Optimal Progressive Pension Systems and Heterogeneous Labor Market Risk*

Leanne Nam†

July 20, 2026

Abstract

Large differences in employment stability are a key feature of the labor market, with unstable careers leading to reduced pension entitlements and lower life-cycle wealth accumulation. This paper studies how optimal pension systems should account for heterogeneity in employment stability. I develop a novel life-cycle model that features heterogeneous job-separation risk with endogenous human capital accumulation, consumption-saving behavior, and retirement decisions in a frictional labor market. Calibrated to the U.S. economy, the optimal pension system is substantially more progressive than the current Social Security system, with an optimal benefit schedule that approaches a benefit floor. Abolishing the Social Security cap further increases progressivity; together, these reforms generate welfare gains equivalent to 0.94% of lifetime consumption for new labor market entrants. Although higher progressivity distorts human capital investment and reduces labor supply, the insurance gains dominate. The optimal pension system provides substantial insurance against heterogeneous career risk and reduces consumption inequality over the life cycle.

Keywords: Pension System, Employment Stability, Life-Cycle Earnings Risk
JEL: E24, H21, H55, J64

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

†University of Mannheim, leanne.nam@uni-mannheim.de, 68161 Mannheim, Germany

1 Introduction

Heterogeneity in employment stability is a salient feature of the labor market. While some workers hold lifetime jobs, others face unstable work histories and frequently experience unemployment. These differences have important implications not only for life-cycle labor market outcomes but also for old-age security, as interrupted employment histories reduce pension entitlements and hinder retirement savings. A growing body of evidence documents substantial and persistent dispersion in employment stability across workers (Hall, 1982; Jung and Kuhn, 2019; Molloy et al., 2024; Kuhn et al., 2025; Gregory et al., 2025; Jarosch, 2023; Ahn et al., 2023). Yet, despite this evidence, little is known about how optimal policy should account for such heterogeneity. This paper aims to fill that gap.

In this paper, I develop a novel life-cycle model that incorporates heterogeneity in employment stability, endogenous human capital accumulation, and consumption–saving and retirement decisions in a frictional labor market. Existing models of pension design typically abstract from the rich labor market structure and endogenous human capital formation. By jointly modeling these components, the framework captures how differences in employment stability propagate through earnings dynamics and wealth accumulation over the life cycle and shape the design of an optimal pension system.

The model is calibrated to the U.S. economy to analyze the optimal degree of pension progressivity. The results show that the welfare-maximizing pension system is substantially more progressive than the current U.S. system, with an optimal benefit schedule that approaches a benefit floor. Raising the Social Security cap yields additional welfare gains, and welfare is maximized when the cap is eliminated altogether. Relative to the current U.S. pension system, the optimal reform generates welfare gains equivalent to 0.94% of lifetime consumption for labor market entrants.

In the model, households make consumption–saving decisions under incomplete financial markets and invest in risky human capital. The model features a frictional labor market in which workers receive job offers both while unemployed and on the job. Upon reaching eligibility for pension benefits, workers optimally choose their retirement age, allowing labor supply to adjust at the extensive margin. A central innovation of the framework is the incorporation of heterogeneity in employment stability, building on a growing literature documenting persistent differences in employment stability across workers (Pinheiro and Visschers, 2015; Kuhn et al., 2025; Jarosch, 2023). Workers differ in job-separation risk at each stage of the life cycle. Unlike in existing work, however, this heterogeneity is embedded in a model with endogenous human capital accumulation and retirement decisions.

Endogenous human capital accumulation is a key driver of earnings dynamics (Becker,

1964; Ben-Porath, 1967) and an important mechanism behind rising life-cycle earnings inequality (Hubmer, 2018; Jung and Kuhn, 2019). In the model, the interaction between employment stability and human capital investment amplifies the long-run consequences of job displacement (Jarosch, 2023). Workers who experience job loss fall down the job ladder, face persistently unstable employment, and accumulate less human capital and wealth over time. As a result, employment-stability heterogeneity generates persistent dispersion in lifetime earnings and savings outcomes, which are central for pension design. In the absence of sufficient progressivity, the pension system transmits unstable employment histories into old age, since benefits depend on pre-retirement earnings trajectories.

Analyzing the optimal pension system, I focus on the classic trade-off between insurance and incentives. A progressive pension system, which provides higher replacement rates to low-income workers, offers insurance against unstable employment histories and the risk of old-age poverty. However, higher progressivity weakens incentives to invest in human capital by lowering the returns to human capital investment. Pension progressivity also affects retirement behavior, as pension wealth is a key determinant of retirement decisions. In particular, higher progressivity reduces labor supply at the extensive margin by encouraging earlier retirement among low-income workers.

I calibrate the model to the U.S. economy and the institutional features of the Social Security system, targeting empirical life-cycle profiles of labor market transition rates, earnings (mean and variance), the tenure distribution, and wealth-to-income ratios. The calibrated model closely matches these key moments. Using this framework, I study the optimal design of the U.S. pension system that maximizes ex-ante long-run welfare in a partial-equilibrium setting. The model shows that heterogeneity in employment stability generates substantial inequality in labor market outcomes and is a key driver of dispersion in earnings, wealth accumulation, and consumption over the life cycle.

When the Social Security cap is fixed at its current level, increasing pension progressivity and moving toward a benefit schedule that closely resembles a benefit floor is optimal, generating a welfare gain equivalent to 0.19% of lifetime consumption for labor market entrants. The welfare analysis further shows that substantial gains can be achieved through moderate reforms of the Social Security cap. Raising the cap by 34% increases welfare by 0.33% of lifetime consumption, while doubling the cap raises welfare by 0.48%. Welfare continues to increase as the cap is relaxed further, with the highest welfare attained when the cap is abolished altogether. Under the optimal policy, workers at the 25th percentile of the lifetime income distribution experience an increase in their replacement rate from 63% to 80%, while replacement rates at the 75th percentile remain largely unchanged. Eliminating the cap increases payroll tax revenue from high-income workers, allowing the payroll tax rate to fall by 1.07 percentage points. Overall, the optimal pension system generates a welfare gain equivalent to 0.94% of lifetime consumption for labor market

entrants.

What is the role of heterogeneity in employment stability for optimal pension design? To answer this question, I consider a counterfactual economy without cross-sectional heterogeneity in employment stability conditional on age. In this environment, job-separation rates vary over the life cycle but are identical across workers of the same age. Comparing the optimal policy in this economy to that in the baseline model reveals that the progressivity parameter almost doubles in the baseline economy, whereas in the economy without age-specific heterogeneity in employment stability, it increases by only about 60%. This substantial difference highlights the importance of persistent differences in job-separation risk across workers. Taken together, these results demonstrate that heterogeneity in employment stability is a central determinant of the optimal degree of pension progressivity.

The optimal pension system with a higher progressivity level reduces the rise in consumption inequality over the life cycle by nearly one-third relative to the current U.S. Social Security system. It provides insurance to workers who suffer from frequent career interruptions and low lifetime earnings. While increased progressivity introduces distortions—reducing the average human capital stock by 2.5% and lowering the average retirement age by 0.4% of total years worked—these effects are modest and do not outweigh the insurance benefits. This is for two reasons. First, young workers do not reduce their human capital investment as they strive to accumulate human capital to increase their prospective labor earnings. As the return to human capital is largely realized in early and mid-career stages, a reduction in pension-related returns in the distant future has limited impact on the effort decision of young workers. This intuition aligns with Michelacci and Ruffo (2015), who show that an optimal age-specific unemployment insurance system should provide higher insurance to young workers as this creates limited distortions due to the high marginal returns from human capital of the young. Second, changes in retirement decisions depend on the accumulated stock of human capital. Low-productivity workers with relatively low human capital decide to retire earlier, whereas productive workers postpone their retirement in response to the increase in pension progressivity. The delay of retirement of productive workers partly offsets the distortion of retirement incentives on low-productivity workers such that the aggregate distortionary effect remains small.

I analyze the transition from the current pension system to the optimal policy, assuming that the reform is phased in gradually over a 70-year horizon. While all labor market entrants benefit from the reform as they are ex-ante identical, welfare effects among older cohorts vary depending on their employment histories. At age 40, workers with low employment stability (bottom 25% of the employment-stability distribution) experience a welfare gain equivalent to 1.13% of remaining lifetime consumption, whereas workers with high employment stability (top 25%) face a welfare gain of 0.55%. Importantly, the

reform generates positive welfare gains for the vast majority of workers, with welfare losses concentrated among older individuals and those with stable employment.

The remainder of this paper is structured as follows. Section 2 relates this paper to the existing literature. Section 3 presents the life-cycle model, and Section 4 describes the baseline calibration. Section 5 examines how heterogeneity in employment stability shapes inequality in labor market outcomes over the life cycle. Section 6 analyzes the optimal design of the pension system in the baseline economy. Section 7 concludes.

2 Related literature

This paper contributes to three important strands of macroeconomic literature. First, by analyzing the policy implications of heterogeneity in employment stability, it contributes to the growing strand of the literature that provides evidence for heterogeneity in employment stability and labor market dynamics (Ahn et al., 2023; Gregory et al., 2025; Molloy et al., 2024). Among others, Jung and Kuhn (2019) find considerable heterogeneity in employment stability for the United States and show that it is key to explain persistent earnings losses after job displacement in a life-cycle model with search. Using German Social Security data and a model with heterogeneity in employment stability, Jarosch (2023) points out that the interaction between human capital and employment stability is crucial to explain the observed earnings losses. Morchio (2020) shows that unemployment is concentrated among certain groups of workers. Workers significantly differ in their job-separation rates and such differences are large from the start of the career. Kuhn et al. (2025) also provide evidence for employment-stability heterogeneity in the U.S. labor market. Based on a life-cycle model, they show that employment stability at labor market entry significantly affects income, consumption, and wealth accumulation over the whole life. I follow the strand of the literature that assumes that heterogeneity in employment stability is generated due to differences across jobs (Jarosch, 2023; Bilal, 2023; Jung and Kuhn, 2019; Larkin, 2019). Another strand of the literature considers heterogeneity in employment stability as worker types rather than job types (Ahn et al., 2023; Hall and Kudlyak, 2019; Gregory et al., 2025). The model implications in this paper are also consistent with these papers in that worker types generate heterogeneity in employment stability. This is because the model generates very stable and unstable career paths across workers that are strongly persistent and in fact, a worker with a very stable job is observationally indistinguishable from a very stable worker type within the model. Hence, differences in employment stability in this paper can be interpreted as worker types as in the above-mentioned empirical approaches.

Second, this paper relates to the literature on unemployment risk and inequality in lifetime earnings. Using the life-cycle model, I discuss the interaction between heterogeneity

in employment stability and how unstable career paths affect lifetime earnings of workers, which consequently affect their pension entitlements. These results complement the large strand of empirical literature on the relationship between inequality in lifetime earnings and labor market risks. For example, Bonhomme and Robin (2009) explore inequality in employment and lifetime earnings for France and find that unemployment is a key driver of lifetime earnings inequality. Similarly, Bowlus and Robin (2012) show that unemployment mobility has important consequences on lifetime earnings inequality. Bönke et al. (2015) find that inequality in lifetime earnings has been increasing in Germany over time, with a larger increase for individuals at the bottom of the earnings distribution. Longer unemployment durations are one of the major sources of the increase in lifetime earnings inequality. Haan et al. (2025) find that heterogeneity in labor market outcomes is a major source of inequality in lifetime earnings for workers with the same abilities.

Lastly, this study contributes to the literature on optimal design of pension systems by applying a life-cycle framework which explicitly takes heterogeneity in employment stability into account. The application of a life-cycle model is a frequently used approach in the previous macroeconomic literature on optimal pension systems. Beginning with the work by Feldstein (1974), a large body of literature investigates the effects of pension systems in a life-cycle framework. For example, Hubbard and Judd (1987) consider a model with capital-market imperfections to show that the introduction of Social Security does not increase an economy's welfare as much as in an economy without borrowing constraints. A more recent work is the study by Fehr and Habermann (2008) who analyze the welfare effects of introducing a more redistributive pension system to Germany. They find that the positive liquidity and income insurance effects are large enough such that an increase in progressivity enhances welfare despite its distortionary effects on labor supply. Fehr et al. (2013) analyze the impact of higher progressivity of the German pension system in a model with disability risk and labor supply decisions at the extensive margin. Golosov et al. (2013) take the current form of the U.S. retirement benefit function and determine its optimal structure, assuming that workers have heterogeneous productivities. They find that a more redistributive system increases the economy's welfare by a considerable amount. O'Dea (2018) finds that introducing means-tested income floors for the elderly is beneficial. Some other prominent studies on pension systems are, for example, Krueger and Pischke (1992); De Nardi et al. (2001); Coile and Gruber (2001); Cremer et al. (2004); Gruber and Wise (2004); Bloom et al. (2007); Haan and Prowse (2014); and Borella et al. (2023). This paper differs from the earlier work on pension reforms in that I study heterogeneity in employment stability as a source of inequality in lifetime earnings.

3 Life-cycle model

This section outlines a life-cycle model with heterogeneity in employment stability, endogenous human capital accumulation, and retirement decisions in a frictional labor market. The key innovation of the framework is the joint study of heterogeneous job-separation risk and endogenous retirement behavior within an otherwise standard life-cycle setting. In particular, heterogeneity in employment stability has not been explicitly modeled in the literature on optimal pension design.

In the model, risk-averse agents maximize expected lifetime utility. The utility function is additively separable in utility from consumption and disutility from providing effort to accumulate human capital. The labor supply of an employed worker amounts to one unit of time. A worker's age is denoted by j.

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

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

The period budget constraint of an agent is given by

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

where a and h denote the risk-free asset and the stock of human capital of the agent, respectively. Moreover, r denotes the risk-free rate in the economy, and $y^{net}(w_j, h_j, \epsilon) = y(w_j, h_j, \epsilon) - T_y(y(w_j, h_j, \epsilon)) - T_p(y(w_j, h_j, \epsilon))$ where y^{net} denotes the labor income of the current period y including transfers minus income taxes $T_y(y(w_j, h_j, \epsilon))$ and the pension

contribution $T_p(y(w_j, h_j, \epsilon))$ as functions of labor income. The assets are restricted to be non-negative ($a_j \geq 0$) implying a borrowing limit of zero.

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

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

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

which was previously rejected.

The investment decision with regard to human capital is a choice of an effort level $t \in [0,1]$ which entails a quadratic disutility κt^2.[1] For a given level of effort t, the realization of human capital investment is stochastic. More specifically, assuming that human capital levels are discrete and that h^+ denotes the immediate successor and h^- the immediate predecessor of h, the law of motion for human capital is

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

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

3.1 Recursive formulation of the decision problem

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

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

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

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

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

[1] I also consider an alternative model where human capital accumulation corresponds to the Ben-Porath (1967) model. While all model results of the alternative model remain largely unchanged and the data moments are also matched well, the earnings variance over the life cycle is slightly larger in the alternative model and does not exhibit concavity as observed in the data. See Appendix A.4 for details.

of an employed agent is as follows:

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

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

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

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

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

where π_u is the probability of receiving a job offer and declining benefits are captured by a transition from w to w^- where w^- denotes the next lower wage level. Note that an unemployed worker receives unemployment benefits $y(w, h, n)$ which are reduced in the next period if the worker remains unemployed. At the search stage, an unemployed worker faces the value function

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

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

Between periods $T^W - T^C$ and T^W, workers have the option to leave the labor force and enter the retirement phase. At the beginning of each of these periods, workers observe a shock ε, drawn from a logistic distribution with location parameters μ_j and a scale parameter σ. After observing this shock, they decide whether to remain in the labor force and continue working or to move to the retirement phase. This shock captures

health shocks and a preference shock related to the value of leisure, representing utility derived from leaving the labor market and entering retirement. Given the realization of the shock, if lifetime utility from retirement is larger than the expected utility of remaining in the labor force, agents decide to retire. Hence, agents face the following discrete choice problem

$$V_{\max}(a, w, \lambda, h, j) = \max\{V(a, w, \lambda, h, j), V^r(a, w, h, j_r) + \varepsilon_j\}, \quad \varepsilon_j \sim \text{Logistic}(\mu_j, \sigma)$$

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

$$\begin{aligned}\mathbb{E}\left[V_{max}(a, w, \lambda, h, j)\right] =& pV(a, w, \lambda, h, j) + (1-p)V^r(a, w, h, j_r) \\ &- \sigma\left((1-p)\log(1-p) + p\log(p)\right) + \mu(1-p),\end{aligned} \tag{2}$$

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

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

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

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

The current model abstracts from mortality risk, bequest motives, and the early claiming penalties in the Social Security system. This simplification is intended to isolate and clearly highlight the core mechanisms relevant for analyzing the optimal design of the pension system. Appendix A.14 presents several extensions of the model that incorporate these additional features. Importantly, the main results presented in the following sections remain robust under these extended specifications.

3.2 Payroll tax finance of the pension system

The pension system is financed by payroll taxes. That is, the expected present value of retirement benefits obtained by all workers in the economy has to be compensated by the

expected revenues from payroll taxes levied on employed workers. Hence, the condition

$$\sum_{i \in \mathcal{R}} \omega(\bar{y}) = \sum_{i \in \mathcal{W}} \tau_p \cdot \tilde{y}(w, h, \epsilon) + G \tag{3}$$

must be satisfied where

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

The cap^p denotes the maximum level of taxable earnings. The parameter τ_p in Equation (3) represents the payroll tax rate. $\mathcal{R}$ denotes the set of retirees, and $\mathcal{W}$ denotes the set of working-age individuals. Payroll taxes are levied only on the labor earnings of employed agents during the working phase; unemployed individuals and retirees do not pay payroll taxes. The left-hand side of Equation (3) aggregates pension benefits paid to all retirees alive in a given period, while the right-hand side aggregates payroll tax revenues collected from all working individuals with government transfers G.[2]

4 Calibration

This section describes the calibration of the model. It first presents the data sources, followed by the externally calibrated parameters, the calibration of the pension system, and the internally estimated parameters, including those governing retirement behavior and those estimated using the simulated method of moments. The section concludes by assessing the model's fit to the data.

4.1 Data

For the calibration of the model, I use three data sources: the Survey of Consumer Finances (SCF), the Panel Study of Income Dynamics (PSID), and the Current Population Survey (CPS). The SCF provides detailed information on households' financial positions, including earnings and wealth. The PSID enables me to track individuals' labor market histories over time, while the CPS provides information on labor market transition rates.

The SCF is a triennial household survey that offers comprehensive and representative data on income, wealth, and household balance sheets in the United States. In addition to detailed financial information, it includes measures of labor market status. I pool SCF

[2]Introducing government transfers G allows the payroll tax rate τ_p to be calibrated to the U.S. statutory rate of 12.4%. Any remaining financing gap in the pension budget is covered by G. Throughout the analysis of the optimal pension system, G is held fixed, and any changes in aggregate pension benefits received by retirees must be offset by corresponding adjustments in aggregate payroll tax revenues.

waves from 1992 to 2016 and restrict the sample to households with employed heads. Since the model abstracts from self-employment, households with self-employed heads are excluded. I further drop observations with implausible wage values—specifically, wages below 75% of the minimum wage—and trim the distribution by excluding the top 1% of households by wealth and earnings, as well as the bottom 1% by earnings and households with negative wealth. This ensures that the calibration focuses on the segment of the distribution the model is designed to capture.

The PSID is the longest-running longitudinal household survey in the United States, beginning in 1968. Surveys were conducted annually until 1997 and biennially thereafter. Detailed information on household wealth has been collected at five-year intervals since 1985 and biennially since 1999. The PSID is used to discipline earnings dynamics and life-cycle patterns. For labor market transition rates, I rely on estimated life-cycle profiles from Jung and Kuhn (2019), constructed using data from the CPS.

