

DISCUSSION PAPER SERIES

IZA DP No. 13887

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

Moritz Kuhn
Gašper Ploj

NOVEMBER 2020

DISCUSSION PAPER SERIES

IZA DP No. 13887

Job Stability, Earnings Dynamics, and Life-Cycle Savings

Moritz Kuhn

University of Bonn, CEPR and IZA

Gašper Ploj

University of Bonn

NOVEMBER 2020

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

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

Labor markets are characterized by large heterogeneity in job stability. Some workers hold lifetime jobs, whereas others cycle repeatedly in and out of employment. This paper explores the economic consequences of such heterogeneity. Using Survey of Consumer Finances (SCF) data, we 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 the structural model, we explore the consequences of heterogeneity in job stability at the individual and macroeconomic level. At the individual level, we 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 find welfare gains of 1.6 percent of lifetime consumption for labor market entrants from a secular decline in U.S. labor market dynamism.

JEL Classification: J64, E21, E24

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

Corresponding author:

Moritz Kuhn
Universität Bonn
Adenauerallee 24-42
53113 Bonn
Germany
E-mail: mokuhn@uni-bonn.de

* We thank Christian Bayer, Thomas Hintermaier, Philip Jung, Per Krusell, Iourii Manovskii, Kurt Mitman, Giuseppe Moscarini, and Aysegul Sahin 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.

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 et al., 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 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 et al., 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 et al., 2008), earnings losses following job displacement (Jacobson et al., 1993a), the distribution of earnings growth (Guvenen et al., 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 et al. (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 (Jung and Kuhn, 2018; Jarosch, 2015). 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 et al., 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 et al. (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 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 3. Section 4 explores the individual consequences of heterogeneity in job stability. In Section 5, we study the macroeconomic decline in U.S. labor market dynamism. Section 6 concludes.

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 (Bewley, undated; Aiyagari, 1994; Huggett, 1993) and models of labor market mobility (Mortensen and Pissarides, 1994; Burdett and Mortensen, 1998). Existing models of consumption-saving behavior or labor market dynamics treat labor market dynamics and consumption-saving choices largely as orthogonal: models of consumption-saving behavior typically consider wages as an ex-

ogenous stochastic process, and models of labor market dynamics typically abstract from human capital investment and consumption-saving decisions. Only recently, a strand of research emerged that combined models of consumption-saving and labor market behavior (Lise, 2012; Krusell et al., 2010, 2017; Hubmer, 2018; Larkin, 2019; Cajner et al., 2020). We add to this literature by exploring the consequences of heterogeneity in job stability. Our paper connects the part of the literature that focuses on macroeconomic dynamics, as in Krusell et al. (2010), with microeconomic behavior, as in Lise (2012).

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 et al. (2020) extend the model in Krusell et al. (2010, 2017) to a life-cycle setting and explore the consequences of tax changes for labor supply.

Our labor market model builds on Jung and Kuhn (2018), who develop a life-cycle search model to demonstrate that heterogeneity in job stability is key to account for earnings losses following job displacement. 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 et al. (2019) explore life-cycle earnings dynamics and document large heterogeneity in life-cycle nonemployment spells. They emphasize that this heterogeneity is key to account for life-cycle earnings dynamics. Additional evidence for heterogeneity in job stability comes from Morchio (2020), who documents large heterogeneity in unemployment within cohorts of U.S. workers.

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

Our work also relates to research on heterogeneity in earnings risk, as in Low et al. (2010) and Karahan and Ozkan (2013). Low et al. (2010) explore a model with labor market search, employment risk, and consumption-saving decisions. They abstract from heterogeneity in job stability, and earnings dynamics are predominantly governed by an exogenous stochastic productivity process. Karahan and Ozkan (2013) estimate a stochastic earnings process with age-dependent parameters and find that the variance and persistence of the process vary with age. They find that the welfare consequences of market incompleteness are substantially lower in a model with an age-varying income process compared to a model with age-invariant income risk.

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.

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 Ríos-Rull, 2016; Bricker et al., 2017; Kuhn et al., forthcoming). Besides the detailed information on household income and wealth, the SCF also offers information on household members' labor market situation. 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 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 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 provides the estimates of the empirical measures of job stability and their relationship to wealth-to-income ratios. In Figure 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 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 economi-

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

³We follow Bricker et al. (2017) and Kuhn and Ríos-Rull (2016) for the construction of these variables.

⁴Employer 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.


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

Figure 1b offers a second interesting observation. While the relationship in the raw data appears almost perfectly linear (Figure 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 to converge back to their target wealth-to-income ratio if no further job loss will occur.

Figures 1c and 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 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 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.

Figure 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.

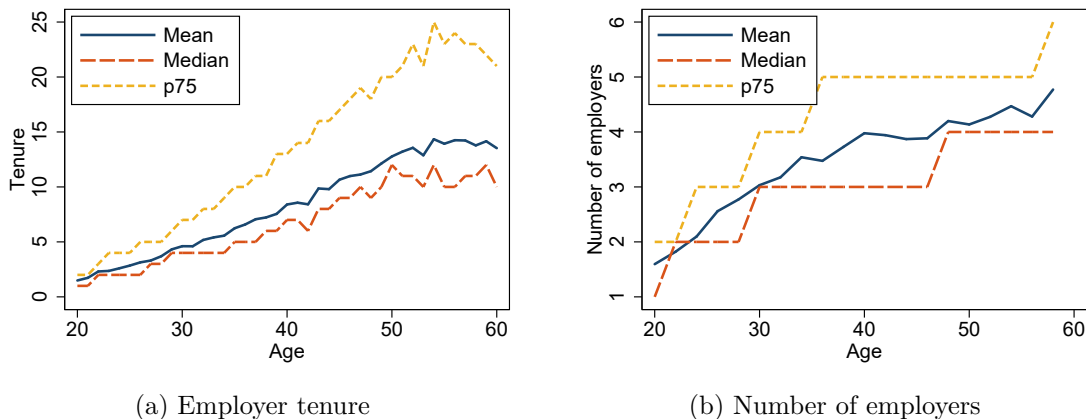
2.2 Heterogeneity in job stability

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 et al. (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 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 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 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.

The increasing life-cycle dispersion of job stability mirrors the widely studied increase in wage inequality with age (Heathcote et al., 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

Figure 2: Tenure and number of employers over the life cycle



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


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 $\mathbb{E}[T]$:

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

We derive in Appendix 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 3: Employment inequality




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

Figure 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 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 4a) and job loss due to firm closure (Figure 4b). We remove year and industry fixed effects in both cases.⁵ In Figure 4a, we observe large heterogeneity in job loss across employers, with the least stable employer having job loss rates that are twice as large as the most stable employers. Such differences in observed job loss could be the result of differences in job-to-job transitions or worker quits. Figure 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.

