table 3.12. Interest rates on loans

	2001	2004	2007	2010	2013
Credit cards	14.3	12.4	13.6	14.6	15.0
Mortgage debt	7.5	6.0	6.3	5.7	4.6
Home equity loan	8.7	6.5	7.7	5.4	5.8
Line of credit	8.8	5.5	8.1	5.2	4.4
Other installment loan	12.0	14.5	11.6	11.9	10.5

Notes: This table reports interest rates, in percent, paid on different types of loans. Sample includes SCF respondents who own and actively manage a business. Observation are weighted using the provided survey weights. Data from the Survey of Consumer Finances.

in response to local demand shocks occurs through firm creation, seem to validate this view.

3.5 Conclusions

This paper uses a combination of data sources to provide new insights on how home equity financing and mortgage credit affect firm and job creation. Although the housing market and its effects on the macroeconomy have recently attracted a lot of attention, surprisingly little of it has been focused on the importance of mortgage/home equity financing for entrepreneurship. I add to this literature by documenting some new facts on the use of home equity for startup funding and by showing that there exists substantial heterogeneity across several dimensions. I proceed by showing that the decline in mortgage credit availability during the Great Recession affected firm creation, but the overall effect is limited - my back-of-the-envelope calculation shows that tighter mortgage credit access decreased firm creation by approximately 51,000 firms and reduced job creation by roughly 297,000 jobs. Furthermore, I show that this effect in fact manifests itself through the credit channel, with smaller, home equity dependent firms in areas with higher reliance on bank funding being the most affected.

Lastly, I look at the potential reasons behind the smaller-than-expected estimated effect. Using the data from the Survey of Consumer Finances I find evidence which points towards the conclusion that firms were able to mitigate the decline in mortgage credit availability by adjusting their funding structure. As the use of mortgage and home equity debt decreased during the Great Recession, firm funding became more reliant on credit card and store debt. Since these types of debt carry a higher interest rate and have a shorter maturity, the cost of capital increased and the funding structure became more fragile. Unfortunately, for data availability reasons I am not able to explore this channel in more detail. Further inquiry into the responsiveness and (non-)fragility of the funding structure of new firms with respect to the aggregate financial shocks remains to be a topic of possible future research.

3.A Appendix

Table 3.A.1. Use of home equity and employment size by total amount of startup capital

Total amount of startup capital	Use of home equity	Average number of employees
Less than $5,000	0.021	2.7
$5,000 - $9,999	0.056	2.5
$10,000 - $24,999	0.082	2.9
$25,000 - $49,999	0.171	2.8
$50,000 - $99,999	0.235	3.1
$100,000 - $249,999	0.276	3.7
$250,000 - $999,999	0.253	5.8
$1,000,000 or more	0.153	24.3

Notes: This table shows in the second column the fraction of newly created firms that used home equity as a source of funding, and in the third column the average employment. Both variables are cross-tabulated by total amount of initial startup capital. Sample includes newly created firms in 2007 with positive employment. Observation are weighted using the provided survey weights. Data from the U.S. Bureau of Census, Survey of Business Owner, 2007 Public Use Microdata Sample.

Table 3.A.2. Industry heterogeneity in the effect of credit shock on firm creation: CZ-level shock

	(1) All industries	(2) Without construction	(3) Construction	(4) Without non-tradables	(5) Without non-tradables and construction
Δ Credit Availability	0.089***	0.087**	0.089	0.106***	0.106***
	(3.49)	(3.27)	(1.18)	(3.67)	(3.49)
State FE	Yes	Yes	Yes	Yes	Yes
NAICS FE	Yes	Yes	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes	Yes
Observations	32311	29609	2702	26150	23448
Adjusted R^2	0.021	0.021	0.030	0.014	0.014

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equations 3.5 and 3.6. Standard errors are heteroscedasticity robust and where indicated clustered at the county level.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.A.3. Heterogeneous effects of credit availability on firm creation: CZ-level shock

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Credit Availability	0.089***						
	(3.49)						
Δ CA × HSE high		0.170**			0.216***		
		(2.94)			(3.65)		
Δ CA × HSE low		-0.134^{+}			-0.055		
		(-1.90)			(-0.74)		
Δ CA × HE high			0.086**			0.065*	
			(2.91)			(2.13)	
Δ CA × HE low			0.092**			0.075*	
			(2.97)			(2.36)	
Δ CA × HE high × HSE high				0.234***			0.279***
				(3.56)			(4.16)
Δ CA × HE low × HSE high				0.106			0.153*
				(1.45)			(2.08)
Δ CA × HE high × HSE low				-0.127			-0.048
				(-1.60)			(-0.57)
Δ CA × HE low × HSE low				-0.140^{+}			-0.061
				(-1.72)			(-0.72)
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	No	No	Yes	Yes	Yes
Industry composition	No	No	No	No	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32311	12204	32311	12204	12204	32283	12204
Adjusted R^2	0.021	0.105	0.064	0.105	0.108	0.067	0.108

Notes: This table reports OLS regression estimates of change in firm creation, computed using equation 3.3, on the estimated change in credit availability, computed using equations 3.5 and 3.6. HSE high (low) corresponds to a variable that indicates if a county has an above (below) median housing supply elasticity provided by Saiz (2010). HE high (low) corresponds to a variable that indicates if a NAICS sector has an above (below) median reliance on home equity financing as found in SBO. Standard errors are heteroscedasticity robust and where indicated clustered at the county level. Demographic controls include the pre-recession county-level percentage of the white population, median household income, percentage of owner-occupied housing, percentage with less than high school diploma, percentage with only a high school diploma, unemployment rate, poverty rate, and percentage of the urban population. Industry composition contains the pre-recession share of total county employment that is in each of the 23 two-digit industries.
t statistics in parentheses, $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