4.2 Functional forms and externally calibrated parameters

This subsection specifies the functional forms and presents externally calibrated parameters. One model period corresponds to one quarter. Agents derive logarithmic utility from consumption so that $u(c) = \log(c)$. The yearly interest rate is set to 0.035. Human capital is discretized on a grid $h \in \{h_1, ..., h_{N_h}, h^*\}$ with $h_1 = 1$ and $h_{N_h} = 6.5$. Human capital levels between these two grid points are equidistant in log space. To capture the right tail of the earnings distribution, the last grid point h^*, which represents the highest level of human capital, is set to $h^* = 25$. The probability of attaining the next higher level of human capital when exerting effort t is

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

where the probability p_H starts from a value of $\bar{p}_H$ and decreases in age in a geometric fashion. Upon reaching the human capital level h_{N_h}, the probability to reach the highest level of human capital h^* is an exogenous parameter p_H^* which is independent of age. In the calibrated model, the implied elasticity of human capital investment with respect to the income tax parameter ranges from -0.28 for the full sample to -0.35 for prime-age workers. These values are consistent with estimates in the literature; for example, Hendricks (2004) reports an elasticity of -0.295.

Agents enter the labor market at age 20 with the lowest level of human capital $h_1 = 1$ and with an asset level of $a_0 = 0$. The replacement rate for unemployment benefits is set to 0.4 as in Shimer (2005). The total span of the life cycle $T^W + T^R$ is set to 70 years such that workers live up to age 90. Following the U.S. Social Security legislation, workers

can start receiving retirement benefits at age 62.[3] Between ages 62 and 67, workers make retirement decisions. At age 67, all workers who decided to remain in the labor force up to this age are shifted to the retirement phase.

Wages and job-separation probabilities are discretized on grids with $[\underline{w}, \overline{w}]$ and $[\underline{\lambda}, \bar{\lambda}]$, respectively. The lowest quarterly separation probability is $\underline{\lambda} = 0.006$ which corresponds to lifetime jobs with an expected job duration of 42 years. The highest separation probability is set to $\bar{\lambda} = 0.35$. The job-offer distribution consists of a joint distribution of job-separation probability and wage. See Appendix A.5.1 for a detailed description of the job-offer distribution. In Appendix A.5.4, I also discuss the assumption that job-arrival rates are identical for workers conditional on employment status.

The labor earnings in the model are subject to progressive earnings taxation where

$$T_y(y) = y - (1 - \nu)y^{1-\tau_y}. \tag{4}$$

Following Heathcote et al. (2017), I calibrate the progressivity parameter of the income taxation τ_y to 0.181. In addition, the payroll tax rate that finances the pension system is set to 12.4%, consistent with the U.S. Social Security system.

I show in Appendix A.6 that the earnings process assumed in this paper matches a range of untargeted empirical moments related to earnings changes associated with job mobility, including job-to-job wage gains (Topel and Ward, 1992) and earnings losses upon displacement (Jacobson et al., 1993). In addition, I validate the model's earnings dynamics against key reduced-form empirical moments documented by Blundell et al. (2008) and Guvenen et al. (2021). The model successfully reproduces these established features of earnings risk, providing further support for the credibility of the assumed income process.

4.3 U.S. pension system

The pension system is calibrated to the institutional features of the U.S. Social Security system. I first specify the parametric benefit class and the role of the earnings cap, then describe how the model approximates the U.S. benefit schedule, and finally explain how average lifetime earnings are approximated within the model.

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

$$\omega(\bar{y}) = \begin{cases} \phi \cdot (\bar{y})^{1-\gamma} & \text{if } \bar{y} < cap^p \\ \overline{\omega} & \text{if } \bar{y} \geq cap^p \end{cases} \tag{5}$$

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

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

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

An increase in the cap also raises overall pension progressivity, while the cap enters both the contribution and the benefit scheme of the pension system. In the benefit function, a higher cap reduces progressivity as pension benefits increase up to the cap. In contrast, on the contribution side, a higher cap increases progressivity by expanding the share of earnings subject to payroll taxation. Overall, a higher cap strengthens pension progressivity because workers near and above the cap pay payroll taxes on the additional taxable earnings, while the associated increase in benefits is limited due to low marginal increase in pension benefits for high earners.

In order to calibrate the pension system in the model to the U.S. pension system, I consider the U.S. Social Security legislation for 2019. Based on this, I set the Social Security cap to \$132,900 and the first and second bendpoints to \$11,112 and \$66,996, respectively. Then, I apply the retirement benefit formula given as follows:[4]

$$\Omega(\bar{y}) = \begin{cases} 0.9\bar{y} & \text{if } \bar{y} < bp_1, \\ 0.9bp_1 + 0.32(\bar{y} - bp_1) & \text{if } bp_1 \leq \bar{y} < bp_2, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(\bar{y} - bp_2) & \text{if } bp_2 \leq \bar{y} < cap^p, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(cap^p - bp_2) & \text{if } \bar{y} \geq cap^p \end{cases} \tag{6}$$

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

[4]The bendpoints are converted to annual values, and the retirement benefit formula follows the documentation http://www.ssa.gov/OACT/COLA/piaformula.html.

estimated parameters are $\phi = 0.7$ and $\gamma = 0.44$. Figure A.1 in Appendix A.1 shows the U.S. Social Security benefit function and the approximating function in the model.

To assess the responsiveness of human capital investment in the model to pension progressivity, I conduct an experiment varying the progressivity parameter γ. Although this elasticity is difficult to estimate empirically, it is informative to compare it to the elasticity with respect to income taxation. The estimated elasticity of human capital investment with respect to γ is -0.39 for the full sample and -0.48 for prime-age workers. These magnitudes are close to, though somewhat larger in absolute value than, the corresponding elasticity with respect to the income tax progressivity parameter.

I use the final level of human capital and wages of individual workers attained at the end of the working phase to infer the lifetime earnings of workers. Human capital is a direct measure of lifetime earnings as workers accumulate human capital throughout their entire career and labor earnings are a function of human capital. The final level of human capital also contains information on workers' earnings histories because human capital accumulation is only possible in employment. Thus, it is informative about workers' average labor earnings and their duration of employment during the working phase. Importantly, the U.S. Social Security system calculates retirement benefits based on the 35 highest years of labor earnings. This feature provides implicit insurance, as years with low earnings or unemployment spells are excluded from the benefit calculation. I account for this mechanism in the model by approximating lifetime earnings using the 35 highest annual earnings over the working life. To approximate lifetime earnings as defined in the U.S. Social Security system, I compute, for each worker, the average of the 35 highest annual labor earnings and regress this measure on a cubic polynomial in final labor earnings, which are determined by human capital and wages. In fact, the final level of human capital and wages attained by workers until retirement explain around 90% of variation in lifetime earnings. Appendix A.5.6 discusses the accuracy of this approximation of average lifetime earnings in detail.

4.4 Estimation of remaining parameters

The remaining parameters are estimated internally to match a set of targeted life-cycle moments. The parameters governing labor market dynamics, human capital accumulation, and savings behavior are estimated by a simulated method of moments, while the parameters of the logistic shock that governs retirement decisions are calibrated to match the retirement age distribution and the retirement elasticity. This subsection describes the estimation, reports the estimated values, and discusses how each parameter is identified. The fit between model and data is assessed in Section 4.5.

I match model-implied profiles to the life-cycle evolution of labor market transition

Table 1: Estimated parameters with SMM

Parameter	Value	Description
β	0.994	Quarterly discount factor
κ	0.337	Utility cost of effort
π_e	0.406	Probability of a job offer when employed
π_u	0.810	Probability of a job offer when unemployed
ψ_w	0.557	Marginal distribution of w^*
ψ_λ	0.486	Marginal distribution of $1-\lambda^*$
θ	0.338	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.055	Skill upgrading probability
ρ	0.985	Persistence of skill upgrading probability
p_H^*	0.064	Probability to move to h^*

rates, the mean and variance of earnings, the tenure distribution (mean, median, and 75th percentile), and the wealth-to-income ratio. Table 1 reports the estimated parameter values and their corresponding targets.

The job-offer arrival rates, π_e and π_u, are identified from the life-cycle profiles of job-to-job and unemployment-to-employment transitions, respectively. The joint distribution of wage offers and job stability, $f(w,\lambda)$, is disciplined by the joint evolution of earnings, tenure, and separation rates over the life cycle. The parameters governing human capital accumulation are identified from the level, curvature, and dispersion of life-cycle earnings profiles, while the utility cost of human capital investment is pinned down by the flattening of earnings growth later in life. Finally, the discount factor is identified using the life-cycle profile of the wealth-to-income ratio.

Retirement decision. Recall that an agent chooses to retire if

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

where μ denotes the location parameter and σ denotes the scale parameter of the logistic distribution. I estimate on average $\mu = 14.06$ and $\sigma = 2.1$. The shock ε drawn from the logistic distribution can be interpreted as allowing for heterogeneity in preferences for leisure. It can also be interpreted as deriving an unexpected utility gain from retirement, for example due to health problems, which are an important cause for retirement (Dwyer and Mitchell, 1999; Disney et al., 2006; Jones et al., 2010; Galama et al., 2013; Trevisan

and Zantomio, 2016).

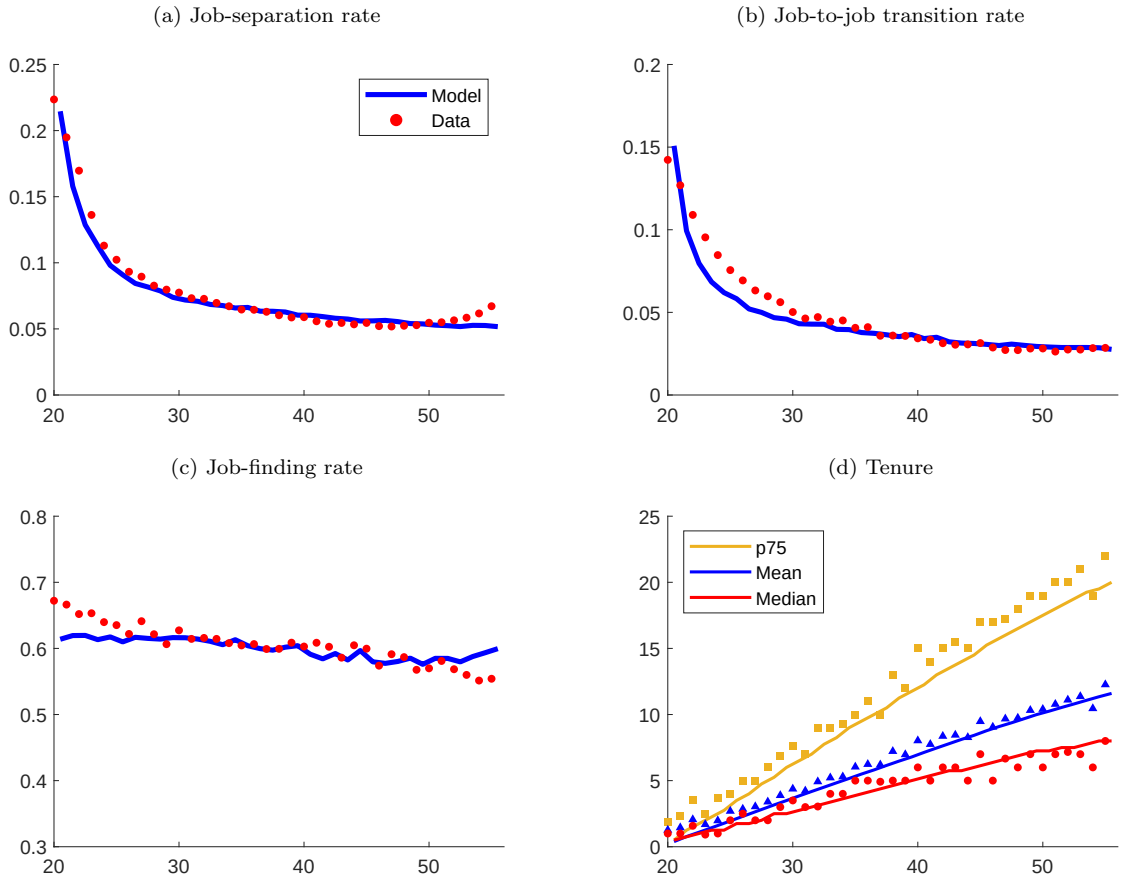
The retirement pattern in the United States is characterized by pronounced spikes at the early eligibility age of 62 and before the full retirement at age 67 under the current Social Security system, with a substantial share of workers retiring in the intervening age range and a smaller fraction delaying retirement beyond the full retirement age. The model captures these key features of the U.S. retirement distribution well, as explained in Section 4.5 below. In Appendix A.14.1, I develop a version of the model with early application penalty for retirement benefits.

Retirement elasticity. I define retirement elasticity as the percentage change in the retirement hazard relative to a change in the generosity of retirement benefits and apply the estimated retirement elasticity by Coile and Gruber (2007). Using the Health and Retirement Study, Coile and Gruber (2007) analyze the effect of pension incentives on retirement behavior. They find an elasticity of retirement with respect to retirement benefits of 0.16 which I target for the retirement elasticity in the model. The policy change takes place for workers at the age of 55 holding the degree of progressivity of the pension system constant at the baseline level. These changes in retirement decisions fall into the range of retirement elasticities typically found in other empirical studies. For example, Moulton and Stevens (2015) follow the methodology of Coile and Gruber (2007) and obtain similar responses in the retirement probability for an increase in Social Security wealth. Brown (2013) uses a quasi-experimental design to provide estimates of the price elasticity of lifetime labor supply. The results in that paper indicate that a 10% change in retirement benefits leads to an adjustment of retirement age by less than two months. Burtless (1986) and Krueger and Pischke (1992) also investigate by how much a change in the Social Security benefit level affects retirement behavior and find a minor role for policy changes. In terms of order of size, the calibrated retirement elasticity is also similar to studies on policy experiments applying structural models. For example, Kimball and Shapiro (2002) use a model of consumption and labor supply to study retirement behavior and find that reducing benefits by 10% induces a postponement of retirement by between 0.1 and 0.5 for workers around age 50. This corresponds to an elasticity of total years worked relative to a change in retirement benefits of between -0.025 and -0.125.

Appendix A.5 provides additional details on the estimation procedure, sensitivity analyses, and parameter identification. For the robustness analysis of the optimal pension system, I also estimate the model separately for college-educated and non-college workers. The corresponding calibration results are reported in Appendix A.10. In what follows, I present the fit of the calibrated model to the targeted life-cycle moments.

4.5 Model fit

Figure 1: Labor market transitions and tenure.



Notes: This figure displays quarterly life-cycle labor market transition rates and tenure in years by age. Panel a shows the average separation rate, Panel b the job-to-job transition rate, and Panel c the job-finding rate. Panel d presents the tenure distribution, including the mean, median, and 75th percentile. In the model, tenure is defined as the duration of continuous employment in the same job, uninterrupted by separations or job-to-job transitions. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical transition rates are constructed using data from the Current Population Survey, and tenure profiles are obtained from the Panel Study of Income Dynamics.

Figure 1 compares the life-cycle profiles of job-separation, job-to-job, job-finding rates, and the tenure distribution in the model and in the data. Both job-separation and job-to-job rates strongly decline between ages 20 and 30 and continue to decline until age 50. The model profiles fit the empirical profiles well with a steep decrease of job-separation rate and job-to-job transition rate at the beginning of the life cycle, followed by a flattening of the profiles after age 30. The empirical profile of job-finding rates also exhibits a decreasing trend over the life cycle which is closely matched by the model. In Figure 1d, I compare the empirical profiles of mean, median, and the 75th percentile of the tenure distribution with their model counterpart. The increasing gap over the life cycle between the 75th percentile, the mean, and the median of the tenure distribution indicates that the tenure distribution gets more dispersed over time. Heterogeneity in employment stability

generates differences in tenure across workers, and these differences grow over the life cycle both in the data and in the model.

Figure 2 displays the profiles of mean log earnings, earnings variance, and wealth-to-income ratio. The model is able to match the strong increase in earnings at the beginning of the life cycle and the decreasing income growth over time when workers get closer to the end of prime age. The model also matches the steep increase of the earnings variance as well as the life-cycle profile of wealth-to-income ratio.

Finally, the model is matched to the distribution of retirement ages in U.S. data. The model captures the key features of the U.S. retirement distribution well (see Figure A.4 in Appendix A.5.2). In particular, it successfully reproduces the prominent peaks at the early and full retirement ages observed in the data.

5 Heterogeneity in employment stability and life-cycle dynamics

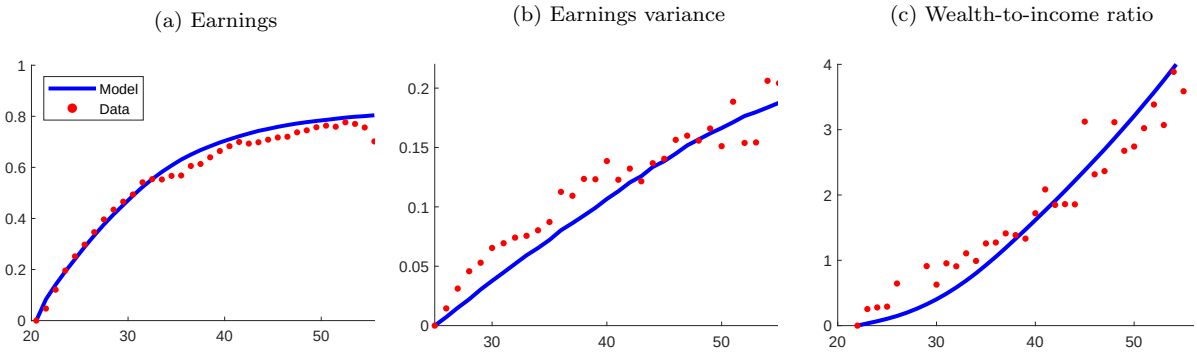
Using the calibrated model, this section studies how heterogeneity in employment stability shapes inequality in labor market outcomes and shows the life-cycle implications of heterogeneity in employment stability for the baseline calibration of the model.

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

5.1 Inequality in labor market outcomes

Figure 3a illustrates how the two groups of agents differ in terms of their average job-separation rate over the life cycle. Comparing the two profiles in Figure 3a, on average, the top quartile already finds more stable jobs at labor market entry. Moreover, the average separation rate of the top quartile remains close to the most stable job which

Figure 2: Earnings and wealth.



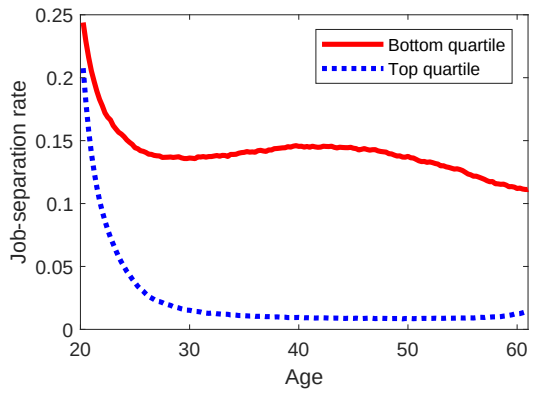
Notes: This figure presents quarterly life-cycle profiles of earnings and the wealth-to-income ratio by age. Panel a displays average log labor earnings, Panel b shows the variance of log labor earnings, and Panel c reports the wealth-to-income ratio. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical profiles of earnings and the wealth-to-income ratio are constructed using data from the Survey of Consumer Finances, while the variance of log earnings is obtained from the Panel Study of Income Dynamics.