Figure 4: Heterogeneity in job destruction rates by firm age



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

⁵Appendix 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.

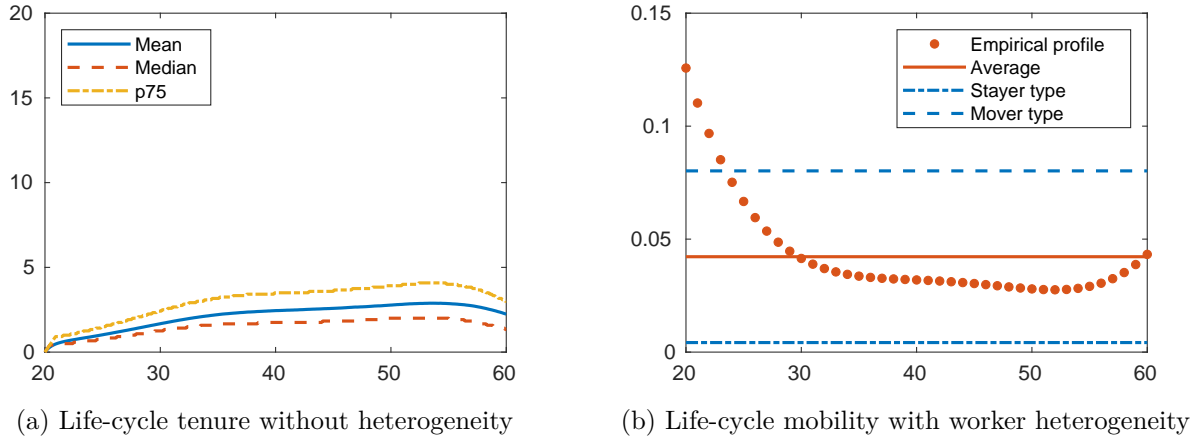
We provide two Monte Carlo simulations to revisit the extent of heterogeneity in job stability and its sources. Figure 5a revisits the extent of heterogeneity and shows a simulation of the tenure distribution for a representative-worker case where we only feed in age heterogeneity in job stability but 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 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 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 5b revisits the question on the sources of heterogeneity based on a second Monte Carlo experiment. It considers a stylized case with workers of a mover type and a stayer type. Workers of the mover type have low job stability, whereas workers of the stayer type have high job stability. For simplicity, we assume that both groups are of equal size. The life-cycle profiles of transitions into nonemployment of the two types are shown as a dashed blue line. By construction, the transition rates are a fixed worker characteristic and do not change over the life cycle, thereby resulting in flat age profiles. The average profile (solid red line) corresponds to the unconditional average of the transition rates, and the red dots show the empirical profile. The average profile and its empirical counterpart show a life-cycle pattern that appears strongly inconsistent. We find a strongly declining empirical profile in the first 10 years of working life, whereas the simulated profile is flat and does not show any life-cycle variation. In Appendix 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.⁶

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

⁶The 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 5: Life-cycle profiles and heterogeneity



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

saving behavior that explicitly introduces heterogeneity 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 phenomenon on the welfare and earnings dynamics of new labor market entrants.

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 et al., 2015). Workers start their life in the working phase that lasts for T^W periods. At the end of the working phase, workers move to the transition phase that is of stochastic length with expected duration T^T . In the end, workers leave the transition phase to the retirement phase that lasts for T^R periods. In each period before the retirement phase, a worker is either employed or 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

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

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.⁷ 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.

⁷Although 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.

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 current or last wage w , the separation probability λ if employed, and her level of human capital h . We formulate separate value functions for employed and 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,

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

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:

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

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

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

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

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

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

where N_w is the number of wage realizations 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

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

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

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

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

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 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

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

We solve the model using backward induction and grid search for the consumption-saving and effort choice decisions. We provide further details on the numerical implementation in Appendix A.6.

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 $h_{i,t} \in \{h_1, \dots, h_{N_h}, h^*\}$, 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.⁸

⁸This structure is reminiscent of the earnings process in Castaneda et al. (2003). The difference here is that the high-income state in our calibration will be much less extreme than the one in Castaneda et al. (2003) and much more persistent. We follow similar ideas already proposed in Jung and Kuhn (2018) or Hubmer (2018).

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 $[\underline{w}, \bar{w}]$ and job stability $[1 - \bar{\lambda}, 1 - \underline{\lambda}]$. We set $N_w = 5$, $\underline{w} = 1$, and $\bar{w} = 1.85$, in line with the empirical support of mean log earnings, and use equidistant grid points in logs. For job stability $1 - \lambda$, we set $N_\lambda = 10$ and set $\bar{\lambda} = 0.35$ 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.⁹ 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

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

minimizes the difference between model moments and empirical moments. For the empirical moments, we use the life-cycle profiles of (log) earnings (mean and variance), labor market transition rates, tenure (mean, median, 75th percentile), and of the wealth-to-income ratio. For labor market transition rates, we rely on estimated life-cycle profiles from Jung and Kuhn (2018) based on CPS data.¹⁰ In Appendix 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 presents the model parameters together with their estimated values.

Table 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.¹¹ 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

¹⁰We 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.

¹¹Utility 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.

rates is qualitatively and quantitatively consistent with the calibration in Hornstein et al. (2011) for the United States. The shape parameters of the marginal distributions, ψ_w and ψ_λ , determine the relative frequency of the different wage and job stability levels in the offer distribution. They imply that one-third of job offers come with the lowest wage, and less than one out of twelve job offers come with the highest wage level. For job stability, we get that less than one out of 20 jobs are the most stable lifetime jobs, whereas almost one out of six job offers are of the least stable type, lasting in expectation for one quarter only. The copula parameter θ implies a positive correlation between wages and job stability. If wages and job stability were independent, the probability of the least-stable lowest-paying job would be 5.1 percent, but given θ , it is almost 50 percent higher with 7.4 percent. For the most-stable best-paying job, the offer probability is 0.9 percent, highlighting that stable and high-paying jobs are hard to find. Such a correlation between wages and job stability is also in line with the empirical evidence in Jung and Kuhn (2018) that high-wage jobs are more stable (lower separation rates). Figure 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 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.