Adelino, Manuel, Song Ma, and David Robinson. 2017. "Firm Age, Investment Opportunities, and Job Creation." *Journal of Finance* 72 (3): 999–1038. [150]

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2015. "House prices, collateral, and self-employment." *Journal of Financial Economics* 117 (2): 288–306. [124]

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2018. "Perception of House Price Risk and Homeownership." *National Bureau of Economic Research Working Paper Series* No. 25090: [143]

Amromin, Gene, Mariacristina De Nardi, and Karl Schulze. 2017. "Household Inequality and the Consumption Response to Aggregate Real Shocks." *National Bureau of Economic Research Working Paper Series* No. 24073: [122, 128, 130]

Autor, David H, and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–97. [138]

Avery, Robert B, Raphael W Bostic, and Katherine A Samolyk. 1998. "The role of personal wealth in small business finance." *Journal of Banking & Finance* 22 (6): 1019–61. [124]

Bahaj, Saleem, Angus Foulis, and Gabor Pinter. 2017. "Home Values and Firm Behaviour." Working paper 1724. [124]

Bahaj, Saleem, Angus Foulis, Gabor Pinter, and Paolo Surico. 2019. "Employment and the Collateral Channel of Monetary Policy." Working paper. Bank of England. [124]

Berger, Allen N, and Gregory F Udell. 1995. "Relationship Lending and Lines of Credit in Small Firm Finance." *Journal of Business* 68 (3): 351–81. [124]

Black, Sandra E, and Philip E Strahan. 2002. "Entrepreneurship and Bank Credit Availability." *Journal of Finance* 57 (6): 2807–33. [123]

Chodorow-Reich, Gabriel. 2014. "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129 (1): 1–59. [123, 131]

Corradin, Stefano, and Alexander Popov. 2015. "House Prices, Home Equity Borrowing, and Entrepreneurship." *Review of Financial Studies* 28 (8): 2399–428. [121, 124]

Davis, Steven J., John Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. MIT Press. [131]

Decker, Ryan A., Meagan Mccollum, and Gregory B. Upton. 2018. "Firm Dynamics and Local Economic Shocks : Evidence from the Shale Oil and Gas Boom." [150]

Dell'Ariccia, Giovanni, Deniz Igan, and Luc Laeven. 2012. "Credit Booms and Lending Standards: Evidence from the Subprime Mortgage Market." *Journal of Money, Credit and Banking* 44 (2-3): 367–84. [133]

Di Maggio, Marco, and Amir Kermani. 2017. "Credit-Induced Boom and Bust." *Review of Financial Studies* 30 (11): 3711–58. [123]

Fairlie, Robert W., and Harry A. Krashinsky. 2012. "Liquidity constraints, household wealth, and entrepreneurship revisited." *Review of Income and Wealth* 58 (2): 279–306. [124]

FED. 2016. "Small Business Credit Survey: Report on Employer Firms." Working paper. Atlanta: Federal Reserve Bank of Atlanta. [121]

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen. 2020. "Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence from the Great Recession and "Normal" Economic Times." *American Economic Journal: Economic Policy* 12 (1): 200–25. [123, 131, 132]

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda. 2012. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95 (2): 347–61. [121]

Hurst, Erik, and Annamaria Lusardi. 2004. “Liquidity Constraints, Household Wealth, and Entrepreneurship.” *Journal of Political Economy* 112 (2): 319–47. [124]

Jensen, Thais Lærkholm, Søren Leth-Petersen, and Ramana Nanda. 2014. “Housing Collateral, Credit Constraints and Entrepreneurship - Evidence from a Mortgage Reform.” Working Paper 20583. National Bureau of Economic Research. [123]

Kerr, Sari, William R. Kerr, and Ramana Nanda. 2015. “House Money and Entrepreneurship.” Working Paper 21458. National Bureau of Economic Research. [124]

Kerr, William R., and Ramana Nanda. 2009. “Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship.” *Journal of Financial Economics* 94 (1): 124–49. [123, 124]

Kerr, William R., and Ramana Nanda. 2010. “Banking Deregulations, Financing Constraints, and Firm Entry Size.” *Journal of the European Economic Association* 8 (2-3): 582–93. [123, 124]

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump*.” *Quarterly Journal of Economics* 128 (4): 1687–726. [134]

Mian, Atif, and Amir Sufi. 2014. “What Explains the 2007-2009 Drop in Employment?” *Econometrica* 82 (6): 2197–223. [134, 141]

Mondragon, John. 2015. “Household Credit and Employment in the Great Recession.” [123, 133]

Ono, Arito, and Iichiro Uesugi. 2009. “Role of Collateral and Personal Guarantees in Relationship Lending: Evidence from Japan’s SME Loan Market.” *Journal of Money, Credit and Banking* 41 (5): 935–60. [124]

Robb, Alicia M., and David T. Robinson. 2014. “The Capital Structure Decisions of New Firms.” *Review of Financial Studies* 27 (1): 153–79. [121, 123, 124, 134, 143]

Saiz, Albert. 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics* 125 (3): 1253–96. [123, 134, 143, 145, 155]

Schmalz, Martin C, David A Sraer, and David Thesmar. 2017. “Housing Collateral and Entrepreneurship.” *Journal of Finance* 72 (1): 99–132. [124]

Schott, Immo. 2015. “Start-ups , House Prices , and the US Recovery.” [123, 124]

Tolbert, Charles M., and Molly Sizer. 1996. “U.S. Commuting Zones and Labor Market Areas: A 1990 Update.” [138]