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

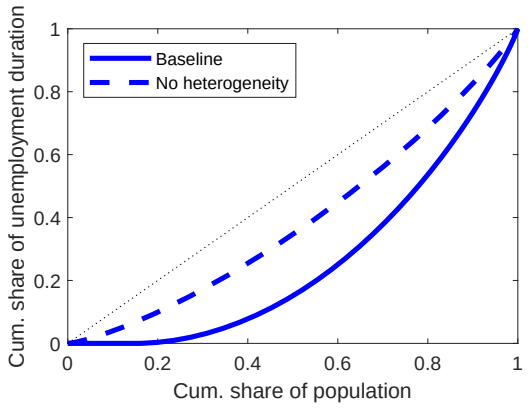
This heterogeneity in employment stability translates into a large inequality in employment history. Figure 3b shows the Lorenz curve for cumulative unemployment duration of prime-age workers between ages 25 and 55. The solid line displays the Lorenz curve for the baseline economy and the dashed line a counterfactual economy without age-specific heterogeneity in employment stability. More specifically, the counterfactual economy is constructed by assuming that all workers have the same age-dependent separation rate from their jobs. The age-dependent separation rate corresponds to the average separation rate conditional on age in the baseline economy with heterogeneity in employment stability. That is, in the counterfactual economy, the separation rates of workers vary over the life cycle, but there is no cross-sectional heterogeneity in employment stability conditional on age. Comparing the baseline economy and the counterfactual economy, the Lorenz curves indicate that heterogeneity in employment stability is a key driver of inequality in employment history. In the baseline economy, about 40% of all workers account for just 8% of total cumulative unemployment duration of all workers, whereas the Lorenz curve

Figure 3: Inequality in labor market history.

(a) Job-separation rate



(b) Lorenz curve of cumulative unemployment duration



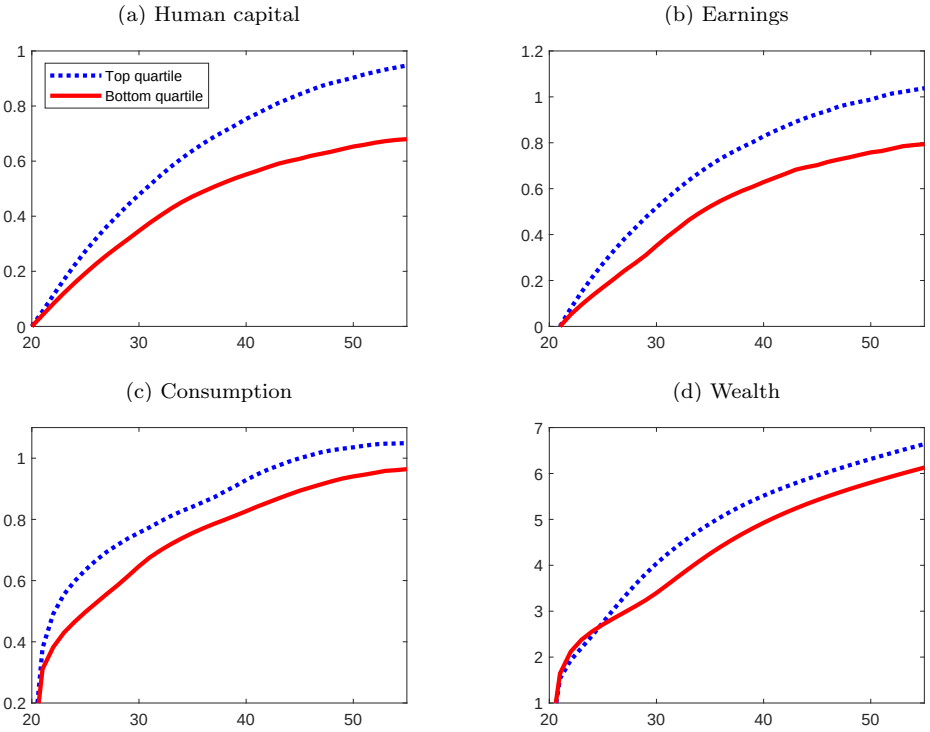
Notes: Panel a shows the average life-cycle profile of job-separation rate. The solid line displays the average life-cycle profile of job-separation rate of workers with most unstable jobs (bottom 25% of employment-stability distribution). The dashed line represents workers with most stable jobs (top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

of the counterfactual does not exhibit a conspicuous curvature.

While there are no fixed worker types in this model, there is a high persistence of the initial employment stability over the life cycle. More specifically, I find that a worker with an unstable job at age 30 (below the median of the employment-stability distribution at age 30) will remain in the unstable group at age 54 (below the median of the employment-stability distribution at age 54) with a probability of 70%.[5] Conversely, a worker with a stable job at age 30 will remain in the stable group at age 54 with a probability of 67%. If the model did not produce any persistence of the initial employment stability for workers, the transition probabilities between the types would be 50%. In Appendix A.7, I show that these transition probabilities are consistent with empirical transition probabilities derived from the Panel Study of Income Dynamics. While the model of this paper does not assume fixed worker types regarding employment stability, the model generates stable and unstable career paths that remain persistent over workers' life cycle. The model's implications are therefore consistent with the empirical literature that considers heterogeneity in employment stability as worker types rather than job types (Ahn et al., 2023; Hall and Kudlyak, 2019; Gregory et al., 2025).

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

Figure 4: Average life-cycle profiles by employment stability.



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

5.2 Life-cycle consequences of employment stability

Figure 4a shows the profiles of average human capital level for workers in the top and bottom quartile of the average life-cycle employment-stability distribution. The profiles are normalized by the initial levels and expressed in log deviations. A higher average job-separation rate directly affects the human capital accumulation process since workers lose the opportunity to invest in human capital upon job loss and the expected return on human capital decreases in job instability. As agents with low employment stability have fewer opportunities to invest in human capital, the profile of the bottom quartile exhibits a lower growth than the profile of the top quartile. Consequently, the gap in the average human capital between the two groups of agents rises over the life cycle. In particular, starting from the same initial level of human capital, the top quartile has, on average, a 22% higher stock of human capital at the end of the working phase.

A similar pattern is observed in Figure 4b for average earnings and in Figure 4c for average consumption. Initially, both groups start from the same level of average earnings, but the profiles quickly diverge during the first five years in which the gap in average separation rate also widens. The earnings and consumption profiles of the bottom quartile

feature a smaller growth over the life cycle.

Figure 4d shows the average life-cycle profiles of wealth. For workers in the top quartile, wealth starts to grow strongly 10 years after labor market entry. The profile of the bottom quartile also increases throughout life, but the growth in wealth is dampened by low earnings growth. Towards the end of the working phase, workers in the bottom quartile have significantly lower assets than agents in the top quartile. Since a high job-separation probability leads to fewer opportunities to invest in human capital and to more frequent job losses, workers in the bottom quartile have, on average, lower earnings which lead to lower savings. Moreover, workers dissave upon job loss in order to smooth consumption, which in turn decreases asset accumulation. The qualitative results remain unchanged when workers are assigned to the stable (top quartile) and unstable group (bottom quartile) using the employment-stability distribution at age 30. In Appendix A.7, I show that using the labor market history in the first 10 years after the labor market entry instead of the entire labor market history up to age 54 yields a similar result as in Figure 4 due to the high persistence of employment stability over the life cycle.

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

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

6 Optimal progressive pension system

Starting from the baseline economy described in the previous section, I derive the optimal pension system. Holding aggregate pension expenditures fixed at the baseline level, I search over the progressivity parameter γ and the level of the Social Security cap to determine the combination that maximizes welfare in the economy. Specifically, I set up a grid over the parameter γ and the cap and, for each pair, compute the corresponding benefit parameter ϕ and the payroll tax rate τ_p that achieve budget balance for the government keeping aggregate pension benefits paid to retirees equal to those in the baseline economy.

The analysis therefore focuses on the optimal allocation of a fixed pool of pension funds. While the long-run sustainability of pension financing is clearly important, the primary objective of this paper is to understand how benefits should be designed in the presence of heterogeneous employment risk. Given the central role of employment stability in shaping lifetime earnings and savings outcomes, the key question is how the pension system can best provide insurance against persistent labor market risk. The next subsection explains the welfare measure used in the policy evaluation.

6.1 Welfare measure

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

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

where $\widetilde{\beta} = \frac{1-\beta^{T^W+T^R+1-j}}{1-\beta}$ and V_b and V_a denote the expected lifetime utility in the baseline economy and in the economy under an alternative pension system, respectively. For $j = 1$, this welfare measure compares the ex-ante expected lifetime utility at labor market entry in the baseline economy to that under the alternative policy. It therefore incorporates expectations over all possible future states and all relevant information over the life cycle. In what follows, the optimal pension system is defined as the system that maximizes ex-ante long-run welfare of households.

6.2 Measure of progressivity level

The degree of pension progressivity depends on several components: the benefit formula, the payroll tax schedule, and the distribution of lifetime earnings in the economy. The progressivity parameter γ in the benefit formula governs how replacement rates vary with average lifetime earnings. However, this parameter alone is not sufficient to characterize the overall progressivity of the pension system. The effective degree of redistribution depends on the net replacement rate, defined as the ratio of retirement benefits to the net average lifetime earnings after deducting payroll taxes, as well as on the Social Security cap, which affects both contributions and benefits. To capture all components that affect the progressivity level of the pension system, I construct the following measure of progressivity.

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

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

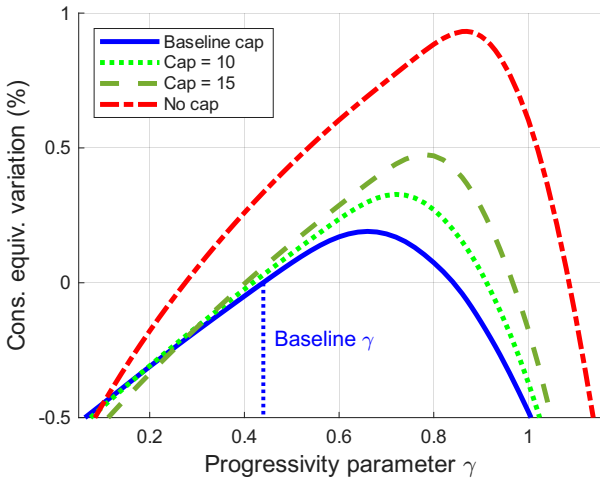
By construction, $\zeta = 0$ if all individuals receive the same net replacement rate. If the pension system is progressive so that workers in the lower quartile receive a strictly higher net replacement rate than those in the upper quartile, then $\zeta > 0$. Larger values of ζ indicate greater redistribution from high- to low-lifetime earners.

6.3 Optimal pension system

Figure 5 reports welfare changes, measured in consumption-equivalent variations, for different combinations of the pension system's progressivity parameter γ and the Social Security cap. A central finding is that, at the baseline level of the Social Security cap, increasing γ and thereby making the pension system more progressive raises welfare. The concave relationship between welfare and the progressivity parameter γ highlights the central trade-off in pension design between redistribution and incentive distortions. Higher pension progressivity provides insurance against unstable employment histories and low lifetime earnings, but it also distorts human capital investment and retirement decisions.

The welfare analysis further shows that substantial gains can be achieved through moderate reforms of the Social Security cap. Raising the cap by 34% yields a welfare gain equivalent to 0.33% of lifetime consumption, while doubling the cap increases welfare by 0.48%. Comparing across different cap levels, the figure shows that welfare increases as the Social Security cap is relaxed, and that fully eliminating the cap leads to the highest welfare gain.

Figure 5: Welfare and changes in progressivity level.



Notes: This figure displays changes in welfare, measured in consumption-equivalent variation, resulting from variations in the progressivity parameter γ and the Social Security earnings cap. The solid line shows welfare changes as a function of γ under the baseline Social Security cap (7.48). The dotted and dashed lines depict welfare changes when the cap is increased to 10 (a 34% increase relative to the U.S. baseline) and 15 (a 100% increase), respectively. The dash-dotted line shows welfare changes when the Social Security cap is entirely removed.

Table 2 reports the policy parameters for the baseline economy and the welfare-maximizing pension system (also see Table A.4 in Appendix A.8). In the baseline economy, the pension system is characterized by $\gamma = 0.44$ and $\phi = 0.70$, with a payroll tax rate of 12.4%. Comparing all policy configurations, welfare is maximized by a pension system without a Social Security cap and with the progressivity parameter set to $\gamma = 0.87$. Relative to the baseline value of 0.44, the optimal progressivity parameter nearly doubles, resulting in a progressivity level of 42.19%. This policy delivers a welfare gain of 0.94% in consumption-equivalent variation for a labor market entrant. The payroll tax rate consistent with budget balance falls to 11.33%, while the corresponding value of ϕ increases to 1.08.

While the welfare analysis above is based on a calibration that targets moments for the overall population, an important question is whether the optimal policy conclusions are sensitive to heterogeneity across education groups. In particular, it is well documented that non-college workers face substantially higher job-separation rates on average than college-educated workers. To assess the implications of this heterogeneity as a robustness check, I recalibrate the model separately for college and non-college workers and repeat the welfare analysis for each group in Appendix A.10. The resulting optimal policy parameters are remarkably similar to those obtained in the benchmark calibration. In both cases, it remains optimal to eliminate the Social Security cap and to adopt a pension benefit schedule that is nearly flat in pre-retirement earnings. However, it is important to note that a more complete analysis would model both education groups jointly within a unified framework. Nevertheless, the robustness exercise suggests that incorporating educational

Table 2: Optimal pension system in the baseline economy and a counterfactual without age-specific heterogeneity in employment stability.

(a) Baseline economy

	Parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline policy	0.44	0.70	7.48	27.61	12.40	-
Optimal policy	0.87	1.08	–	42.19	11.33	0.94

(b) Counterfactual economy without age-specific heterogeneity

	Parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline policy	0.44	0.70	8.83	25.49	12.40	-
Optimal policy	0.70	0.93	–	34.16	11.29	1.02

Notes: Panel 2a shows the welfare results for the baseline model. Panel 2b shows the welfare results for a counterfactual economy without age-specific heterogeneity in employment stability. In the counterfactual model, separation rates are different across ages, but there is no heterogeneity conditional on age. The results for the optimal policy are obtained by removing the Social Security cap and varying the parameter γ starting from the baseline economy. For a given level of γ, the parameter ϕ and the payroll tax rate are always adjusted so as to satisfy budget balance for the government. Columns 2-4 show the parameters γ and ϕ, which determine the shape and the level of pension benefits, respectively, and the Social Security cap. Column 5 presents the degree of progressivity achieved by the combination of policy parameters and Column 6 the payroll tax rate in the economy that achieves budget balance for the pension system. Column 7 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

heterogeneity is unlikely to materially affect the main policy conclusions.

To assess the importance of heterogeneity in employment stability for pension design, I compare the baseline economy to a counterfactual economy without age-specific heterogeneity in employment stability. In Section 5, I showed that heterogeneity in employment stability generates persistent differences in employment histories and accounts for a large fraction of inequality in lifetime earnings. To quantify its importance for pension design, I repeat the welfare analysis in a counterfactual economy without age-specific heterogeneity in employment stability. In this environment, separation rates vary over the life cycle, but there is no age-specific heterogeneity. The counterfactual model is recalibrated to match the same empirical moments as in the baseline specification. However, as shown in Appendix A.11, removing age-specific heterogeneity in separation risk—while allowing only age-dependent or age- and education-dependent separation rates—fails to generate the observed dispersion in employment histories. This dispersion, which is a salient feature of empirical labor market mobility, is the central motivation for the optimal policy analysis of the paper.

Table 2b displays the policy parameters and welfare results for both the optimal and the baseline specification of the pension system. While an increase in the degree of progressivity also leads to welfare gains in the counterfactual economy without age-specific heterogeneity in employment stability, the optimal degree of pension progressivity increases much more strongly in the baseline economy than in the counterfactual economy. In the baseline economy, the progressivity parameter γ almost doubles, rising from 0.44 to 0.87, whereas in the economy without age-specific heterogeneity, it increases by only about 60%. Hence, a disregard of heterogeneity in employment stability leads to a significantly lower degree of optimal progressivity for the pension system.

The policy analysis in this section is conducted in a partial-equilibrium setting where interest rates and wages are fixed. In Appendix A.16, I analyze the optimal pension system when general-equilibrium effects on interest rates and wages are taken into account. Moreover, the results remain robust when the model is extended to incorporate mortality risk, a bequest motive, and early pension penalties in the Social Security system (see Appendix A.14.1). In the baseline analysis, the income tax system and the unemployment insurance policy are held constant. To evaluate the interaction between the pension system and other policy instruments, I extend the analysis in Appendix A.14 by varying the degree of income tax progressivity and the generosity of unemployment insurance benefits. The results show that the welfare gains from implementing the optimal pension system decline as income tax progressivity and the generosity of the unemployment insurance system increase. Nevertheless, the qualitative conclusion remains robust: increasing pension progressivity is always welfare-improving relative to the current system and the welfare gains from a more progressive pension system remain substantial.

In the analysis above, the baseline model assumes that all workers are ex-ante identical in terms of initial wealth, human capital, and labor market status. As a result, the optimal pension reform generates identical welfare changes across individuals in the baseline economy. Introducing ex-ante heterogeneity in state variables, however, may have important implications for the welfare effects and design of the optimal pension system. In Appendix A.13, I examine how the welfare gains from the optimal pension system vary with workers' initial job-separation risk, wealth, and human capital. In Appendix A.14.4, I further extend the model to allow for human capital depreciation. Across these alternative specifications, the qualitative policy conclusions remain unchanged, indicating that the main results on the optimal design of the pension system are robust.

6.4 Welfare decomposition

To better understand the sources of the welfare gains, I decompose the overall welfare effect into a tax component, an insurance component, a retirement component, and a

residual distortionary component. I begin by isolating the role of the lower payroll tax rate under the optimal pension system. To do so, I construct a counterfactual economy in which the optimal pension parameters are implemented, but the payroll tax rate is fixed at its baseline level. Let the welfare gain in this economy be denoted by CEV_{tax}. The contribution of the tax channel is then given by the difference between the total welfare gain and this counterfactual welfare gain, $\text{CEV} - \text{CEV}_{\text{tax}}$.

Next, I isolate the insurance value of the reform. Starting from the previous counterfactual, I endow all workers with the average level of human capital at the end of the working phase. At this level of human capital, replacement rates are identical under the baseline and optimal pension systems, eliminating the redistribution associated with differences in labor market outcomes. By subtracting this welfare change (CEV_{ins}) from the tax component, I obtain the insurance component.

To quantify the retirement channel, I further consider a counterfactual economy in which workers are not allowed to choose their retirement age. The welfare change under this assumption (CEV_{ret}) is then subtracted from the insurance component which yields the retirement component. The welfare loss arising from distortionary effects is then given by the difference between the total welfare effect and the sum of all components.

The reduction in the payroll tax rate under the optimal pension system generates a welfare gain of 0.5%. Abolishing the Social Security cap increases contributions from high-income workers, which permits a lower payroll tax rate while maintaining budget balance. The decline in the tax rate particularly benefits young workers, for whom payroll taxes are especially distortionary: earnings are low early in life, and precautionary savings motives are strong, so reducing payroll taxation alleviates liquidity constraints and raises consumption when marginal utility is high.

Allowing for an endogenous retirement decision generates an additional welfare gain of 0.67%. The extensive-margin adjustment of labor supply enables workers to respond optimally to changes in pension wealth induced by the reform. The insurance component of the optimal pension system accounts for the largest welfare gain, amounting to 1.99%. By strengthening redistribution toward individuals with persistently low earnings and unstable employment histories, the system provides enhanced insurance against adverse labor market outcomes. These gains are partially offset by distortionary effects, which reduce welfare by 2.22%. Under the optimal system, workers invest less in human capital and supply less labor on average, reflecting weaker marginal returns to earnings under the more progressive benefit structure. In sum, the welfare gains from improved insurance and reduced payroll distortions dominate the incentive costs, making a more progressive pension system welfare-enhancing.