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 6 shows the empirical life-cycle profiles for separation, job-to-job, and job-finding rates and their model counterparts. Looking at the separation rate in Figure 6a, we see that its evolution is matched very closely. Model and data show a strong decline up to age 30 and constantly falling separation rates between ages 30 and 50. Figure 6b shows that the model also matches the life cycle of job-to-job rates very well, with only a slightly steeper decline of job-to-job rates between ages 20 and 30 compared to the data. 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 6c are matched well in level and trend and generally show life-cycle variation. Finally, Figure 6d shows the life-cycle profiles of mean, median, and the 75th percentile of the tenure distribution. The model closely matches the empirical increase and heterogeneity in job stability. Importantly, this demonstrates that the model is jointly consistent with high average transition rates (Figures 6a and 6b) and high job stability for most workers (Figure 6d).¹² Appendix Figure 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.¹³

Figure 7 turns to the life-cycle dynamics of earnings and wealth. Looking at the life-cycle profile of mean log earnings in Figure 7a, we find that the model matches the steep increase in earnings after labor market entry and the flattening out after age 40. It closely matches the large average increase of roughly 0.8 log points over the life cycle but shows slightly less concavity in comparison to its empirical counterpart. Figure 7b shows the life-cycle profile of the wealth-to-income ratio as a measure of wealth accumulation. Again, we find a close fit between model and data. Wealth-to-income ratios in model and data rise from slightly above 0 at age 20 to approximately 3.5 at age 55.¹⁴ The empirical profile is slightly less convex than

¹²For consistency, here we consider transition rates and tenure levels from CPS data. In Appendix A.7, we show that wage and tenure data from the SCF data align closely to the CPS levels.


¹³For 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.

¹⁴SCF 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.

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 7c and 7d is not. In our empirical analysis (Section 2), we document a positive correlation between job stability and wealth accumulation after controlling for age effects. Figure 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 particular, 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.

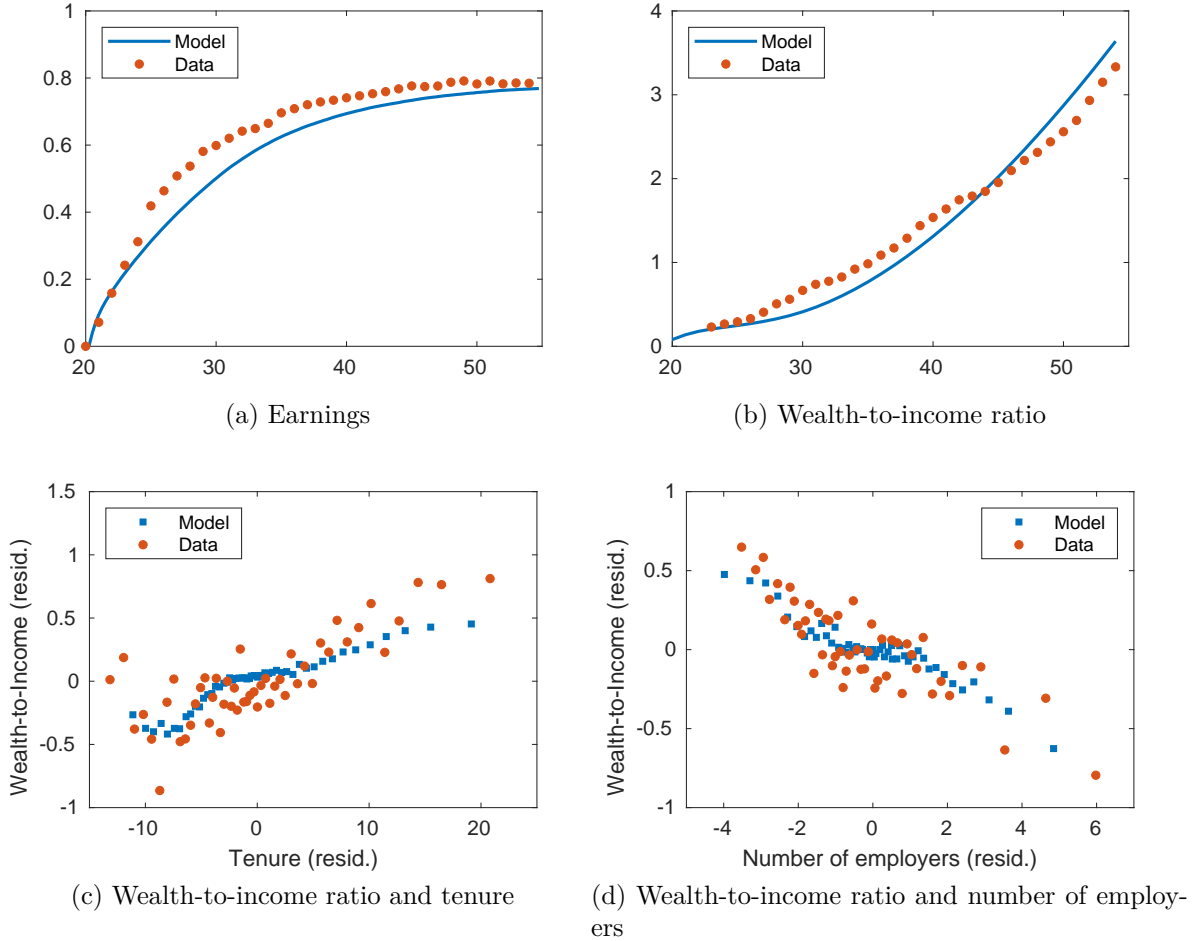
Figure 7d shows that the model also accounts for the negative relationship between wealth-to-income ratios and number of employers, demonstrating that job loss and job-to-job dynamics and their relationship to wealth accumulation are consistently accounted for by the model mechanism and align quantitatively with the data. By contrast, we show in Appendix 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 6: Transition rates and tenure



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


Figure 7: Earnings and wealth



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 8). Figure 8a shows the life-cycle increase in log earnings variance and the close match between model and data. We see the typical almost linear increase in the variance by age in both the model and the data. Labor market search models oftentimes struggle to account for this increase, as discussed in Lise (2012), Jung and Kuhn (2018), or Hubmer (2018). Augmenting models of job search with differences in human capital accumulation provides one way to account for the observed increase (Jung and Kuhn, 2018; Hubmer, 2018). While we build on this approach, we further refine it by endogenizing the human capital accumulation decision. As we 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 8: Earnings, consumption, wealth, and 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 2 for details.

Figure 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 8a and 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 8c. We use the

measure for employment inequality introduced in the empirical analysis of Section 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 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 et al. (2008), or Heathcote et al. (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 et al. (2019). We also decompose earnings growth and dispersion over the life cycle and demonstrate that our decomposition is consistent with the results in Topel and Ward (1992) on early career wage growth and resolves the tension highlighted in Hornstein et al. (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 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.

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.¹⁵

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 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 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

¹⁵For welfare results in Section 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 A.12.

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 9: Job stability, precautionary savings, and consumption growth



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 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 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 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 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 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.¹⁶

Looking at income in Figure 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

¹⁶For employment rates in Figure 11a, we show percentage point differences relative to the average 25-year-old worker.

Figure 10: Consequences of differences in job stability on income and 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.

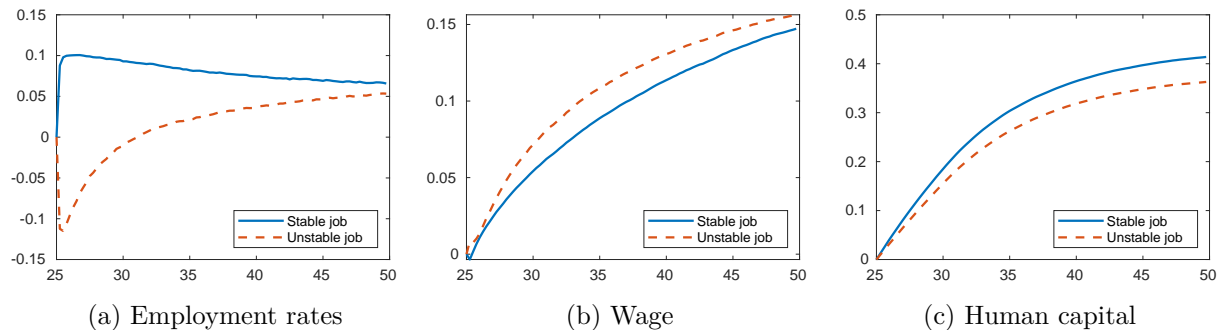
of observed income differences and, as discussed below, are on impact the key driver of the income divergence. Over time, we find not only that the bad start to the labor market leads to quickly diverging incomes but also that the decline remains persistent for the remainder of working life. At age 40, 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 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 10a). 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 11: Decomposition income effects from differences in job stability



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 percent-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.

Figure 11 decomposes income dynamics into an effect from employment rates, human capital, and wages. Figure 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 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 11b and 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 regarding wages in Figure 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 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 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 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.¹⁷

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 et al., 1993b) 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 40.¹⁸ Empirical studies document that such job displacements lead to large and persistent earnings losses for workers (Jacobson et al., 1993b; Couch and Placzek, 2010; Davis and von 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 (Jung and Kuhn, 2018; Jarosch, 2015). Models without heterogeneity in job stability struggle to account for the persistence in earnings losses (Low et al., 2010), a fact we also highlight in Appendix A.9 where we show earnings losses for a model without job stability heterogeneity and how the resulting earnings losses are only transitory.

Figure 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 that upon job loss, workers are unlikely to immediately find a well-paying job through off-the-job search. Most of the job offers that workers receive come with low wages and high separation rates. On-the-job search allows workers to catch up by climbing the wage ladder toward better-paying jobs; however, the speed of convergence reduces substantially after the first five years. Looking at the evolution of human capital, we

¹⁷It is important to note that productivity differences stemming from unobserved differences in human capital are empirically part of frictional wage dispersion.

¹⁸This 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 et al. (1993b)).

Figure 12: Cost of displacement



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

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 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 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. 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.

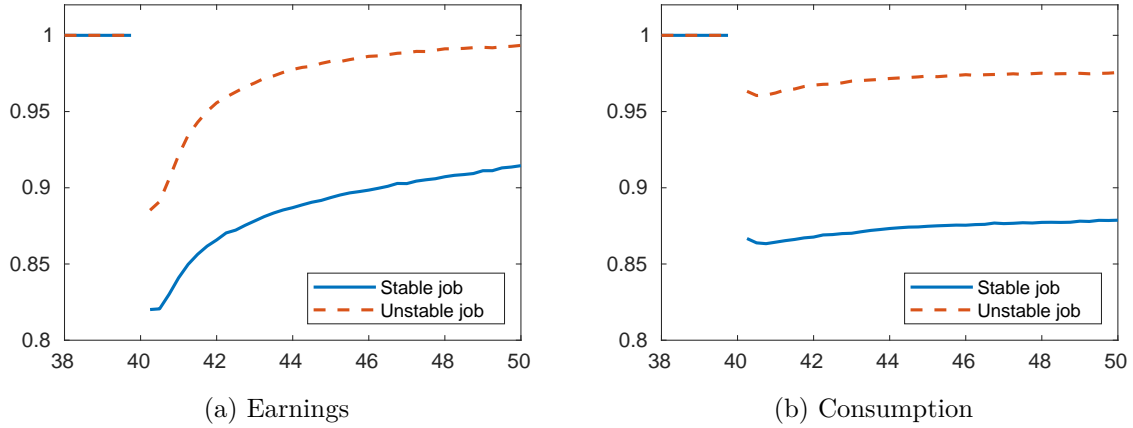
These consumption and labor market dynamics also provide intuition for the positive relationship between wealth and tenure implied by the model and observed in the data (Figure 7). Workers who lose their jobs experience significant decreases in income, and at the same time, their tenure drops to zero. 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 7.

Figure 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.

On impact, losing the stable or unstable job leads to earnings losses of 12 percent and 18 percent, respectively (Figure 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

Figure 13: Effects of displacement by job stability



Notes: This figure shows the evolution of earnings and consumption of workers who become unemployed at age 40 relative to the control group. Workers with 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.

from heterogeneity in job stability are only transitory (Appendix 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 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 con-

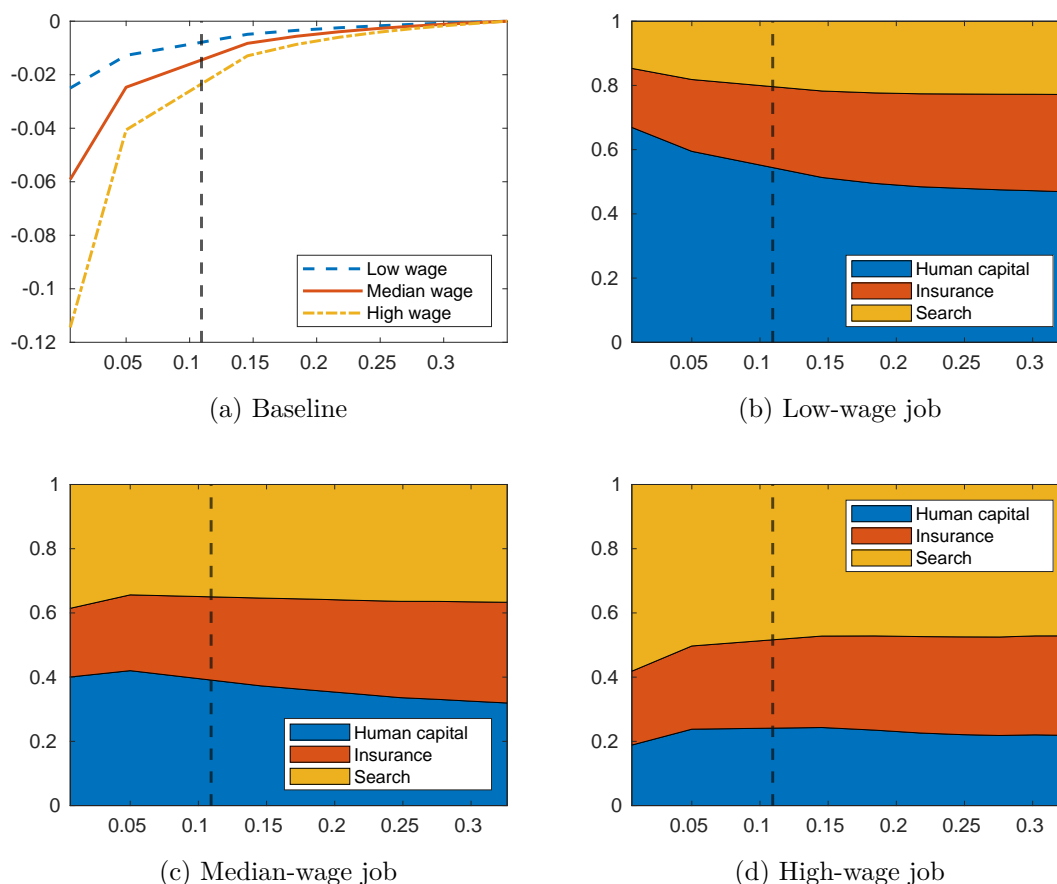
sumption in excess of the persistent earnings drop results from the differences in employment rates that lead to a larger drop in income compared to earnings. The employment effect is substantially larger for workers who lose a stable job as employment rates starting from a stable job are much higher 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.