Importantly, the model assumes log utility, implying a relatively low degree of risk aversion. As a result, the estimated insurance gains from the optimal pension system can

be interpreted as a lower bound: under higher levels of risk aversion, the welfare benefits of increased pension progressivity would likely be even larger.

6.5 Retirement incentive distortion

The optimal pension system affects workers' retirement decisions and therefore introduces a distortionary effect. This section examines this distortion in greater detail by comparing the average retirement age in the baseline economy with that under the optimal policy.[6] Overall, the labor supply at the extensive margin is reduced as the mean retirement age declines from 64.34 to 64.13. On average, workers choose to retire earlier in response to the increased progressivity of the pension system.

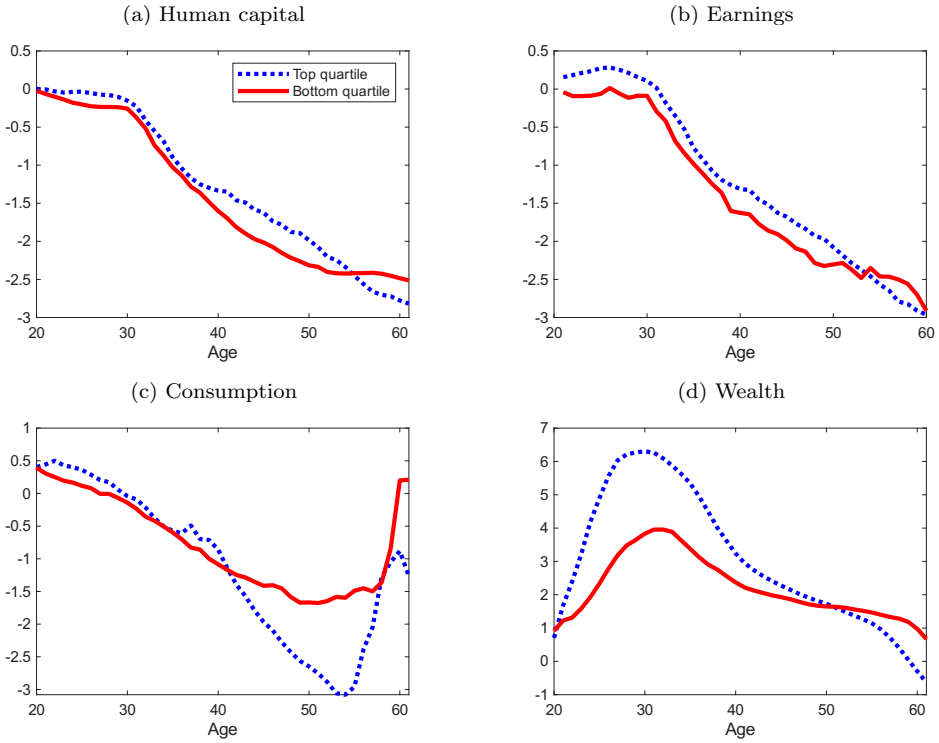
However, the change in retirement behavior varies systematically with workers' human capital. To illustrate this heterogeneity, I compute the mean retirement age for two groups: workers whose final stock of human capital lies below the median of the distribution, and workers whose human capital exceeds the median. Workers in the below-median group retire earlier under the optimal policy than in the baseline economy, with the average retirement age declining from 64.04 to 63.54. In contrast, workers above the median remain in the labor force longer, raising their average retirement age from 64.84 to 65.17.

Once workers become eligible for retirement benefits, they face a choice between retiring and continuing to work. This decision is shaped by two opposing forces. On the one hand, remaining in the labor force generates additional earnings and allows workers to accumulate further human capital, which can increase future retirement benefits. On the other hand, continued employment delays retirement and shortens the period over which benefits are received, thereby reducing the present value of pension wealth. The relative importance of these forces varies across workers depending on their employment status, human capital, and wage level. As a result, the effect of working an additional period on lifetime resources is theoretically ambiguous.

The increase in pension progressivity affects workers differently depending on their level of human capital. For workers with high human capital stocks, the reform reduces pension wealth relative to labor earnings, increasing the incentive to remain in the labor force. As a result, these workers delay retirement in order to continue earning labor income. By contrast, workers who accumulate relatively little human capital experience an increase in pension wealth relative to their potential labor earnings. Since their expected gains from continued employment are limited, it becomes optimal for them to retire earlier. Consequently, higher pension progressivity induces heterogeneous retirement responses across workers, with high-human-capital workers delaying retirement and low-human-capital workers retiring earlier.

[6]The results are summarized in Table A.6 in Appendix A.12.

Figure 6: Relative change in average life-cycle profiles by employment stability.



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

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

6.6 Life-cycle dynamics under the optimal progressive pension system

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

Figure 6 shows the relative deviation from the baseline economy of human capital, labor earnings, consumption, and wealth in the economy under the optimal pension system. The

first finding is that on average, the increase in pension progressivity discourages human capital investment over the life cycle. In Figure 6a, starting at age 30, the relative change in the profiles of average human capital declines steadily until the early retirement age, implying that workers accumulate less human capital under the optimal policy than under the baseline specification.[7] The gap in human capital level grows over life and amounts to -2.5% for agents in the top and -2.8% for agents in the bottom quartile at the early retirement age. Because achieving a high level of human capital is associated with a lower replacement rate in the retirement benefit formula than in the baseline economy, the return on human capital decreases for workers in the top quartile. Consequently, agents employed in stable jobs for whom human capital investment is otherwise large strongly decrease their effort provision for human capital accumulation. A similar reason accounts for lower human capital investment of workers in the bottom quartile.

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

The intuition why the progressivity level of the pension system has little effect on human capital investment of young agents is similar to the idea of Michelacci and Ruffo (2015). They show that an optimal unemployment insurance system should provide higher benefits to younger agents since, in contrast to older workers, the moral hazard problem is small for the young. Michelacci and Ruffo (2015) explain that, in line with the findings in this paper, having a job is important for young workers since employment not only increases current income, but it also provides the opportunity to accumulate human capital and therefore higher income growth. This small distortion in human capital investment among young workers is consistent with empirical evidence suggesting that many individuals are not fully aware of how Social Security benefits are structured

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

(Bairoliya and McKiernan, 2021). Although the model assumes rational expectations, young workers' behavior remains largely unaffected because their short-term labor market prospects matter more than distant changes in retirement benefits.

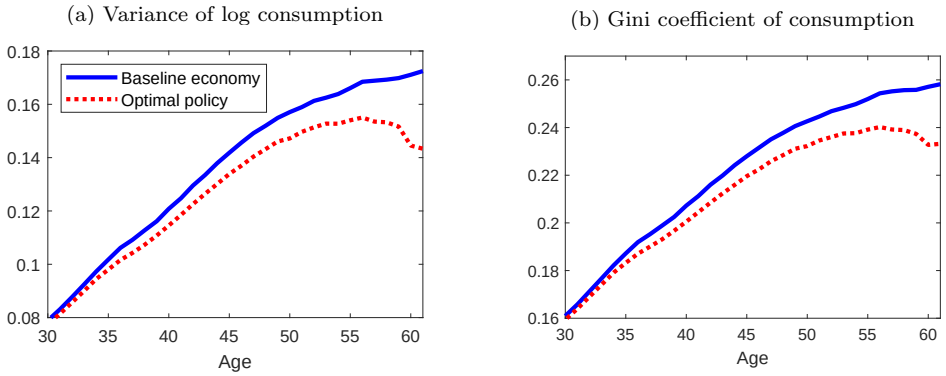
The change in the shape of the average human capital profile consequently affects the average life-cycle profiles of earnings and consumption displayed in Figure 6b and Figure 6c. The first observation is that the relative changes in the profiles of average earnings closely follow the shape of the profiles in Figure 6a. For both groups of workers, the decrease in the average human capital implies lower average earnings compared to the baseline economy. The average consumption of workers in the top quartile decreases towards the later stage of the working phase. Also for the bottom quartile, the average consumption decreases because the average earnings become lower. However, the average consumption strongly increases towards the end of the working phase in the new economy. The reason is that the optimal pension system leads to an increase in pension wealth for workers with unstable jobs. The wealth effect induces these workers to raise their consumption close to retirement. This finding indicates that the increase in pension progressivity provides insurance to workers who face on average the highest level of job instability throughout life by redistributing resources to these workers. Because the pension system becomes more redistributive, agents in the bottom quartile are able to increase their consumption relative to their pre-retirement earnings in the late stage of the life cycle.

Figure 6d reflects this change in the consumption-saving behavior. For both groups of workers, the average asset level increases at the beginning of the life cycle, but declines and reaches the baseline level at the end of the working phase. Workers in the bottom quartile of the employment-stability distribution reduce their asset accumulation close to retirement as more resources are available to them in the retirement phase. Life-cycle consumption smoothing implies that these agents reduce their retirement savings such that the life-cycle profile of average assets becomes flatter under the optimal pension system.

Figure 7 displays that the optimal pension system reduces consumption inequality over the life cycle. The profile for variance of log consumption in the baseline economy almost linearly increases over life by more than 50% reaching a value of 0.17 before retirement. In the economy under the optimal pension system, the profile becomes flatter and has a concave shape. At the end of the working phase, the variance of log consumption reaches a value of 0.14. The optimal pension system therefore dampens the rise in consumption inequality by almost one-third relative to the baseline economy.

Comparing the profiles of the Gini coefficient of consumption in Figure 7b, a similar change as for the profile of variance of log consumption is observed. Whereas the Gini coefficient increases linearly in the baseline model and reaches a value of 0.26 before retirement, the optimal pension system dampens the increase in the Gini coefficient over

Figure 7: The life-cycle profiles for variance of log consumption and Gini coefficient of consumption.



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

life. In the economy under the optimal pension system, the Gini coefficient of consumption at the end of the working phase decreases to 0.23.

6.7 Transition to the optimal pension system

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

More specifically, the progressivity parameter γ, the scale parameter ϕ, and the Social Security cap are gradually increased by a constant percentage over a 20-year horizon. Throughout the transition, the tax rate adjusts endogenously in every year to ensure period-by-period budget balance, so that any changes in pension benefits paid in a given year are fully financed by taxes collected in that same year. The whole transition path including the adjustments of tax rates over time takes approximately 70 years. The evolution of both the policy parameters and the tax rate along the transition path is reported in Appendix A.15. In year zero, the baseline pension system is in place. Beginning in year one, the government implements the new policies with higher parameter values. As in Auerbach and Kotlikoff (1987), all workers have perfect foresight along the transition path regarding the pension system that will be in effect when their cohort reaches retirement age.

The welfare changes for the different cohorts are shown in Figure 8. These welfare

changes are evaluated in the period the government credibly announces the full transition path toward the optimal pension system. As the policy parameters are gradually changed for each cohort and the remaining lifetime until retirement differs across cohorts, the transition to the optimal policy generates heterogeneous welfare effects across age groups. Most workers younger than age 48 experience positive welfare gains, which converge to 0.94% for labor market entrants in the new steady state under the optimal pension system.

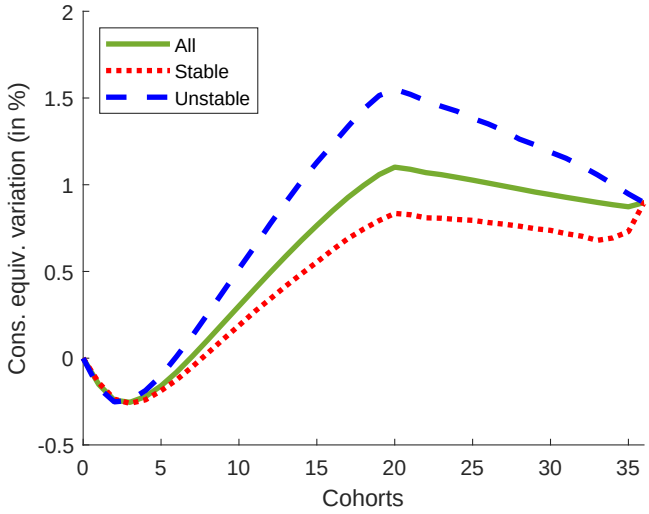
To analyze the differential welfare effects along the employment-stability distribution, I consider the bottom quartile (unstable) and top quartile (stable) of the employment-stability distribution of each cohort. As in the previous sections, this distribution considers the mean job-separation rates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

Older cohorts close to retirement when the transition to the optimal policy takes place experience negative welfare effects during the transition. As shown in Figure 8, cohorts 0-5 (households at age 50-55) exhibit welfare losses in terms of consumption-equivalent variation, reflecting their limited ability to adjust labor supply and savings in response to the reform. The welfare losses for older cohorts arise because these workers have little time to benefit from the long-run insurance and redistribution effects of higher progressivity. Overall, the results highlight that the transitional welfare effects of higher progressivity are strongly heterogeneous across cohorts and employment stability, with job-specific risk playing a central role in shaping welfare outcomes.

7 Conclusion

This paper studies how heterogeneity in employment stability shapes the optimal design of pension systems and quantifies the welfare gains from implementing the optimal pension system.

Figure 8: Welfare along the transition path.



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

A key contribution is the development of a life-cycle model that integrates heterogeneous labor market risk with endogenous human capital accumulation, consumption–saving behavior, and retirement decisions in a frictional labor market. The framework provides a unified structure in which labor market risk interacts with earnings dynamics, savings choices, and retirement behavior over the life cycle. Capturing these interactions is crucial for understanding how heterogeneity in employment stability shapes optimal pension design. Heterogeneity in employment stability generates persistent dispersion in lifetime earnings and savings outcomes through the endogenous human capital channel: Workers with unstable employment lose the opportunity to climb the job ladder, accumulate less human capital and wealth over the life cycle, leading to long-run consequences of job separations.

The model is calibrated to the U.S. economy and the institutional features of the Social Security system, targeting empirical life-cycle profiles of labor market transition rates, earnings (mean and variance), the tenure distribution, and wealth-to-income ratios. The calibrated model closely matches these key moments. Using the calibrated model, I show that heterogeneity in employment stability translates into a large inequality in labor market outcomes and is a key driver of inequality in lifetime earnings.

The numerical analysis shows that increasing pension progressivity toward a benefit schedule that closely resembles a benefit floor maximizes ex-ante welfare. Welfare gains are further enhanced by raising the Social Security cap: increasing the cap by 34% generates a welfare gain equivalent to 0.33% of lifetime consumption for labor market

entrants. Eliminating the cap altogether yields the largest welfare improvement, increasing welfare by 0.94% in consumption-equivalent terms. While greater pension progressivity reduces human capital investment and induces modest adjustments in retirement behavior, these distortionary effects are small relative to the insurance gains. The optimal reform substantially reduces life-cycle consumption inequality and provides valuable insurance against persistent employment instability and low lifetime earnings.

Overall, the analysis shows that the optimal pension system is substantially more progressive than the current U.S. Social Security system, highlighting the importance of heterogeneous labor market dynamics for the design of old-age insurance. Persistent differences in employment stability shape earnings trajectories, wealth accumulation, and retirement behavior, thereby influencing the optimal degree of insurance. Accounting for labor market dynamism and the associated life-cycle interactions is therefore important when evaluating pension reforms and should also be considered in the design of other social insurance policies.

References

Ahn, H. J., B. Hobijn, and A. Şahin (2023): "The Dual U.S. Labor Market Uncovered," Working Paper 31241, National Bureau of Economic Research.

Auerbach, A. J. and L. J. Kotlikoff (1987): "Evaluating fiscal policy with a dynamic simulation model," *The American Economic Review*, 77, 49–55.

Bairoliya, N. and K. McKiernan (2021): "Revisiting retirement and social security claiming decisions," *Available at SSRN*, 3896031.

Becker, G. S. (1964): "Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education." University of chicago press.

Ben-Porath, Y. (1967): "The production of human capital and the life cycle of earnings," *Journal of Political Economy*, 75, 352–365.

Bilal, A. (2023): "The Geography of Unemployment," *The Quarterly Journal of Economics*, 138, 1507–1576.

Bloom, D. E., D. Canning, R. K. Mansfield, and M. Moore (2007): "Demographic change, social security systems, and savings," *Journal of Monetary Economics*, 54, 92–114.

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

Bonhomme, S. and J.-M. Robin (2009): "Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990–2000," *The Review of Economic Studies*, 76, 63–92.

Borella, M., M. De Nardi, and F. Yang (2023): "Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labour Supply?" *The Review of Economic Studies*, 90, 102–131.

Bowlus, A. J. and J.-M. Robin (2012): "An International Comparison of Lifetime Inequality: How Continental Europe Resembles North America," *Journal of the European Economic Association*, 10, 1236–1262.

Brown, K. M. (2013): "The link between pensions and retirement timing: Lessons from California teachers," *Journal of Public Economics*, 98, 1–14.

Burdett, K. and D. T. Mortensen (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–273.

Burtless, G. (1986): "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *The Review of Economic Studies*, 53, 781–805.

Bénabou, R. (2000): "Unequal Societies: Income Distribution and the Social Contract," *American Economic Review*, 90, 96–129.

——— (2002): "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70, 481–517.

Bönke, T., G. Corneo, and H. Lüthen (2015): "Lifetime Earnings Inequality in Germany," *Journal of Labor Economics*, 33, 171–208.

Coile, C. and J. Gruber (2001): "Social Security Incentives for Retirement." National Bureau of Economic Research, Working Paper 7830.

——— (2007): "Future Social Security Entitlements and the Retirement Decision," *The Review of Economics and Statistics*, 89, 234–246.

Cremer, H., J.-M. Lozachmeur, and P. Pestieau (2004): "Social security, retirement age and optimal income taxation," *Journal of Public Economics*, 88, 2259–2281.

Davis, S. J. and T. M. von Wachter (2011): "Recessions and the Cost of Job Loss," Working Paper 17638, National Bureau of Economic Research.

De Nardi, M. (2004): "Wealth Inequality and Intergenerational Links," *The Review of Economic Studies*, 71, 743–768.

De Nardi, M., S. Imrohoroglu, and T. Sargent (2001): "Saving and pension reform in general equilibrium models," *Oxford Review of Economic Policy*, 17, 20–39.

Disney, R., C. Emmerson, and M. Wakefield (2006): "Ill health and retirement in Britain: a panel data-based analysis," *Journal of Health Economics*, 25, 621–649.

Dwyer, D. S. and O. S. Mitchell (1999): "Health problems as determinants of retirement: are self-rated measures endogenous?" *Journal of Health Economics*, 18, 173–193.

Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2022): "Job Search Behavior Among the Employed and Non-Employed," *Econometrica*, 90, 1743–1779.

Fehr, H. and C. Habermann (2008): "Risk Sharing and Efficiency Implications of Progressive Pension Arrangements," *The Scandinavian Journal of Economics*, 110, 419–443.

Fehr, H., M. Kallweit, and F. Kindermann (2013): "Should pensions be progressive?" *European Economic Review*, 63, 94–116.

Feldstein, M. (1974): "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy*, 82, 905–926.

Galama, T., A. Kapteyn, R. Fonseca, and P.-C. Michaud (2013): "A health production model with endogenous retirement," *Health Economics*, 22, 883–902.

Golosov, M., A. Shourideh, M. Troshkin, and A. Tsyvinski (2013): "Optimal Pension Systems with Simple Instruments," *American Economic Review*, 103, 502–507.

Gregory, V., G. Menzio, and D. G. Wiczer (2025): "The Alpha Beta Gamma of the Labor Market," *Journal of Monetary Economics*, 150, 103695, Labor Market Dynamics: Drivers and Consequences.