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 et al., 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 14: Welfare costs of job instability for different types of jobs



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.

Figure 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 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.¹⁹ In the next step, we additionally endow the worker who already has high human capital with high wealth.²⁰ For this worker, we again conduct the comparative statics experiment of setting her at the least-stable job. Now, financial market incompleteness does not 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

¹⁹Specifically, 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.

²⁰We define high wealth as the highest wealth level attained at the end of working life during our simulations of the baseline model.

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 14b to 14d show the decomposition of the welfare effect from Figure 14a into the three components at the three different wage levels.²¹ We find 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 14b). At the median wage (Figure 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.²²

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 et al., 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

²¹We construct the search component as the residual so that the three decomposition components always sum to 100 percent of the total effect.

²²We repeat the same welfare analysis in Appendix 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.


0.6 percentage points between 1994 and 2004, starting from around 2.8 percent per month.²³ 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 et al., 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 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 et al. (2020). The decline in the separation rates and job-to-job rates follows the estimated declines of monthly rates in Fujita (2018) and Fujita et al. (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.²⁴ Figure 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”).

²³Fujita et al. (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.

²⁴We 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 $\lambda_j = \underline{\lambda} + \left(\frac{j-1}{N_\lambda-1}\right)^{0.6} \times (\bar{\lambda} - \underline{\lambda})$. See footnote 9 for details of the baseline economy. Figure A.13 shows the marginal distributions of separation rates for the baseline economy and the less dynamic economy.

Figure 15: Comparison of labor market life-cycle profiles



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.


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 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.²⁵ 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

²⁵We 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.

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.²⁶ 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.

Figure 16: Comparison of earnings, wage, and human capital dynamics



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 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 16a shows the ratios of the average profiles where a number larger than one implies

²⁶If 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.

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 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 4). When we repeat the experiment from Figure 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 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.

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.

References

- AGUIAR, M. AND E. HURST (2013): “Deconstructing life cycle expenditure,” *Journal of Political Economy*, 121, 437–492.
- AIYAGARI, S. R. (1994): “Uninsured idiosyncratic risk and aggregate saving,” *The Quarterly Journal of Economics*, 109, 659–684.
- BAYER, C. AND M. KUHN (2018): “Which Ladder to Climb? Wages of workers by job, plant, and education,” .
- BEWLEY, T. (undated): “Interest bearing money and the equilibrium stock of capital,” Tech. rep., mimeo.
- BILAL, A. (2019): “The Geography of Unemployment,” .
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98, 1887–1921.
- BRICKER, J., L. J. DETTLING, A. HENRIQUES, J. W. HSU, L. JACOBS, K. B. MOORE, S. PACK, J. SABELHAUS, J. THOMPSON, AND R. A. WINDLE (2017): “Changes in US family finances from 2013 to 2016: Evidence from the Survey of Consumer Finances,” *Fed. Res. Bull.*, 103, 1.
- BURDETT, K. AND D. T. MORTENSEN (1998): “Wage differentials, employer size, and unemployment,” *International Economic Review*, 257–273.
- CAGETTI, M. (2003): “Wealth Accumulation over the Life Cycle and Precautionary Savings,” *Journal of Business & Economic Statistics*, 21, 339–353.
- CAJNER, T., I. GÜNER, AND T. MUKOYAMA (2020): “Gross Worker Flows over the Life Cycle,” .
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *The Quarterly Journal of Economics*, 112, 1–55.
- CASTANEDA, A., J. DIAZ-GIMENEZ, AND J.-V. RIOS-RULL (2003): “Accounting for the US earnings and wealth inequality,” *Journal of political economy*, 111, 818–857.
- COUCH, K. A. AND D. W. PLACZEK (2010): “Earnings losses of displaced workers revisited,” *American Economic Review*, 100, 572–89.

- DALY, M., D. HRYSHKO, AND I. MANOVSKII (2016): “Improving the measurement of earnings dynamics,” Tech. rep., National Bureau of Economic Research.
- DAVIS, S. J. (2008): “The Decline of Job Loss and Why It Matters,” *American Economic Review*, 98, 263–267.
- DAVIS, S. J. AND T. VON WACHTER (2011): “Recessions and the Costs of Job Loss,” *Brookings Papers on Economic Activity*, 1.
- DEATON, A. AND C. PAXSON (1994): “Intertemporal choice and inequality,” *Journal of political economy*, 102, 437–467.
- DÍAZ-GIMÉNEZ, J., A. GLOVER, AND J.-V. RÍOS-RULL (2011): “Facts on the distributions of earnings, income, and wealth in the United States: 2007 update,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 34, 2–31.
- FALLICK, B. AND C. A. FLEISCHMAN (2004): “Employer-to-employer flows in the US labor market: The complete picture of gross worker flows,” .
- FRIEDMAN, M. (1957): *Theory of the consumption function*, Princeton university press.
- FUJITA, S. (2018): “Declining labor turnover and turbulence,” *Journal of Monetary Economics*, 99, 1–19.
- FUJITA, S., G. MOSCARINI, AND F. POSTEL-VINAY (2020): “Measuring Employer-to-Employer Reallocation,” Tech. rep., National Bureau of Economic Research.
- GOURINCHAS, P.-O. AND J. A. PARKER (2002): “Consumption Over the Life Cycle,” *Econometrica*, 70, 47–89.
- GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2019): “What do data on millions of US workers reveal about life-cycle earnings dynamics?” *working paper*.
- HAGEDORN, M. AND I. MANOVSKII (2010): “Search frictions and wage dispersion,” *Manuscript, University of Zurich and University of Pennsylvania*.
- HALL, R. E. (1982): “The Importance of Lifetime Jobs in the US Economy,” *American Economic Review*, 72, 716–724.

- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010a): “Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006,” *Review of Economic dynamics*, 13, 15–51.
- (2010b): “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006,” *Review of Economic Dynamics*, 13, 15–51.
- HORNSTEIN, A., P. KRUSELL, AND G. L. VIOLANTE (2011): “Frictional Wage Dispersion in Search Models: A Quantitative Assessment,” *American Economic Review*, 101, 2873–2898.
- HUBMER, J. (2018): “The job ladder and its implications for earnings risk,” *Review of Economic Dynamics*, 29, 172–194.
- HUGGETT, M. (1993): “The risk-free rate in heterogeneous-agent incomplete-insurance economies,” *Journal of economic Dynamics and Control*, 17, 953–969.
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993a): “Earnings Losses of Displaced Workers,” *The American Economic Review*, 83, 685–709.
- (1993b): “Earnings losses of displaced workers,” *The American economic review*, 685–709.
- JAROSCH, G. (2015): “Searching for Job Security and the Consequences of Job Loss,” .
- JUNG, P. AND M. KUHN (2018): “Earnings Losses and Labor Mobility Over the Life Cycle,” *Journal of the European Economic Association*, jvy014–jvy014.
- KARAHAN, F. AND S. OZKAN (2013): “On the persistence of income shocks over the life cycle: Evidence, theory, and implications,” *Review of Economic Dynamics*, 16, 452 – 476.
- KREBS, T., M. KUHN, AND M. L. WRIGHT (2015): “Human capital risk, contract enforcement, and the macroeconomy,” *American Economic Review*, 105, 3223–72.
- KRUSELL, P., T. MUKOYAMA, R. ROGERSON, AND A. ŞAHIN (2017): “Gross worker flows over the business cycle,” *American Economic Review*, 107, 3447–76.
- KRUSELL, P., T. MUKOYAMA, AND A. SAHIN (2010): “Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations,” *The Review of Economic Studies*, 77, 1477.

- KUDLYAK, M. AND F. LANGE (2014): “Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories,” *IZA DP*.
- KUHN, M. (2013): “Recursive Equilibria in an Aiyagari-Style Economy with Permanent Income Shocks,” *International Economic Review*, 54, 807–835, [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12018](https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12018).
- KUHN, M. AND J.-V. RÍOS-RULL (2016): “2013 Update on the US earnings, income, and wealth distributional facts: A View from Macroeconomics,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 37, 2–73.
- KUHN, M., M. SCHULARICK, AND U. STEINS (forthcoming): “Income and Wealth Inequality in America, 1949-2016,” *Journal of Political Economy*.
- LARKIN, K. P. (2019): “Job Risk, Separation Shocks and Household Asset Allocation,” in *2019 Meeting Papers*, Society for Economic Dynamics, 1058.
- LISE, J. (2012): “On-the-Job Search and Precautionary Savings,” *The Review of Economic Studies*, 80, 1086–1113.
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2010): “Wage Risk and Employment Risk over the Life Cycle,” *American Economic Review*, 100, 1432–67.
- MEGHIR, C. AND L. PISTAFERRI (2004): “Income variance dynamics and heterogeneity,” *Econometrica*, 72, 1–32.
- MICHELACCI, C. AND H. RUFFO (2015): “Optimal Life Cycle Unemployment Insurance,” *American Economic Review*, 105, 816–59.
- MODIGLIANI, F. AND R. BRUMBERG (1954): “Utility analysis and the consumption function: An interpretation of cross-section data,” *Franco Modigliani*, 1, 388–436.
- MOLLOY, R., C. SMITH, AND A. K. WOZNIAK (2020): “Changing Stability in U.S. Employment Relationships: A Tale of Two Tails,” Working Paper 26694, National Bureau of Economic Research.
- MOLLOY, R., R. TREZZI, C. L. SMITH, AND A. WOZNIAK (2016): “Understanding declining fluidity in the US labor market,” *Brookings Papers on Economic Activity*, 2016, 183–259.


- MORCHIO, I. (2020): “Work Histories and Lifetime Unemployment,” *International Economic Review*, 61, 321–350.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job creation and job destruction in the theory of unemployment,” *The review of economic studies*, 61, 397–415.
- OECD (2015): *Pensions at a Glance 2015*.
- PFEFFER, F. T., R. F. SCHOENI, A. KENNICKELL, AND P. ANDRESKI (2016): “Measuring wealth and wealth inequality: Comparing two US surveys,” *Journal of economic and social measurement*, 41, 103–120.
- PINHEIRO, R. AND L. VISSCHERS (2015): “Unemployment risk and wage differentials,” *Journal of Economic Theory*, 157, 397–424.
- PRIMICERI, G. E. AND T. VAN RENS (2009): “Heterogeneous life-cycle profiles, income risk and consumption inequality,” *Journal of monetary Economics*, 56, 20–39.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25–49.
- STEVENS, A. H. (1997): “Persistent effects of job displacement: The importance of multiple job losses,” *Journal of Labor Economics*, 15, 165–188.
- TOPEL, R. H. AND M. P. WARD (1992): “Job Mobility and the Careers of Young Men,” *The Quarterly Journal of Economics*, 107, 439–479.

A Appendix

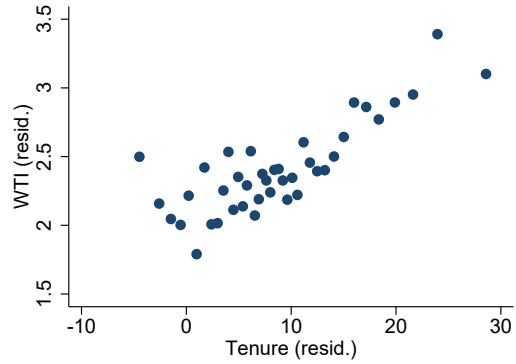
A.1 Employment history and wealth: robustness checks

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 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 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.