Gruber, J. and D. A. Wise (2004): "Social Security Programs and Retirement around the World: Micro-Estimation", University of Chicago Press.

Gulyas, A. and K. Pytka (2021): "Understanding the sources of earnings losses after job displacement: A machine-learning approach," Tech. rep.

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

Haan, P., D. Kemptner, V. Prowse, and M. Schaller (2025): "Insurance, redistribution, and the inequality of lifetime income," *Quantitative Economics*, 16, 565–613.

Haan, P. and V. Prowse (2014): "Longevity, life-cycle behavior and pension reform," *Journal of Econometrics*, 178, 582–601.

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

Hall, R. and M. Kudlyak (2019): "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics," Working Paper 25625, National Bureau of Economic Research.

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

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

Hendricks, L. (2004): "Taxation and Human Capital Accumulation," *Macroeconomic Dynamics*, 8, 310–334.

Hubbard, R. G. and K. L. Judd (1987): "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax," *The American Economic Review*, 77, 630–646.

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

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993): "Earnings Losses of Displaced Workers," *The American Economic Review*, 83, 685–709.

Jarosch, G. (2023): "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, 91, 903–942.

Jones, A. M., N. Rice, and J. Roberts (2010): "Sick of work or too sick to work? Evidence on self-reported health shocks and early retirement from the BHPS," *Economic Modelling*, 27, 866–880.

Jung, P. and M. Kuhn (2019): "Earnings Losses and Labor Mobility Over the Life Cycle," *Journal of the European Economic Association*, 17, 678–724.

Keane, M. P. and K. I. Wolpin (1997): "The Career Decisions of Young Men," *Journal of Political Economy*, 105, 473–522.

Kimball, M. S. and M. D. Shapiro (2002): "The Retirement Elasticity: Theory and Implications." Tech. rep.

Krueger, A. B. and J.-S. Pischke (1992): "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics*, 10, 412–437.

Krusell, P. and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Kuhn, M., L. Nam, and G. Ploj (2025): "Employment Stability, Earnings Dynamics, and Life-Cycle Savings." Tech. rep.

Larkin, K. (2019): "Job risk, separation shocks and household asset allocation," in *2019 Meeting Papers*, Society for Economic Dynamics, 1058.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432–1467.

McCall, J. J. (1970): "Economics of Information and Job Search," *The Quarterly Journal of Economics*, 84, 113–126.

Michelacci, C. and H. Ruffo (2015): "Optimal Life Cycle Unemployment Insurance," *American Economic Review*, 105, 816–859.

Molloy, R., C. L. Smith, and A. Wozniak (2024): "Changing stability in US employment relationships: A tale of two tails," *Journal of Human Resources*, 59, 35–69.

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

Moulton, J. G. and A. H. Stevens (2015): "Reconsidering the Social Security Notch and retirement: Wealth and incentive effects," *Economics Letters*, 132, 65–68.

O'Dea, C. (2018): "Insurance, Efficiency and the Design of Public Pensions," 2018 Meeting Papers 1037, Society for Economic Dynamics.

OECD (2019): *Pensions at a Glance 2019*.

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

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.

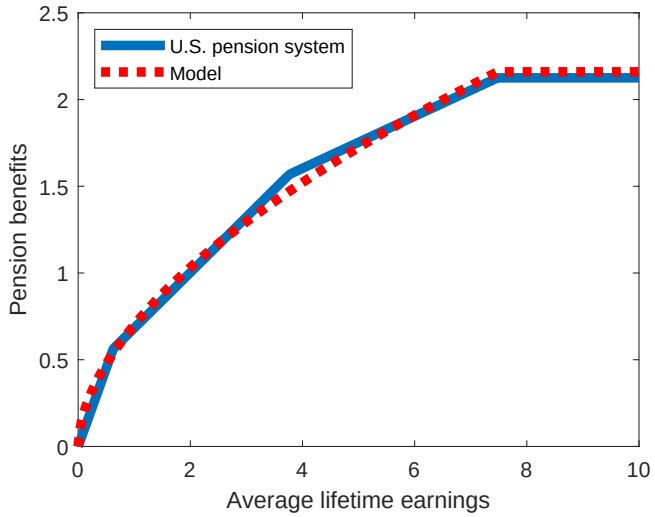
Topel, R. H. and M. P. Ward (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107, 439–479.

Trevisan, E. and F. Zantomio (2016): "The impact of acute health shocks on the labour supply of older workers: Evidence from sixteen European countries," *Labour Economics*, 43, 171–185.

A Appendix

A.1 Calibration of the pension system

Figure A.1: Pension System.



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

Figure A.1 illustrates the piecewise-linear benefit formula of the U.S. Social Security system (solid line). The dotted line represents the model's approximation, given by the functional form

$$\omega(\bar{y}) = \begin{cases} \phi \cdot (\bar{y})^{1-\gamma} & \text{if } \bar{y} < cap^p \\ \overline{\omega} & \text{if } \bar{y} \geq cap^p \end{cases} \tag{7}$$

where cap^p denotes maximum taxable labor earnings and $\bar{\omega} = \phi \cdot (cap^p)^{1-\gamma}$ if average lifetime earnings exceed the cap.

A.2 Logistic distribution and expected utility

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

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

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

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

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

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

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

$$\mathbb{E}V_{\max} = pV_1 + (1-p)V_2 + \int_{\delta}^{\infty} \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon \tag{8}$$

where

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

Integration by parts yields

$$\int_{\delta}^{\infty} \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon = \delta \cdot \frac{\exp\left(-\frac{\delta-\mu}{\sigma}\right)}{1+\exp\left(-\frac{\delta-\mu}{\sigma}\right)} + \sigma \cdot \log\left(1+\exp\left(-\frac{\delta-\mu}{\sigma}\right)\right). \tag{9}$$

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

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

I use these expressions to simplify Equation (9) and get

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

Using this result, Equation (8) can be simplified to

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

which yields Equation (2) in Section 3.1.

A.3 Model solution

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

A.4 Human capital

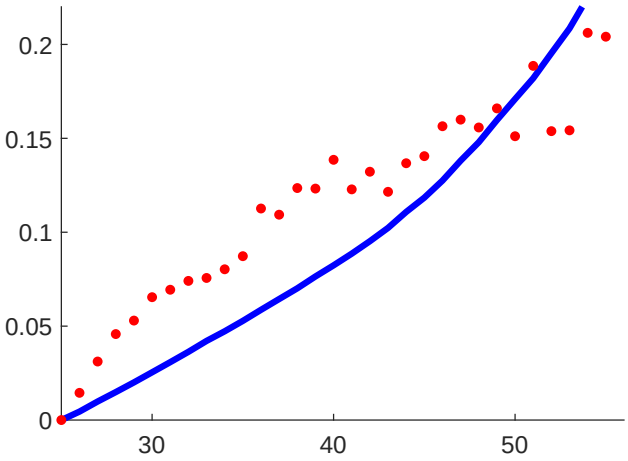
In this section, I discuss an alternative specification for the human capital accumulation process. In particular, I consider a human capital accumulation process that corresponds to the Ben-Porath (1967) model. While I leave all other parts of the model unchanged, I assume that human capital investment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

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

where t is the time invested into human capital.

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

Figure A.2: Earnings variance.



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

A.2 shows that compared to the baseline model, the increase in earnings variance over the life cycle becomes linear up to age 40 and exhibits a convex shape afterwards. The interaction between employment stability and the human capital accumulation explains why the increase in earnings variance is more pronounced if human capital investment entails monetary costs. Workers with unstable employment have, on average, lower income levels and at the same time, these workers have a large precautionary motive to save due to the high income risk they face. As a result, human capital investment is even costlier for workers with unstable employment if it enters the budget constraint. Lower investment leads to lower human capital accumulation of these workers, and consequently, the increase in earnings dispersion becomes stronger in the model with monetary cost of human capital investment.

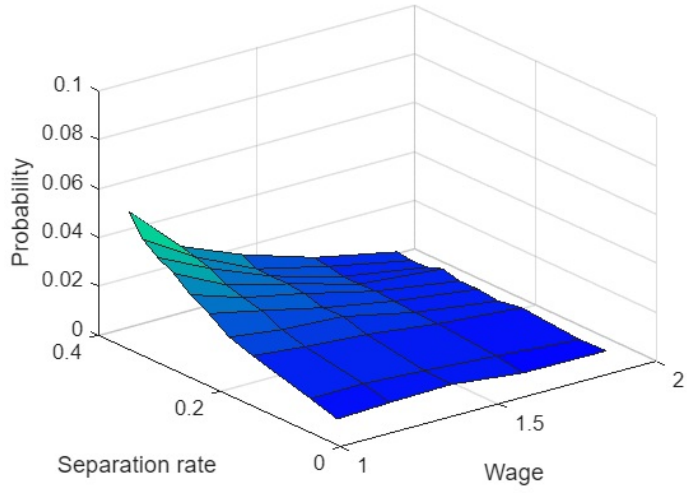
While the model with Ben-Porath (1967) human capital accumulation also matches the data moments well, I assume that human capital investment leads to utility costs for the baseline economy so as to match the concavity of the earnings variance profile over the life cycle.

A.5 Calibration

A.5.1 Job-offer distribution

The number of grid points for wages is set to $N_w = 5$ where $\underline{w} = 1$ and $\overline{w} = 1.85$. All grid points are equidistant in logs. Concerning the job-separation probabilities, it is assumed that $N_\lambda = 10$ with $\underline{\lambda} = 0.006$ and $\overline{\lambda} = 0.35$. All remaining grid points are located non-linearly between these two values. The marginal distributions for wages and

Figure A.3: Joint distribution of wages and separation rates.



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

employment stability are assumed to have a truncated exponential distribution on the supports $[\underline{w}, \overline{w}]$ and $[1-\overline{\lambda}, 1-\underline{\lambda}]$, respectively. The joint distribution of job-separation probability and wage is determined by mapping the supports of these two random variables to the unit interval $[0,1]$. Let $w^* \in [0,1]$ denote the standardized wage level and let $1-\lambda^*$ denote the standardized employment stability. The density of each of these standardized variables is then given by

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

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

Figure A.3 shows the estimated distribution of wages and job-separation rates. In general, there is a positive relationship between employment stability and wages. The most stable jobs have high wages and vice versa. The vast majority of the probability mass lies in the area of high separation rate, reflecting the scarcity of lifetime jobs.

A.5.2 Parameters

I estimate the model parameters using a simulated method of moments. I minimize the sum of squared percentage deviations of the life-cycle profiles produced by the model from the empirical counterparts. The empirical moments include the life-cycle profiles of separation and job-finding rate as well as the transition rate from job to job, mean

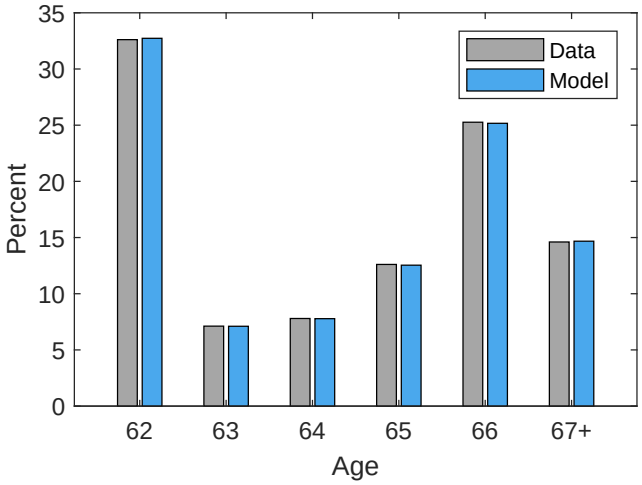
and variance profiles of earnings, the mean, median, and 75th percentile of the tenure distribution, and the wealth-to-income ratio. Let θ_p denote the vector of parameters and let a denote age. Then, the objective function is

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

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

Figure A.4 plots the distribution of retirement age in the data for the United States in 2019 and in the baseline calibration of the model. The retirement pattern in the United States exhibits two huge peaks at age 62 and at age 66; around 32.6% of workers entitled to retirement benefits in 2019 were at age 62, which is the early eligibility age of the current U.S. Social Security rule; around 25.3% were at age 66. The proportion of retired workers between ages 63 and 65 was around 27.5%. Around 14.5% were at age 67 or older. Figure A.4 shows that the specification of the model matches the U.S. retirement pattern well. In particular, the model is able to fit the huge peaks at ages 62 and 66.

Figure A.4: Retirement age distribution in the data and in the baseline model.



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

A.5.3 Identification of parameters

In the following, I discuss the identification of parameters and provide a detailed mapping between parameter blocks and the empirical moments that discipline them. Because all parameters interact in the model, I clarify the economic channels through which parameters shape observable life-cycle patterns and demonstrate how targeted moments discipline them. The explanation is summarized in Table A.1.

The calibration targets three sets of moments that correspond to the main building blocks of the model: labor market dynamics, human capital accumulation, and savings behavior. Labor market moments identify the parameters governing job mobility and sorting across jobs with different wages and separation risks. Life-cycle earnings profiles identify the parameters governing human capital accumulation. Finally, wealth accumulation over the life cycle identifies workers' intertemporal preferences. In the following, I provide a comparative statics experiment where I shift parameter values and analyze how this affects the corresponding life-cycle profiles in the model.

Labor market dynamics. The first block of parameters governs worker mobility and sorting in the labor market. The job-offer arrival rates, π_e and π_u, determine the job-arrival rates with which employed and unemployed workers receive job offers. Higher offer arrival rates increase job mobility and accelerate sorting into better jobs. These parameters are identified using the life-cycle profiles of employment-to-employment and unemployment-to-employment transition rates. Figure A.5 illustrates how job-to-job transition rate and job-finding rate change in response to a change in the job-offer arrival rates. When the

Table A.1: Mapping between model parameters and empirical moments

Parameter		**Key Identifying Moments**
π_e	Job offer arrival rate when employed	Job-to-job transition rate by age
π_u	Job offer arrival rate when unemployed	Job-finding rate by age
ψ_λ	Distribution shape of separation rate	Separation rate; tenure distribution (level)
ψ_w	Distribution shape of offered wages	Mean earnings profile in early age; tenure distribution (dispersion at young age)
θ	Correlation between wages and separation rate	Joint evolution of separations and tenure distribution (dispersion at the end of prime age)
ρ	Persistence of human capital	Earnings dispersion by age
p_H, p_H^*	Human capital upgrading probabilities	Earnings level and shape of earnings dispersion over the life cycle
κ	Utility cost of skill accumulation effort	Late-career flattening of mean earnings
β	Discount factor	Wealth-to-income profile by age

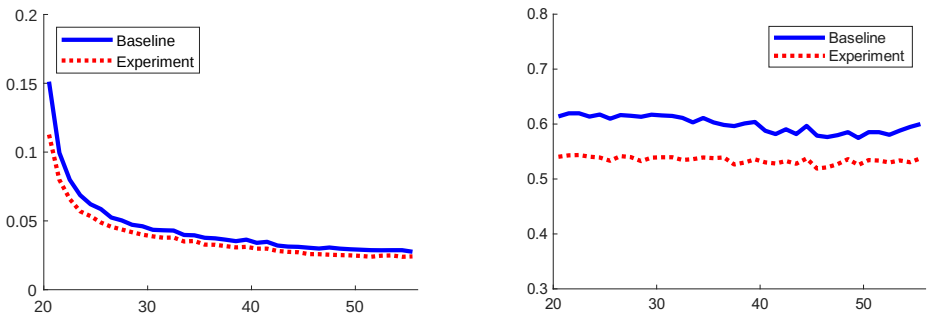
job-arrival rates π_e and π_u are reduced by 10 percentage points, the life-cycle profiles of both the job-to-job transition rate and the job-finding rate shift downward, reflecting the lower frequency of job offers. The decline in the job-to-job transition rate is particularly pronounced for young workers, for whom job mobility plays a central role in climbing the job ladder early in the career. As job offers become less frequent, job-to-job transitions decrease especially at younger ages.

The remaining labor-market parameters determine the joint distribution of wage offers and job stability, $f(w, \lambda)$. The parameter ψ_λ governs the distribution of job stability and is identified from the life-cycle profile of separation rates together with the level of the tenure distribution. The parameter determines how much mass of job offers is concentrated in relatively unstable versus stable jobs. Figure A.6 shows that a distribution with relatively fewer highly stable matches slows down early-career sorting into long-lasting employment relationships. Young workers are more likely to accept jobs with higher separation risk and experience repeated job loss before gradually moving into more stable matches. This generates higher overall separation rates and lower average tenure, particularly at younger ages. Over time, slower sorting also delays the accumulation of long-tenure matches, flattening the tenure profile across the life cycle.

The parameter ψ_w governs the distribution of wage offers and is identified using the slope and curvature of life-cycle earnings growth as well as early-career tenure dispersion. Economically, ψ_w shapes the extent to which workers can climb the job ladder through job-to-job mobility. As illustrated in Figure A.7, a lower arrival rate of high-wage opportunities

Figure A.5: Change in job-arrival rates π_e and π_u

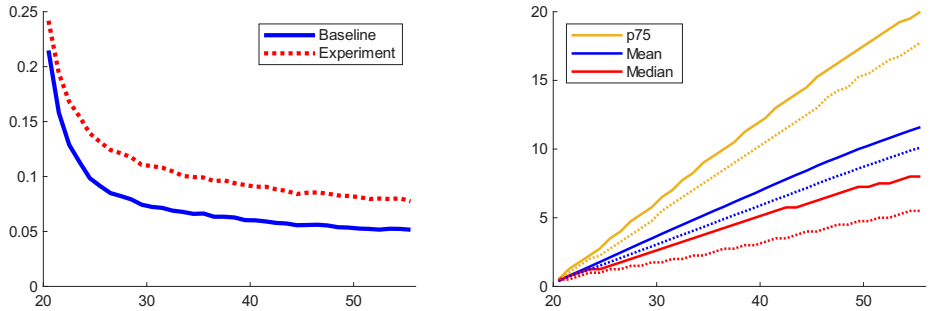
(a) Job-to-job rate to change in π_e (b) Job-finding rate to change in π_u



Notes: These figures document the sensitivity of transition rates to parameter perturbations in the job-arrival probabilities π_e (reduced from 0.41 to 0.31) and π_u (reduced from 0.81 to 0.71).

Figure A.6: Change in job-offer distribution shape parameter ψ_λ

(a) Separation rate (b) Tenure distribution

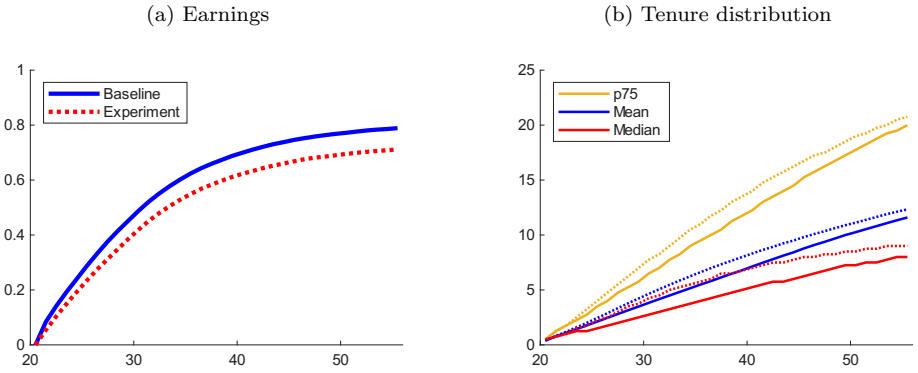


Notes: These figures document the sensitivity of separation rate and tenure distribution to a parameter perturbation in the job-offer distribution shape parameter ψ_λ (increased from 0.49 to 0.94).

slows wage growth and delays movement up the job ladder. Young workers, who rely heavily on job-to-job transitions to achieve rapid wage growth, experience smaller wage gains from mobility. As a result, earnings growth between ages 20 and 30 becomes flatter. In addition, slower wage upgrading delays sorting into better-paying and potentially more stable jobs, which also affects early-career tenure dispersion.