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

$$\begin{aligned} T_H - T_R &= \frac{1}{N} \sum_{i=1}^N \frac{1}{\pi_i} - N \left(\sum_{i=1}^N \pi_i \right)^{-1} = \left(\sum_{i=1}^N \pi_i \right)^{-1} \left(\left(\frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^N \pi_j}{\pi_i} \right) - N \right) \\ &= \left(\sum_{i=1}^N \pi_i \right)^{-1} \left(\left(\sum_{i=1}^N \frac{\frac{1}{N} \sum_{j=1}^N \pi_j}{\pi_i} \right) - N \right) = \left(\sum_{i=1}^N \pi_i \right)^{-1} \left(\left(\sum_{i=1}^N \frac{\bar{\pi}}{\pi_i} \right) - N \right). \end{aligned} \quad (9)$$

Using a second-order approximation of $f(\pi_i) = \frac{\bar{\pi}}{\pi_i}$ around $\bar{\pi}$ and plugging it into equation (9) yields


$$\begin{aligned}
T_H - T_R &= (\bar{\pi})^{-1} \left(\frac{1}{N} \sum_{i=1}^N \left(1 - \frac{\pi_i - \bar{\pi}}{\bar{\pi}} - \left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}} \right)^2 \right)^2 - 1 \right) \\
&= T_R \frac{1}{N} \sum_{i=1}^N \left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}} \right)^2 \\
\frac{T_H - T_R}{T_R} &= \frac{1}{N} \sum_{i=1}^N \left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}} \right)^2 \\
\sigma_E &= 1 + \frac{1}{N} \sum_{i=1}^N \left(\frac{\pi_i - \bar{\pi}}{\bar{\pi}} \right)^2. \tag{10}
\end{aligned}$$

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.

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

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



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

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 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 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.

Compared to the empirical profiles shown in Figure 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.

Figure A.3: Worker types and tenure distribution



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

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

$$V_e^{T,P}(a, w, \lambda, h) = \max_{\{c, a' \geq 0\}} u(c) + \beta \left[\psi V_r(a', w, h, j_r = 1) + \right. \quad (11)$$

$$\left. (1 - \psi) (\pi_e V_e^{T,S}(a', w, \lambda, h) + (1 - \pi_e) V_e^T(a', w, \lambda, h)) \right] \quad (12)$$

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

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

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

where N_w 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

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

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

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

A.6 Model solution and estimation

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 \mathcal{A} , the wage grid by \mathcal{W} , the grid for job destruction probabilities by \mathcal{L} , and the grid for human capital by \mathcal{H} , we construct the state space as the Cartesian product of the separate grids $\mathcal{A} \times \mathcal{W} \times \mathcal{L} \times \mathcal{H} = \{a_1, \dots, a_{N_a}\} \times \{w_1, \dots, w_{N_w}\} \times \{l_1, \dots, l_{N_l}\} \times \{h_1, \dots, h_{N_h}\}$. The upper bounds 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 $[0, 1]$.²⁷ 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

²⁷We standardize the support of wages and job destruction probabilities to allow for an easier numerical implementation of the joint distribution. We discuss the details in Section 3.1.

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.

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

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

where the empirical profiles are denoted with a hat.

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 .

A.6.4 The estimated job-offer distribution

Panel A.4a of Figure 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 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 A.4: Job-offer distribution




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

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

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




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.

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

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



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


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 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 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.

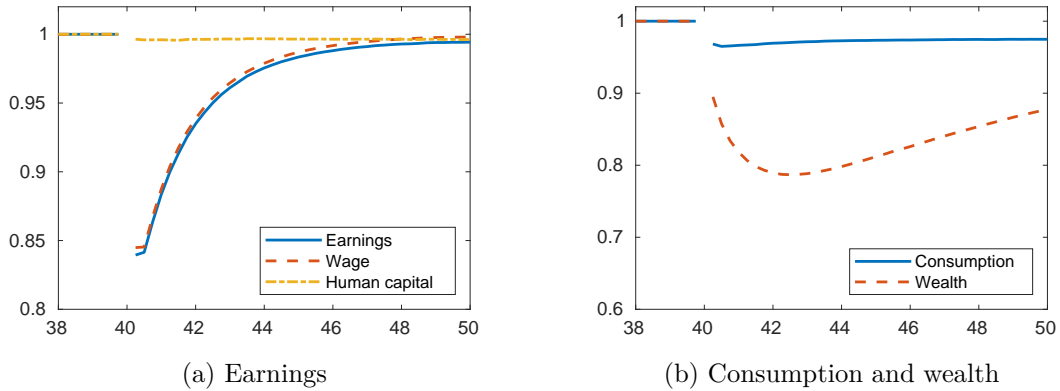
Furthermore, without any cross-sectional heterogeneity in the separation rate, the model also cannot replicate large and persistent earnings losses following displacement. In Figure A.8, we show the cost of displacement for the model without cross-sectional heterogeneity in the separation rate. Contrary to the results from the baseline model in Figure 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 A.7: Tenure and wealth



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

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



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.

A.10 Life-cycle earnings dynamics

Results presented in Section 3.2 demonstrate that the model matches the life-cycle profiles for means and variances. Here we provide additional evidence that the model also provides a good fit along other dimensions. To explore the fit for earnings dynamics, we compare how the model-implied earnings dynamics align with statistical representations of earnings processes as typically estimated in applied work and used to parametrize exogenous earnings dynamics in consumption-saving models. Such a description of earnings dynamics by a reduced-form statistical representation allows for a straightforward comparison of earnings dynamics between model and data. First, we perform a standard decomposition of earnings dynamics into a permanent and transitory component and estimate the variances of the innovation terms (Meghir and Pistaferri, 2004; Blundell et al., 2008; Heathcote et al., 2010a). Second, as emphasized in Guvenen et al. (2019), we look at the higher moments of the distribution of earnings growth. Third, we decompose earnings growth into contributions from human capital accumulation and job switching and demonstrate that the model aligns with the evidence on early career wage growth by Topel and Ward (1992). We close with a discussion of the model predictions for frictional wage dispersion.


In the first case, we estimate the variance of the permanent component of earnings dynamics using simulated earnings series from the model aggregated to an annual frequency. We apply the identification approach, as in Blundell et al. (2008), to the simulated data. We estimate a variance of the permanent component of 0.025 that falls well within the range of empirical estimates. Blundell et al. (2008) estimate time-varying variances of the permanent component ranging from 0.01 to 0.03 for the period from 1980 to 1990.²⁸ Empirical estimates for the variance of transitory shocks are harder to compare as they also comprise the contribution from measurement error that is likely substantial in the data, so it is not surprising that our finding is that the empirical estimates for the variance of the transitory shocks (0.03-0.05) are substantially larger than the model-implied estimate (0.016). We interpret the difference as the contribution from measurement error but also unmodeled earnings components such as bonuses and overtime pay.