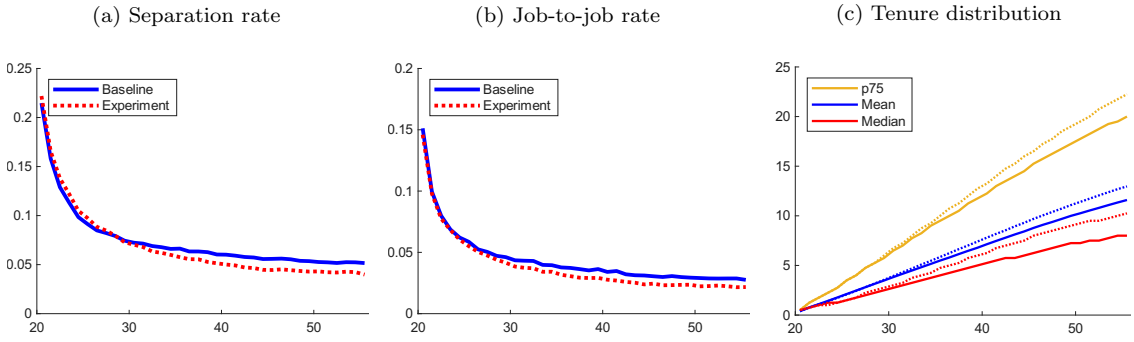
The correlation parameter θ captures the relationship between wages and job stability and is identified from the joint evolution of mobility rates, separations, and tenure dispersion over the life cycle. Economically, θ determines the extent to which higher-paying jobs are also more stable. A stronger positive correlation implies that workers who move up the wage ladder simultaneously sort into jobs with lower separation risk. This generates a tighter link between wage growth, declining mobility, and declining separation rates over the life cycle. It also affects tenure dynamics, as workers who sort into high-paying jobs remain employed in those matches for longer periods, increasing the concentration of tenure at older ages. Figure A.8 illustrates how changes in θ affect these joint patterns. My estimate of the correlation parameter, 0.34, implies a positive

Figure A.7: Change in job-offer distribution shape parameter ψ_w



Notes: These figures document the sensitivity of earnings and tenure distribution to a parameter perturbation in the job-offer distribution shape parameter ψ_w (increased from 0.56 to 1.35).

Figure A.8: Change in parameter θ for correlation between wages and separation rates



Notes: These figures document the sensitivity of separation rate, job-to-job rate, and tenure distribution to a parameter perturbation in the correlation θ between wages and employment stability (increased from 0.34 to 0.74).

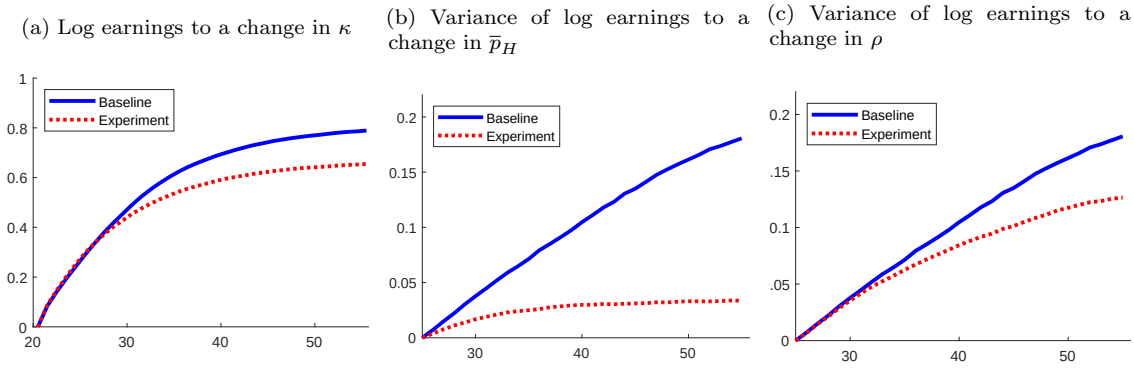
relationship between wages and job stability, or equivalently, a correlation of -0.34 between wages and separation rates. The estimated correlation is remarkably close to the value of -0.33 reported by Jarosch (2023). In both frameworks, higher-paying jobs are more stable, consistent with the presence of a job ladder in which workers progressively sort into better-paid and more durable matches.

Human capital accumulation. The second block of parameters governs earnings growth through human capital accumulation. The parameters $\overline{p}_H$ and ρ determine the level and persistence of skill upgrading and are identified using the life-cycle profiles of mean earnings and earnings dispersion. As shown in Figures A.9b and A.9c, these parameters affect both the growth of earnings and the evolution of within-age inequality over the life cycle.

The utility cost parameter κ governs workers' incentives to invest in human capital and is identified from the late-career flattening of the earnings profile. Figure A.9a illustrates that higher investment costs reduce human capital accumulation and lead to slower

earnings growth at older ages, when wage growth is primarily driven by skill acquisition rather than job mobility.

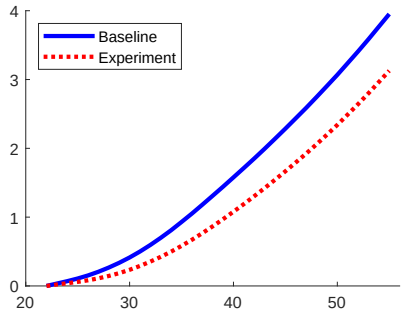
Figure A.9: Change in parameters governing human capital κ, $\overline{p}_H$, and ρ



Notes: These figures document the sensitivity of log earnings and variance of log earnings to parameter perturbations in the parameters governing human capital accumulation κ (increased from 0.34 to 0.54), $\overline{p}_H$ (reduced from 0.055 to 0.035), and ρ (reduced from 0.985 to 0.983).

Preferences and savings. The final parameter is the discount factor β, which governs intertemporal consumption and savings decisions. It is identified using the life-cycle profile of the wealth-to-income ratio. As shown in Figure A.10, a lower value of β increases impatience, leading workers to discount future utility heavily. As a result, agents place greater weight on current consumption and reduce their retirement savings and the life-cycle asset profile shifts downward.

Figure A.10: Change in time discount parameter β



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

A.5.4 Job-offer arrival rates

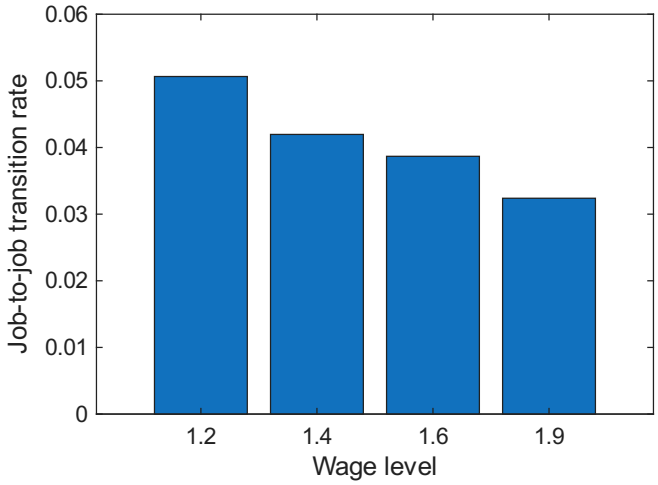
The model assumes that job-arrival rates are identical across employed workers, which raises the question of how well this assumption aligns with the data. Empirically, it is difficult to assess whether job-arrival rates vary across the wage distribution because job

offers themselves are typically unobserved. Instead, the available evidence comes from realized labor market outcomes, such as job-to-job transitions and employment durations. Accordingly, the model is disciplined by matching these observable moments, including job-to-job transition rates and the life-cycle profile of job tenure.

Existing evidence suggests that search behavior differs systematically by wage level. In particular, job search intensity among employed workers declines strongly with current wages (Faberman et al., 2022). At the same time, it is not obvious whether this implies that high-wage workers receive fewer raw job contacts, as opposing forces may be at play: while high-wage workers search less actively, they may benefit from stronger professional networks and receive more unsolicited offers. For the purpose of job mobility, however, the relevant object is not the raw arrival rate of contacts, but the arrival rate of acceptable job offers. In line with standard wage-posting models (e.g., Burdett and Mortensen, 1998), the current model implies that workers with higher current wages have higher reservation wages and therefore accept only a small subset of potential offers. As a result, even if contact rates were the same across workers conditional on employment, the effective arrival rate of acceptable offers declines with wages.

This mechanism is visible in the job-offer distribution in the model. For a worker in the bottom 10% of the wage distribution, the probability that a randomly drawn offer pays at least as much as the current wage is about 68%. For a worker in the top 20% of the wage distribution, this probability declines sharply to 9.6%. Consistent with this pattern, observed job-to-job transition rates decline with the current wage.

Figure A.11: Job-to-job transition conditional on wage.



Notes: This figure shows the quarterly job-to-job transition rate conditional on current wage of workers.

Figure A.11 shows the job-to-job transition rates by wage in the baseline model. Job-to-job transition rate declines in the current wage of workers, indicating that high-wage workers are less likely to move directly between employers. As high-wage workers face

fewer acceptable outside offers, this maps to lower job-to-job mobility for workers with higher wages.

Taken together, these facts imply that effective job-arrival rates of relevant (acceptable) jobs are declining in wages. High-wage and stable jobs are associated with fewer outside options, not because workers are less attentive or irrational, but because high-paying jobs are rare and reservation wages are high. While not directly imposed by the model, lower job-to-job mobility among high-wage workers is a natural outcome of the search-theoretic mechanisms (Burdett and Mortensen, 1998; McCall, 1970).

A.5.5 Employment stability

The modeling of employment stability needs careful discussion to determine whether it is consistent with the data. In this paper, heterogeneity in employment stability arises from workers drawing jobs with different separation rates. Separation events include both involuntary separations (layoffs) and voluntary separations (quits), and the model abstracts from explicitly distinguishing between these two margins. This modeling choice is consistent with approaches in Pinheiro and Visschers (2015), Jarosch (2023), and Larkin (2019).

To assess whether differences in employment stability are driven by higher voluntary quitting among less stable workers, I use data from the Current Population Survey and compute employment-to-unemployment (EU) transitions and employment-to-nonparticipation (EN) transitions, in addition to the total separation rate. I perform this analysis both for the full sample of workers and for workers with tenure of two years or less, who represent a group with relatively low employment stability.

I interpret EU transitions as primarily capturing involuntary separations, while EN transitions provide an imperfect but informative proxy for voluntary separations. If low employment stability were mainly driven by higher voluntary quitting, one would expect substantially higher EN transition rates among low-tenure workers relative to the overall workforce.

The results are summarized in Table A.2. Importantly, I find that the composition of the total separation rate (the share of EU and EN transition rates) varies little between the full sample and low-tenure workers, implying that the composition of separations (quits vs. layoffs) is similar across worker groups with different labor market histories. This pattern suggests that lower job stability in bad jobs is not primarily driven by higher voluntary quitting behavior. The evidence is consistent with interpreting job stability as a job- or firm-specific characteristic reflecting exposure to labor market risk, rather than as the outcome of workers' endogenous quitting decisions.

Table A.2: Separation, EU, and EN rate (in %)

Age	All			Low tenure ($\leq$ 2 years)		
	Total	EU	EN	Total	EU	EN
20-29	4.28	1.95	2.33	5.19	2.35	2.84
30-39	2.16	0.96	1.19	3.51	1.78	1.73
40-49	2.16	0.93	1.24	4.78	2.39	2.39
50-59	2.95	1.16	1.79	5.18	2.16	3.02

Notes: This table shows the average monthly separation rates in four age groups (20-29, 30-39, 40-49, 50-59) for all workers in columns 2-4 and for workers with less than or equal to two years of tenure (low employment stability). The total separation rate denotes separation from employment, combining transitions from employment to unemployment (EU) and from employment to nonparticipation (EN).

A.5.6 Approximation of average lifetime earnings

To evaluate the accuracy of the approximation of average lifetime earnings by the labor earnings at the end of the working phase, I proceed as follows. First, for each worker, I take the 35 years with the highest labor earnings of the worker and compute the average lifetime earnings in these years. Then, I derive the approximations to these average lifetime earnings by regressing them on cubic polynomials of the labor earnings before retirement of each worker. This yields an R^2 statistic of approximately 0.9. The labor earnings in the last period thus approximate the average lifetime earnings well, explaining 90% of the variance in the average lifetime earnings. The fit is even tighter in the lower half of the lifetime earnings distribution: for households below the median, this measure explains 99% of the variation in lifetime earnings.

For computational tractability, the current model does not include a separate state variable that explicitly tracks workers' cumulative lifetime earnings. Instead, human capital serves as a sufficient statistic summarizing workers' earnings histories. In particular, the relationship between lifetime earnings and human capital in the model is very strong. The correlation coefficient between lifetime earnings and human capital is 0.952, explaining 87.8% of the variation in workers' lifetime earnings. This tight relationship even holds in the upper tail of the lifetime earnings distribution, with a correlation coefficient of 0.948 and an R^2 of 0.877. Moreover, I also compute the Spearman rank correlation coefficient that is informative about the preservation of workers' position in the lifetime earnings distribution, which amounts to 0.94. Hence, workers' relative ranking by lifetime income is closely mirrored by their ranking in end-of-career human capital.

This strong relationship reflects the fact that human capital effectively serves as a proxy for lifetime earnings: human capital is proportional to labor earnings in each period and accumulates only during employment. As a result, it embeds information about both

earnings levels and employment histories. Importantly, the model does not rely solely on recent earnings for determining pension benefits when workers enter retirement, but instead uses the accumulated information contained in end-of-career human capital and wages, which closely approximate lifetime earnings. Conceptually, treating human capital as a proxy for lifetime earnings is a close approximation to augmenting the model with an explicit lifetime-earnings state variable, which would likewise require coarse discretization for computational tractability.

A.6 Wages and human capital

A central challenge in evaluating the model's income process is that observed earnings data do not allow wages and human capital to be separately identified. To address this issue, I examine the implications of the model's earnings dynamics and assess whether they are consistent with well-established empirical evidence.

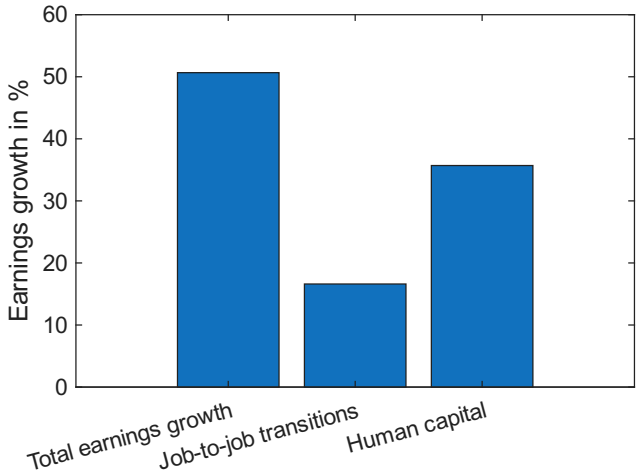
In what follows, I show that the model matches a range of untargeted empirical moments related to earnings changes associated with job mobility, including job-to-job wage gains (Topel and Ward, 1992) and earnings losses upon displacement (Jacobson et al., 1993). I also validate the model's earnings process by showing that it matches key reduced-form empirical moments on earnings dynamics documented in Blundell et al. (2008) and Guvenen et al. (2021).

In the model, human capital represents a transferable component of productivity that carries over across jobs, while wages capture the non-transferable, job-specific component of the current job. The relative magnitude and interaction of the two components are important for how earnings evolve when workers transition between jobs or experience job separations. The empirical moments related to earnings changes upon job mobility are therefore informative about the relative importance of human capital and wages.

First, in the model, workers experience approximately 50% wage growth during the first ten years after labor market entry, of which about 34% is attributable to wage gains from job-to-job transitions. These gains reflect growth in the non-transferable, job-specific component of earnings associated with moving up the job ladder, while the remaining earnings growth is driven by the accumulation of transferable human capital. The decomposition is displayed in Figure A.12. This finding is consistent with Topel and Ward (1992), who document that roughly one-third of early-career wage growth arises from wage increases associated with job changes.

A further piece of evidence supporting the calibration of the comparative importance of both income components is the model's ability to replicate the magnitude and persistence of earnings losses following job displacement. Empirically, both earnings and consumption losses after displacement are large and long-lasting. Following the displacement literature

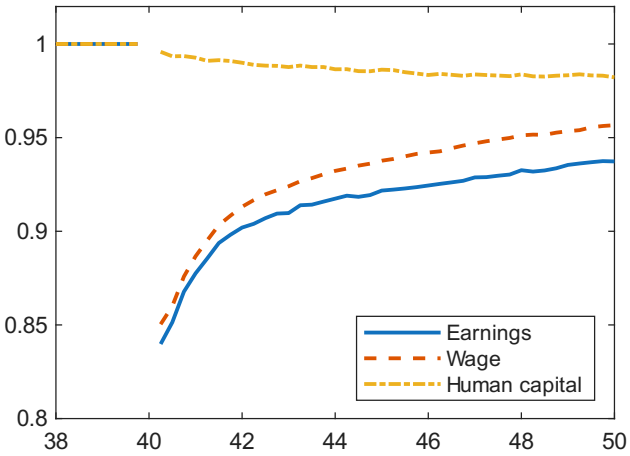
Figure A.12: Earnings growth in early career



Notes: This figure shows the average earnings growth in the first 10 years after labor market entry. Total earnings growth arises from wage gains at job-to-job transitions and from human capital accumulation.

(Jacobson et al., 1993), I analyze the effects of job loss for workers who are identical at age 40. Specifically, I compare a cross section of these workers to an otherwise identical group that experiences displacement at age 40. The results are summarized in Figure A.13.

Figure A.13: Earnings loss upon displacement



Notes: This figure shows the evolution of earnings, wage, and human capital 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.

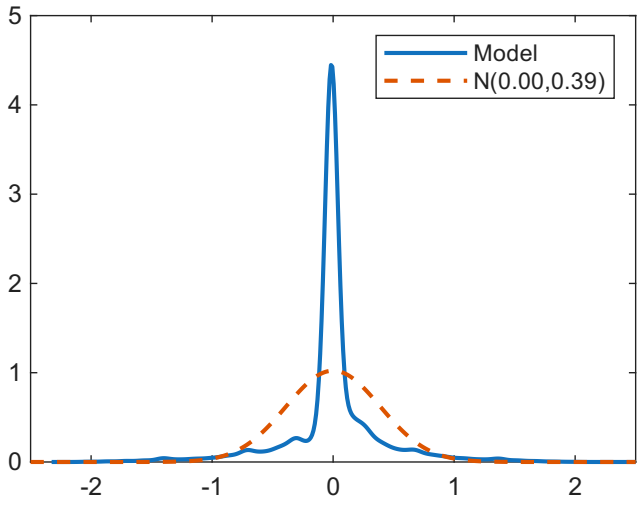
Consistent with empirical evidence documenting substantial and persistent earnings losses after job displacement (Jacobson et al., 1993; Davis and von Wachter, 2011), the model predicts an immediate earnings decline of about 16%. This loss is highly persistent and does not fully recover even ten years after displacement. Importantly, Figure A.13 decomposes this earnings response into wage and human capital components. The majority

of the initial earnings loss is driven by a decline in wages, consistent with the findings of Stevens (1997) and Jung and Kuhn (2019). Displaced workers fall down the job ladder and face difficulties in quickly transitioning into high-paying jobs. This pattern reflects the loss of the job-specific earnings component upon displacement. The non-transferable wage component accounts for roughly 70% of the earnings loss ten years after displacement, consistent with Gulyas and Pytka (2021), who find that declines in firm-specific wage premia explain about 60–70% of post-displacement earnings losses.