In a second step, we consider the findings by Guvenen et al. (2019), who emphasize that earnings growth rates are not normally distributed but exhibit large negative skewness and high excess kurtosis. As has been demonstrated by Hubmer (2018), these patterns can

²⁸Heathcote et al. (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 et al. (2016) for further discussion.

be well explained by a life-cycle version of the job ladder model with a human capital process. The same also applies to our model. Figure 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 A.9: One-year earnings changes

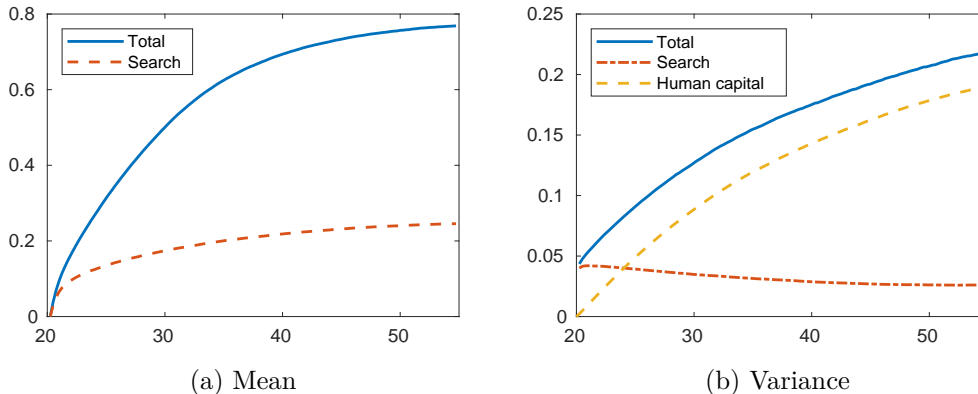


Notes: Figure shows the estimated kernel density of the model-based one-year earnings changes superimposed on Gaussian densities with the same standard deviation. Earnings growth net of average age effect shown. Kernel density estimation with a bandwidth of 0.05 used.

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 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 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 6). Human capital investment accounts for almost the entire increase in earnings once most workers have found stable jobs.

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



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

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

dispersion.

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

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

For this comparison, we split households into income and wealth quintiles. We consider the joint distribution by comparing how households are distributed across wealth quintiles conditional on their income quintile. For wealth dynamics, we compare how households move across wealth quintiles over time. Panels (a) through (c) in Figure A.11 show selected conditional distribution functions for wealth by income quintile from model and SCF data.²⁹ 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

²⁹In Table 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.

five-step wealth transition matrix).³⁰ For data on the joint distribution, we rely on the SCF data. We follow previous research (Díaz-Giménez et al., 2011; Kuhn et al., forthcoming) and rely on PSID data from 1984 to 1999 to trace out individual-level wealth dynamics. The repeated cross sections of the SCF data prevent such an analysis.³¹

Figure A.11: Joint distribution of income and wealth and wealth dynamics



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

Looking at the joint distribution in Figures A.11a to 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

³⁰We report all conditional probabilities in Table A.2.

³¹Pfeffer et al. (2016) compare SCF and PSID data, concluding that except for the very top of the wealth distribution, the two surveys provide consistent wealth distributions for the vast majority of households.

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

Looking at wealth dynamics in Figures A.11d to 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 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

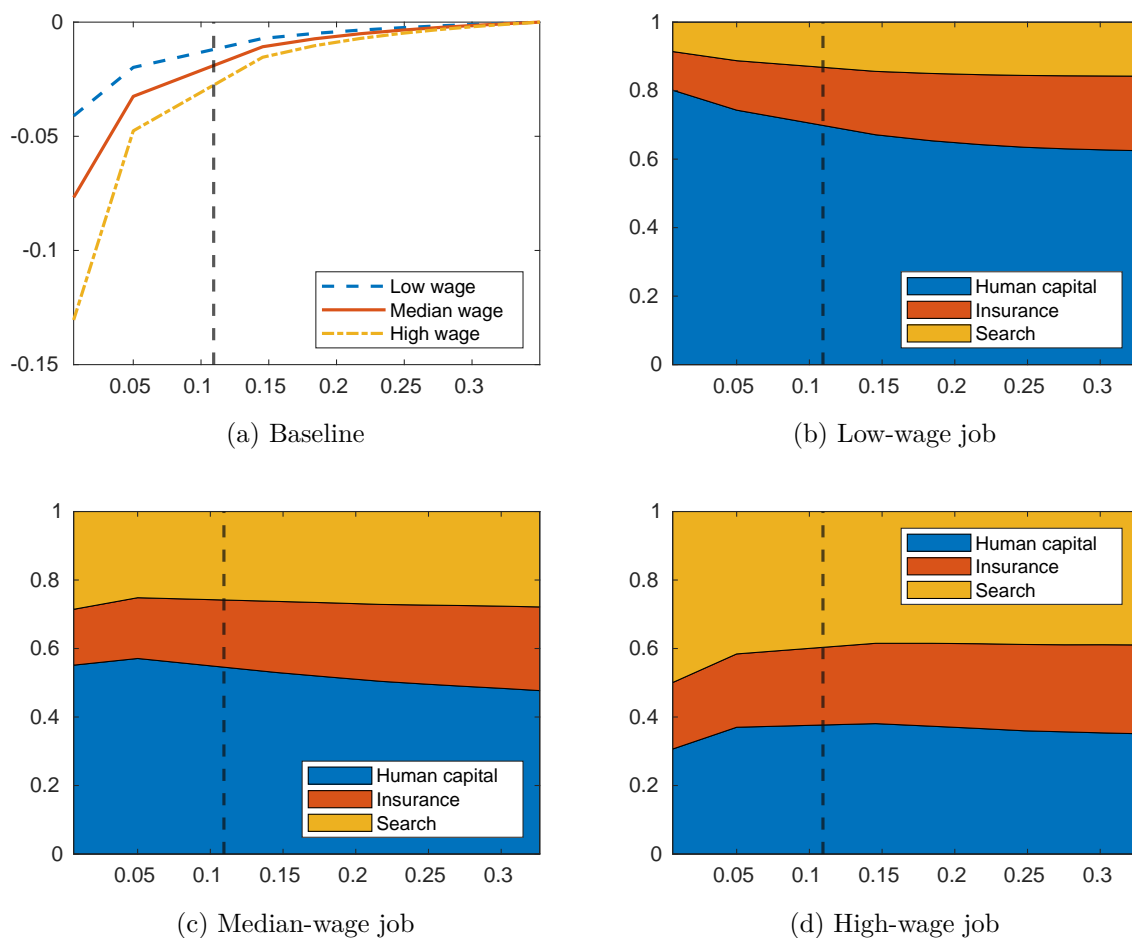
Table A.2: Wealth transition matrix, age 40

	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

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

Section 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 4.3. Figure 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 A.12: Welfare costs of job instability for 20-year-old workers in different job types




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.

A.13 Calibration for declining labor market dynamism

In Section 5, we describe how we calibrate declining labor market dynamism for the U.S. economy. Figure 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 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 A.13 as first-order stochastic dominance of separation rates in the baseline economy compared to the economy with lower mobility.

Figure 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 5.