To further assess the model's fit to earnings dynamics, I compare model-implied earnings processes with reduced-form statistical representations commonly estimated in applied work and used to parametrize exogenous earnings risk in consumption–saving models. First, I decompose earnings into permanent and transitory components and estimate the variances of innovation terms following standard approaches in the literature (Low et al., 2010; Blundell et al., 2008; Heathcote et al., 2010). Second, motivated by Guvenen et al. (2021), I examine higher moments of the earnings-growth distribution.

In the first step, I estimate the variance of the permanent component of earnings using simulated annual earnings series from the model. Applying the identification strategy of Blundell et al. (2008), I obtain an estimated permanent-shock variance of 0.025, which lies well within the range of empirical estimates. For example, Blundell et al. (2008) report time-varying variances between 0.01 and 0.03 for the period 1980–1990.

Figure A.14: One-year earnings changes



Notes: This 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.

In a second step, I evaluate the model in light of Guvenen et al. (2021), who document that earnings growth is non-normal, exhibiting pronounced negative skewness and excess kurtosis. As shown by Hubmer (2018), these features can be accounted for by life-cycle

job-ladder models with human capital accumulation. My model reproduces these patterns. Figure A.14 displays the distribution of one-year earnings growth rates in the model, along with a normal distribution matched in variance. Earnings changes are strongly left-skewed and tightly concentrated around zero.

Taken together, these findings provide important validation of the income process assumed in the model. Although wages and human capital are difficult to identify separately in the data, the model reproduces untargeted empirical moments that are highly informative about the underlying earnings dynamics, including early-career earnings growth from job mobility and the magnitude, persistence, and composition of earnings losses following displacement. This evidence suggests that the model's income process captures the key features of observed earnings dynamics given the underlying labor market structure.

A.7 Persistence of initial employment stability

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

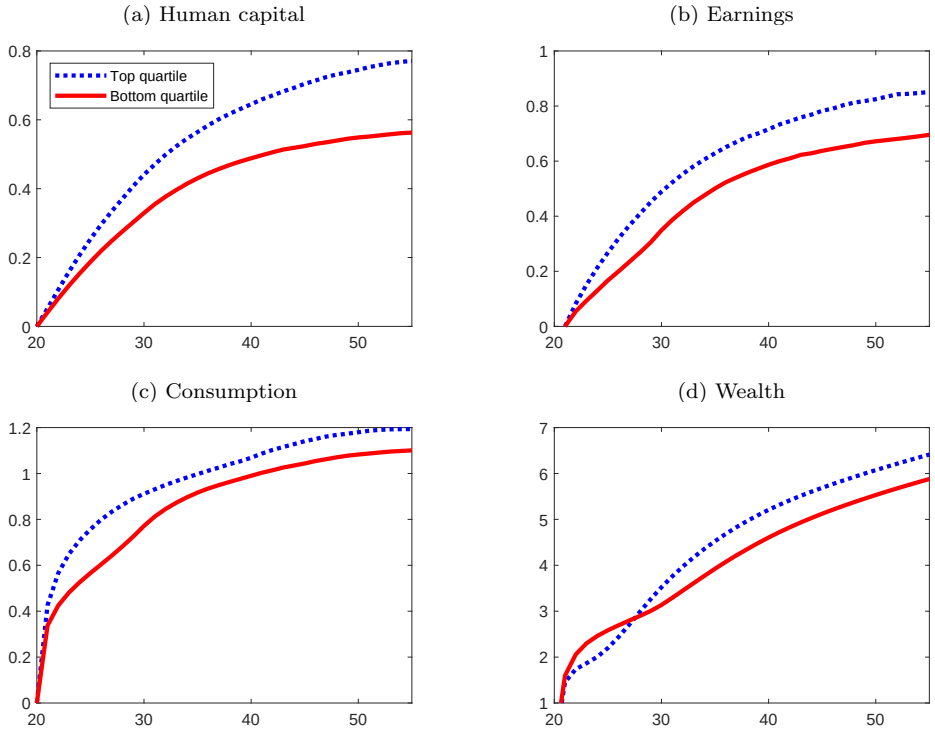
Table A.3: Persistence rate of employment stability at age 30.

	Data	Model
Stable	0.61	0.67
Unstable	0.69	0.70

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

I restrict the sample to household heads between ages 25 and 50 and keep household heads who are not permanently out of the labor force. Based on this sample, I compute the cumulative sum of nonemployment duration at each age of household heads starting at age 25. This yields a distribution of cumulative nonemployment duration at each age. Then, using the median of the cumulative nonemployment duration at age 30, I assign individuals to stable and unstable groups at age 30. Similarly, I use the median of the cumulative nonemployment duration at age 50 to assign individuals to stable and unstable groups. A worker can be either in the stable or unstable group when young (age 30),

Figure A.15: Average life-cycle profiles by employment stability.



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

and either in the stable or unstable group when old (age 50). If employment stability is persistent, workers in the stable (unstable) group at age 30 will remain in the stable (unstable) group at age 50 with a probability $> 50\%$. That is, for a worker in the stable group when young, it is more likely that the worker will remain in the stable group in the future.

The results in Table A.3 indicate that employment stability at age 30 is persistent in the data. A worker with an unstable job at age 30 remains in the unstable group at age 50 with a probability of 70%. The model also produces a high persistence rate of employment stability, while it produces a higher persistence rate for the stable group.

Figure A.15 shows the life-cycle profiles of human capital, earnings, consumption, and wealth-to-income ratios as in Figure 4. The profiles are again normalized by the initial levels and expressed in log deviations. In contrast to Figure 4, however, the top and bottom quartile of the employment-stability distribution are determined based on the first ten years of labor market history rather than the whole working life. Due to persistence of employment stability, workers with unstable jobs at the beginning of the career have, on average, unstable jobs over the whole life cycle. As a consequence, I observe large gaps in

the labor market outcomes as shown in Figure A.15 where the two groups of workers were determined based on the employment stability in the first ten years since labor market entry.

A.8 Policy parameters and welfare results

Table A.4 displays the policy parameters associated with each progressivity level γ for selected values of γ. The parameters in the baseline economy are displayed in the first line where $\gamma = 0.44$ and $\phi = 0.7$ with a payroll tax rate of 12.4%. Comparing the welfare change across all parameter specifications, the optimal policy is a pension system with no Social Security cap and the progressivity parameter set to $\gamma = 0.87$. The corresponding progressivity level is 42.19% and this parameter combination leads to a welfare gain of 0.94% in terms of consumption-equivalent variation for a labor market entrant. The payroll tax rate that achieves budget balance for the pension system decreases to 11.33% and the corresponding level of parameter ϕ increases to 1.08.

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

Model	Parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
Baseline	0.44	0.70	7.48	27.61	12.4	-
No cap	0	0.38	–	18.97	10.97	−0.77
	0.2	0.50	–	19.88	11.03	−0.18
	0.4	0.65	–	25.25	11.12	0.26
	0.6	0.82	–	32.17	11.20	0.61
	0.8	1.01	–	39.64	11.30	0.88
	0.87	1.08	–	42.19	11.33	0.94
	0.9	1.12	–	43.29	11.35	0.92
	1	1.22	–	46.79	11.40	0.63
	1.2	1.38	–	54.10	11.52	−1.22

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

A.9 Distortions on human capital

An important determinant of the optimal pension system is the elasticity of human capital investment with respect to the income tax schedule. This elasticity governs the magnitude of policy-induced distortions in skill accumulation and lifetime earnings. In the model, the implied elasticity of human capital investment with respect to the income tax parameter ranges from -0.28 for the full sample to -0.35 for prime-age workers. These values are consistent with estimates in the literature; for example, Hendricks (2004) reports an elasticity of -0.295.

I also analyze the elasticity of human capital investment with respect to the progressivity of the pension system. Although this elasticity is difficult to estimate empirically, it is informative to compare it to the elasticity with respect to income taxation. Specifically, I vary the parameter γ to assess how households adjust their human capital accumulation decisions in response to changes in pension progressivity. The parameter γ governs the curvature of the pension benefit formula,

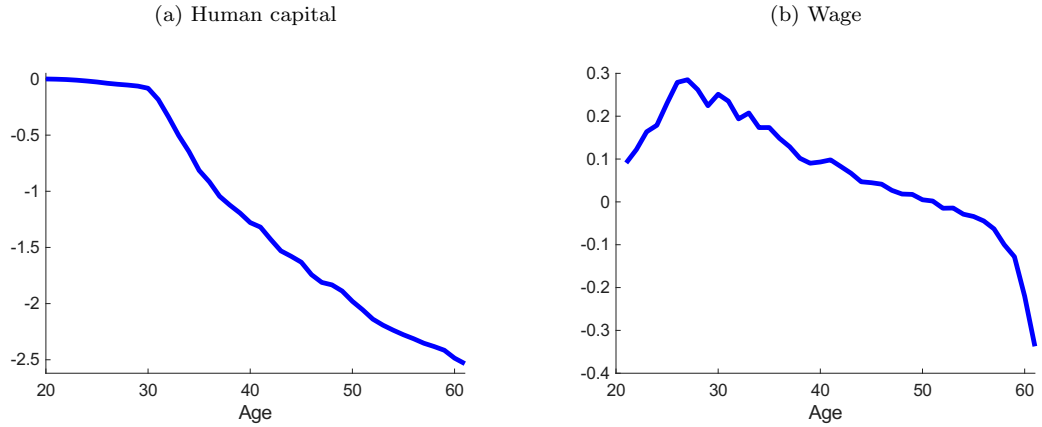
$$\omega(\bar{y}) = \phi \cdot (\bar{y})^{1-\gamma}$$

where ω denotes pension benefits and γ determines the degree of progressivity of the pension system. I estimate an elasticity of human capital investment with respect to γ of -0.39 for the full sample and -0.48 for prime-age workers. These estimates are close to, though somewhat larger in absolute value than, the corresponding elasticity with respect to the income tax parameter.

In the following, I compare the baseline economy with the economy under the optimal pension system to examine how the increase in pension progressivity affects human capital investment and job ladder decisions. Figure A.16 illustrates the resulting changes in the life-cycle profiles of human capital and wages relative to the baseline economy. If wage growth were driven solely by passive experience accumulation, policy-induced distortions to human capital would be substantially smaller. Instead, the model generates sizable and systematic responses of human capital accumulation to changes in progressivity, indicating that forward-looking investment decisions—rather than mechanical experience effects alone—play an important role in shaping earnings dynamics over the life cycle.

Figure A.16a shows that by the end of the prime-age working period, workers in the optimal pension economy accumulate approximately 2.5% less human capital than in the baseline economy, indicating a sizable distortion of human capital investment induced by higher pension progressivity. Increased progressivity weakens incentives to invest in human capital, as future labor income becomes more strongly insured. Changes in progressivity also affect workers' job ladder decisions. As shown in Figure A.16b, workers in the optimal pension economy select high-wage jobs relative to the baseline economy up to about age

Figure A.16: Changes under optimal policy relative to the baseline economy



Notes: This figure shows the percent deviation in average life-cycle profiles of human capital (left panel) and wage (right panel) in the economy under the optimal pension system compared to the baseline economy.

40. After that point, however, average wages fall below those in the baseline, reflecting altered job ladder dynamics over the life cycle. This pattern indicates that the optimal pension system distorts not only human capital accumulation but also job choice and mobility decisions.

These distortions are further reflected in job mobility patterns. On average, job-to-job transition rates increase by 0.06 percentage points for workers aged 20 to 30 under the optimal pension system. With reduced incentives to accumulate human capital, young workers place greater weight on immediate wage gains and switch jobs more frequently, taking advantage of higher insurance against labor market risk.

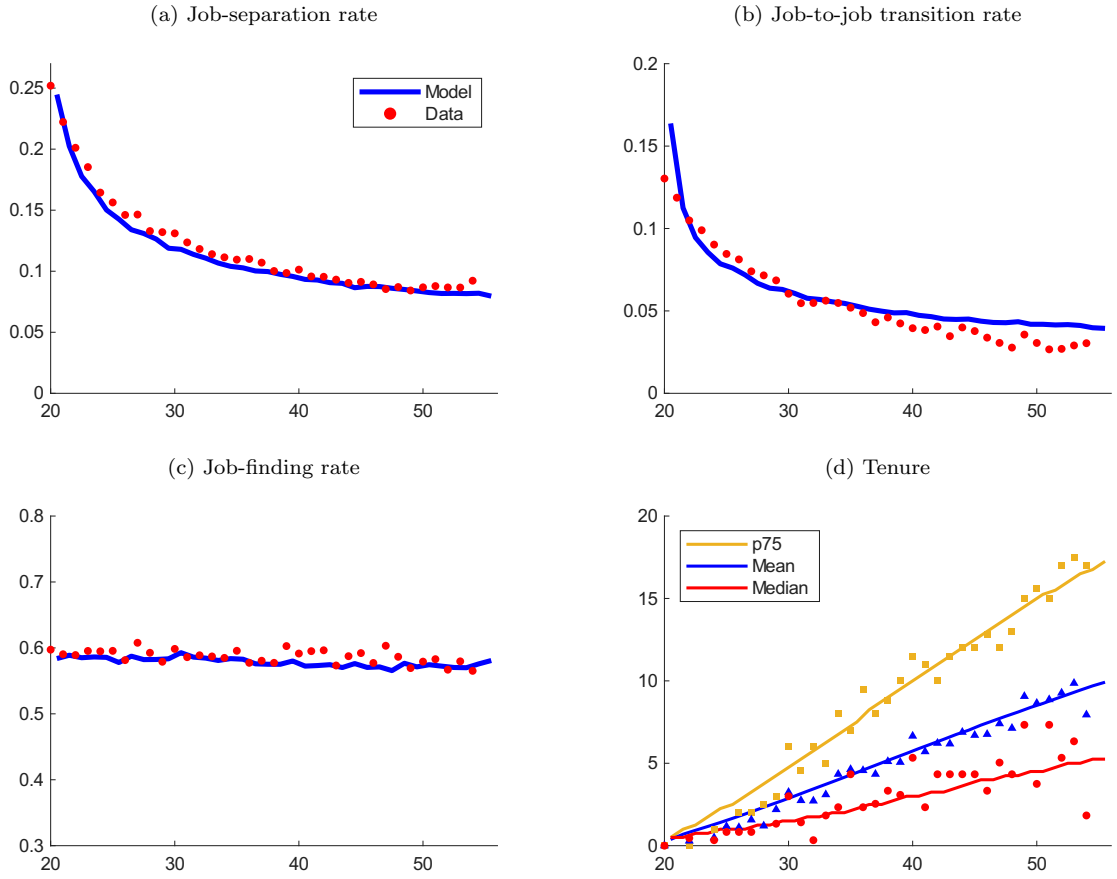
A.10 Optimal pension system for non-college and college-educated workers

In the optimal policy analysis in the main text (Figure 5), the welfare conclusions are based on a calibration that targets moments for the overall population. An important question is how the optimal pension results change once ex-ante heterogeneity across education groups is taken into account. In particular, it is a well-documented empirical fact that non-college workers experience substantially higher job-separation rates than college-educated workers. To examine the implications of this heterogeneity, I recalibrate the model separately for college and non-college workers and repeat the welfare analysis

for each group.[8]

In the following, I present the calibration results from a version of the model that is estimated separately for non-college and college-educated workers. The empirical moments are constructed using the same methodology as in the baseline calibration, but conditioning on educational attainment.

Figure A.17: Non-college workers: labor market transitions and tenure.

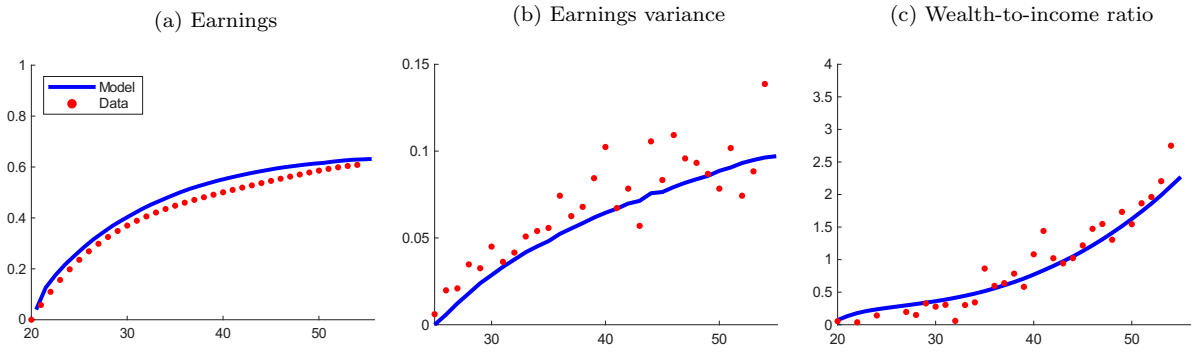


Notes: This figure displays quarterly life-cycle labor market transition rates and tenure in years by age for non-college workers. Panel a shows the average separation rate, Panel b the job-to-job transition rate, and Panel c the job-finding rate. Panel d presents the tenure distribution, including the mean, median, and 75th percentile. In the model, tenure is defined as the duration of continuous employment in the same job, uninterrupted by separations or job-to-job transitions. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical transition rates are constructed using data from the Current Population Survey, and tenure profiles are obtained from the Panel Study of Income Dynamics.

Figure A.17 reports the calibration outcomes for non-college workers. The average separation rate is substantially higher for workers without a college degree than in the full

[8]An alternative approach would be to model both groups jointly in a unified framework when analyzing the optimal pension system. However, such a setup would shift the focus toward redistribution across ex-ante heterogeneous households with different income and wealth levels as college-educated workers have higher average income and wealth. In contrast, the purpose of this paper is to isolate redistribution across workers who differ primarily in their labor market risk and experience different employment shocks over the life cycle.

Figure A.18: Non-college workers: earnings and wealth.



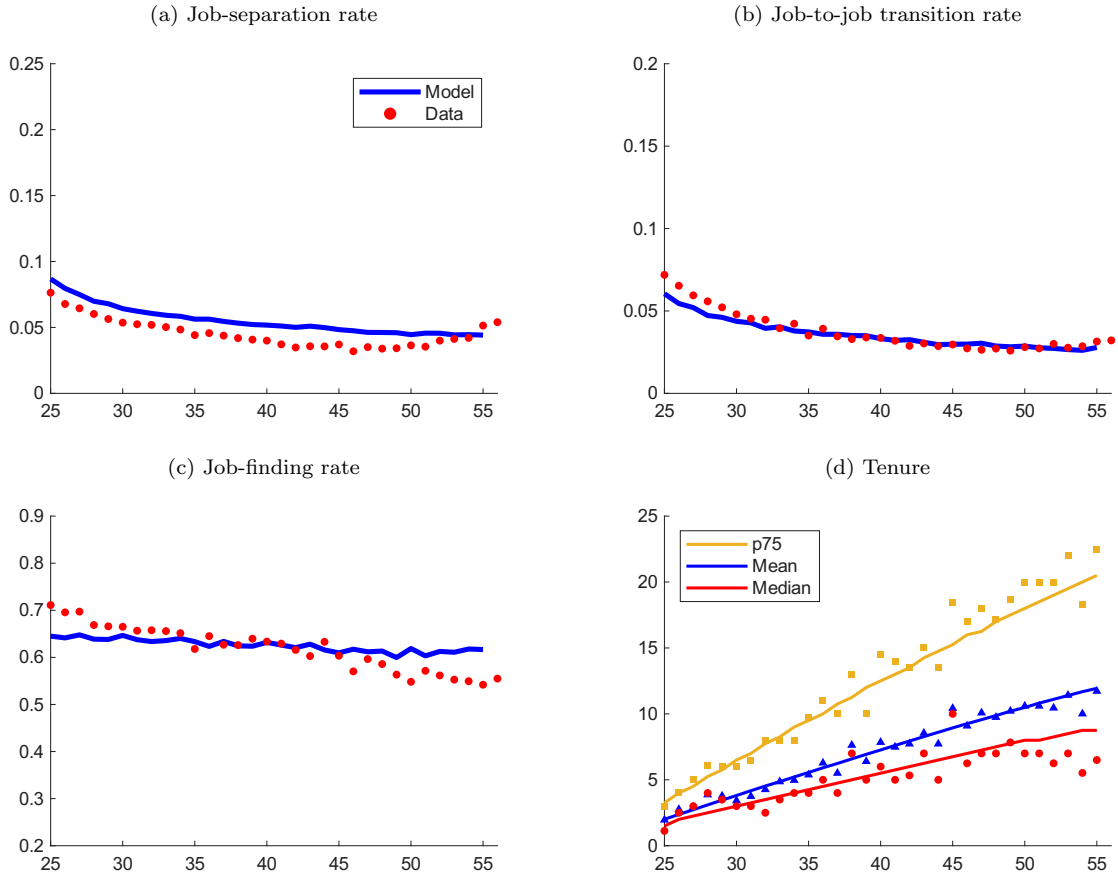
Notes: This figure presents quarterly life-cycle profiles of earnings and the wealth-to-income ratio by age for non-college workers. Panel a displays average log labor earnings, Panel b shows the variance of log labor earnings, and Panel c reports the wealth-to-income ratio. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical profiles of earnings and the wealth-to-income ratio are constructed using data from the Survey of Consumer Finances, while the variance of log earnings is obtained from the Panel Study of Income Dynamics.

sample. Moreover, their job-finding rate is lower throughout the life cycle and exhibits a relatively flat profile. The tenure distribution also widens less over time, indicating that it is more difficult for workers to find stable jobs. The life-cycle profiles for earnings, earnings variance, and the wealth-to-income ratio are shown in Figure A.18. Across all three dimensions, non-college workers exhibit weaker growth over the life cycle relative to the overall population, reflecting lower earnings growth and less wealth accumulation.

For college-educated workers, empirical life-cycle profiles are constructed starting at age 25. Compared to non-college workers, college-educated workers exhibit substantially lower average separation rates and higher job-finding rates throughout the life cycle, as shown by Figure A.19. As a result, the average tenure is higher and their tenure distribution widens more strongly over time, reflecting faster transitions into stable employment. Although college-educated workers tend to secure stable jobs more quickly, there remains considerable heterogeneity in their career trajectories. This is evident in the large dispersion in tenure outcomes across workers. Moreover, despite the lower average separation rate, the empirical profiles indicate that a non-negligible share of college-educated workers still experience unstable employment over the life cycle.

College-educated workers experience substantially higher earnings growth over the life cycle. As shown in Figure A.20, earnings growth between ages 25 and 55 is nearly twice as large as that of non-college workers. Earnings dispersion also increases more strongly over the life cycle among college-educated workers, indicating greater heterogeneity in income trajectories. In addition, college-educated workers accumulate substantially more wealth over their working lives, consistent with the literature on wealth inequality (Kuhn et al., 2025). Their wealth-to-income ratio at the end of prime-age working life is approximately 60% higher than that of non-college workers, reflecting both higher lifetime earnings and

Figure A.19: College-educated workers: labor market transitions and tenure.

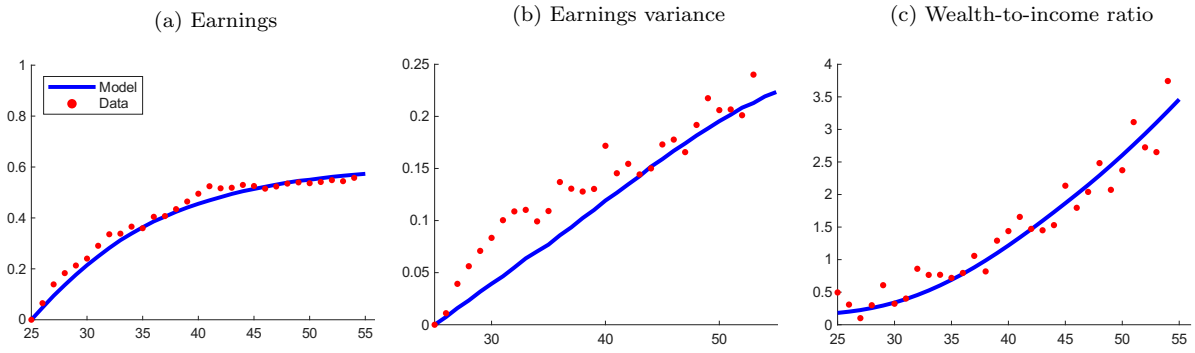


Notes: This figure displays quarterly life-cycle labor market transition rates and tenure in years by age for college-educated workers. Panel a shows the average separation rate, Panel b the job-to-job transition rate, and Panel c the job-finding rate. Panel d presents the tenure distribution, including the mean, median, and 75th percentile. In the model, tenure is defined as the duration of continuous employment in the same job, uninterrupted by separations or job-to-job transitions. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical transition rates are constructed using data from the Current Population Survey, and tenure profiles are obtained from the Panel Study of Income Dynamics.

greater motive for life-cycle savings.

Figure A.21 presents the welfare results under different pension systems for college-educated and non-college workers. The optimal policy parameters in both cases are remarkably similar to those obtained when calibrating the model to the full population. In particular, it remains optimal to eliminate the Social Security cap and to implement a pension benefit schedule that becomes flat in pre-retirement earnings. Interestingly, the welfare gains from the optimal pension reform are substantially larger for college-educated workers. Although college-educated workers face lower average job-separation rates, heterogeneity in labor market risk within this group remains pronounced and the resulting differences in career trajectories are even more dispersed, implying that the optimal pension system continues to provide meaningful insurance and redistribution. For workers without a college degree, the optimal policy parameters are likewise similar,

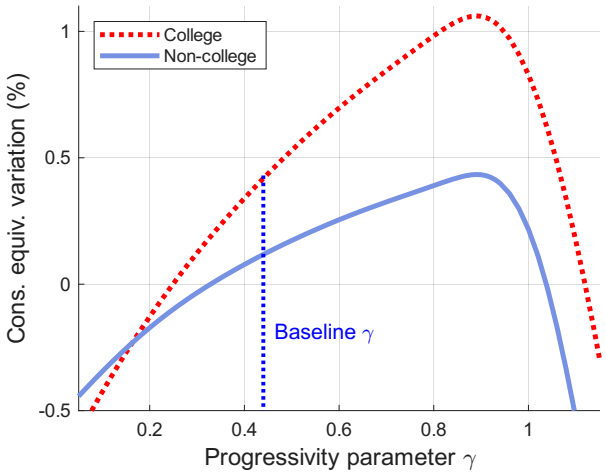
Figure A.20: College-educated workers: earnings and wealth.



Notes: This figure presents quarterly life-cycle profiles of earnings and the wealth-to-income ratio by age for college-educated workers. Panel a displays average log labor earnings, Panel b shows the variance of log labor earnings, and Panel c reports the wealth-to-income ratio. Dots represent the empirical profiles, while solid lines depict the corresponding model-generated profiles. Empirical profiles of earnings and the wealth-to-income ratio are constructed using data from the Survey of Consumer Finances, while the variance of log earnings is obtained from the Panel Study of Income Dynamics.

but the associated welfare gains are smaller. This reflects their lower average earnings, which makes the optimal pension system more costly to implement, as increasing pension progressivity requires higher payroll taxes relative to income.

Figure A.21: Welfare and changes in progressivity level for college and non-college workers.



Notes: This figure displays changes in welfare, measured in consumption-equivalent variation, resulting from variations in the progressivity parameter γ in the no-cap scenario. The dotted line represents welfare changes for college-educated workers, while the solid line shows the corresponding welfare changes for non-college workers.

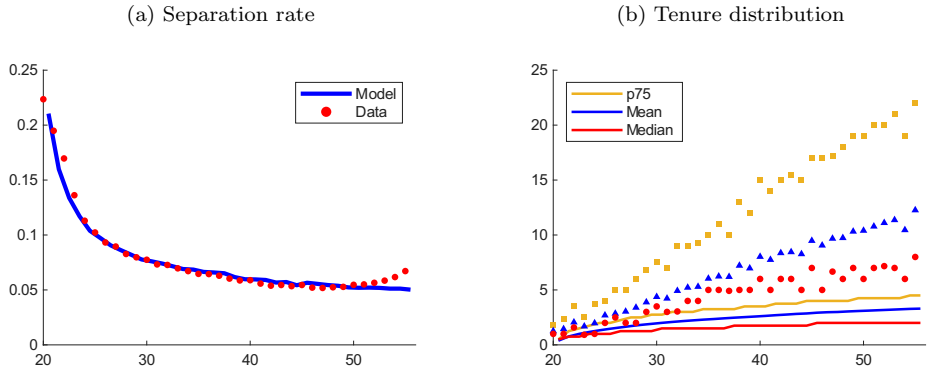
A.11 Alternative models without age-specific heterogeneity in separation rate

In this section, I assess whether models without rich cross-sectional heterogeneity in employment stability can be empirically rejected. To do so, I consider two restricted versions of the model: (i) a model in which separation rates vary only by age, and (ii) a model in which separation rates vary by both age and education.

The results show that removing cross-sectional heterogeneity in separation risk—while allowing for age-dependent or age- and education-dependent separation rates—fails to generate the observed dispersion in workers' employment histories. In particular, these restricted models cannot replicate the substantial within-group heterogeneity in career stability documented in the data. Since this dispersion in employment trajectories is the central empirical moment that motivates the paper and describes a salient feature of empirical labor market mobility, the models without cross-sectional heterogeneity in separation risk are unable to account for the mechanisms underlying the optimal pension results.

In the first restricted specification, separation rates vary only by age. I impose the average age-dependent separation rates derived from the Current Population Survey (CPS), thereby shutting down cross-sectional heterogeneity in job stability. As in the baseline model, the parameters are calibrated using the simulated method of moments, targeting life-cycle profiles of labor market transition rates, mean earnings, earnings dispersion, and the wealth-to-income ratio.

Figure A.22: Baseline model without age-specific heterogeneity in employment stability



Notes: These figures show the separation rate (Panel a) and the tenure distribution (Panel b) over the life cycle in a model without age-specific heterogeneity in employment stability.

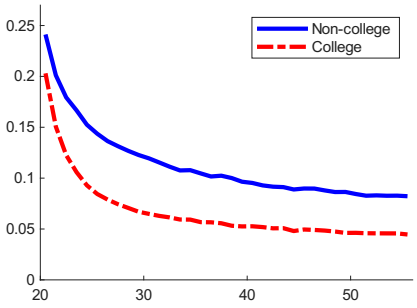
The resulting separation profile is shown in Figure A.22a. By construction, the model matches the average separation rate over the life cycle very well. However, it fails to reproduce the heterogeneity in career paths observed in the data. As shown in Figure A.22b, the tenure distribution remains nearly constant after age 35 and does not widen over the life cycle. This stands in stark contrast to the data, where workers' labor market histories diverge substantially over time, leading to increasing inequality in tenure outcomes over the life cycle. In particular, the model fails to replicate the empirically salient feature of a large share of stable lifetime jobs. At age 55, the 75th percentile of tenure reaches only 4 years in the model, compared to 23 years in the data.

In the second version of the model, education enters as a worker-fixed effect. Non-college workers face substantially higher average separation rates over the life cycle than college-educated workers. I construct empirical life-cycle profiles separately for each education group using data from the Current Population Survey, the Survey of Consumer Finances, and the Panel Study of Income Dynamics. The model is then calibrated separately for college-educated and non-college workers using the simulated method of moments, targeting the group-specific life-cycle moments for labor market transitions, earnings (levels and dispersion), and wealth.

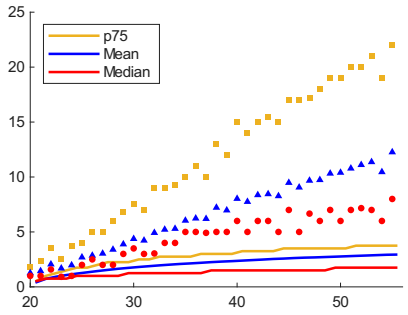
Figure A.23a shows the average separation rates for non-college and college-educated workers. Within each education group, all workers are assumed to face the same age-specific separation rate. As a result, the implied tenure distribution over the life cycle remains very similar to that in the model without age-specific heterogeneity in employment stability (Figure A.22b). In particular, the tenure distribution does not widen over the life cycle and remains far more compressed than in the data. The model-implied tenure distribution deviates markedly from its empirical counterpart across the life cycle. The persistent compression of tenure outcomes indicates that the specification with only age- and education-specific separation rates cannot account for the observed divergence in

Figure A.23: Model with only worker fixed effects

(a) Separation rate



(b) Tenure distribution



Notes: These figures display the separation rate by education (Panel a) and the life-cycle tenure distribution (Panel b) in a model with worker fixed effects but without age-specific heterogeneity in employment stability within the education groups. The model allows for permanent worker heterogeneity across two education groups—college and non-college—but assumes identical age-specific separation rate within each group.

employment histories, which leads me to conclude that this model is rejected by the data.

Table A.5: Gini coefficient of nonemployment duration

	Data	Model		
		Baseline	Baseline, no heterogeneity	Worker fixed effect, no heterogeneity
Gini coefficient	0.56	0.50	0.28	0.30

Notes: This table shows the Gini coefficient of accumulated nonemployment duration at the end of prime-age working life. The data is computed from the Panel Study of Income Dynamics (PSID). The second column corresponds to the model with heterogeneity in employment stability. The third column refers to the model without age-specific heterogeneity in employment stability, while the last column represents the model with worker fixed effects and no age-specific heterogeneity in employment stability.

As a further diagnostic statistic, I analyze the inequality measures of workers' labor market histories. Table A.5 reports the Gini coefficient of accumulated nonemployment duration at the end of prime-age working life, with the empirical estimate derived from the PSID data. The baseline model generates a Gini coefficient that is somewhat lower than, but close to, the value observed in the data. By contrast, both the model without age-specific heterogeneity in separation rates and the model with education as a worker fixed effect produce substantially lower Gini coefficients, indicating much less inequality in employment histories. This further highlights that rich cross-sectional heterogeneity in separation rates is essential for replicating the dispersion in labor market histories observed in the data.

A.12 Retirement age distortion by human capital

Table A.6 reports the average retirement age in the baseline economy and under the optimal pension system. Workers with human capital below the median retire earlier relative to the baseline, whereas those with above-median human capital choose to remain in the labor force longer.

Table A.6: The mean retirement age in the baseline economy and under the ex-ante optimal pension system.

Mean retirement age	Baseline	Optimal pension system
All	64.34	64.13
Human capital below median	64.04	63.54
Human capital above median	64.84	65.17

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

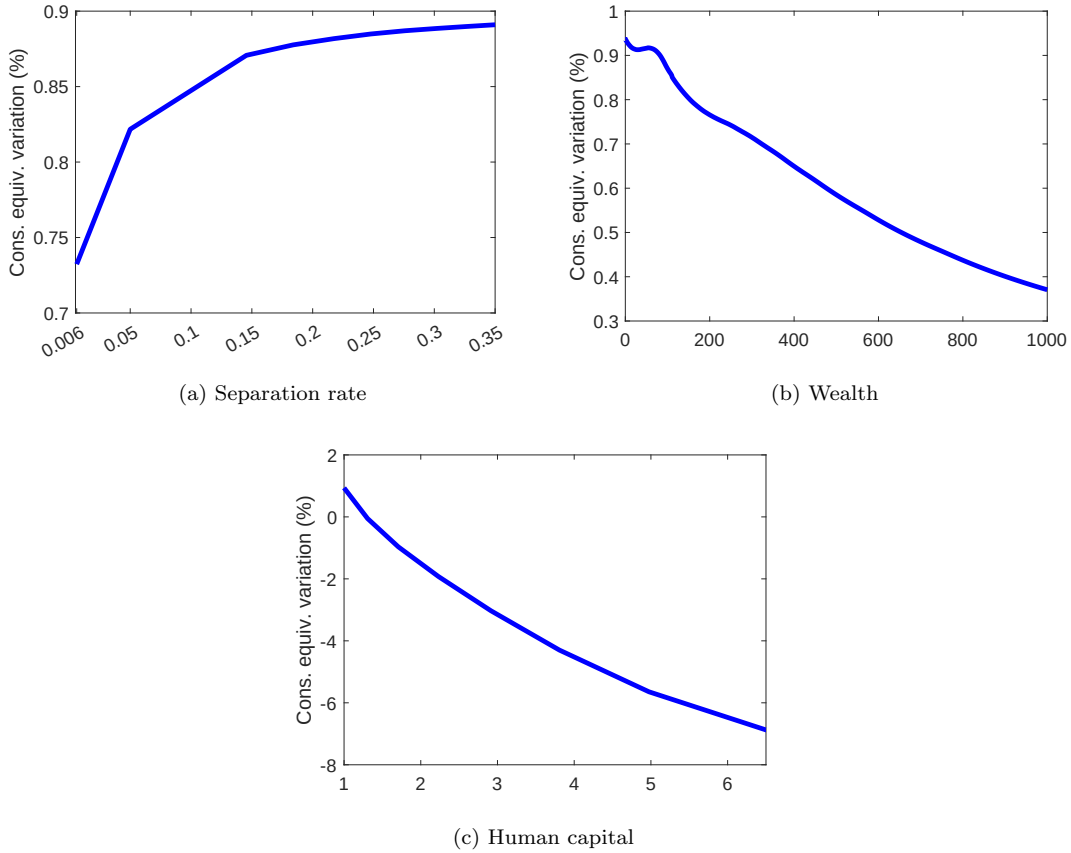
A.13 Initial conditions and welfare effects of pension systems

To analyze the importance of the initial state of workers at labor market entry in studying the optimal progressive pension system, I derive the welfare changes of the ex-ante optimal pension policy as functions of separation rate, wealth, and human capital. More precisely, while holding all other state variables constant at the initial state of a worker at labor market entry, I assess how variations in separation rate, wealth, and human capital each affect the welfare gain from the optimal pension policy. The results are displayed in Figure A.24.

Figure A.24a shows that the welfare change increases in job-separation rate, that is, a worker with an initially unstable job gains more from an increase in pension progressivity. The redistributive effect of the optimal policy provides insurance against future poor labor market outcomes and leads to a welfare gain for a worker with an unstable job. For a worker with a very stable job, however, prospective earnings and human capital growth are larger. The expected gain from higher pension progressivity therefore strictly decreases in employment stability.

Figure A.24b presents the welfare change as a function of the wealth level. Interestingly, the function is non-monotonic: the welfare change exhibits a hump close to zero wealth, but then, decreases steadily in the asset level. Starting from zero wealth, higher wealth allows for a better ability of consumption smoothing in the presence of the payroll tax rate. Because young workers have relatively low labor earnings and they save a large proportion of their earnings for precautionary reasons, an increase in wealth leads to

Figure A.24: Welfare change as a function of the initial state variables of separation rate, wealth, and human capital.



(a) Separation rate

(b) Wealth

(c) Human capital

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

Finally, Figure A.24c shows the welfare change as a function of human capital. Starting from the initial level, a larger stock of human capital implies a decrease in welfare. This points out that workers with low human capital gain from an increase in pension progressivity, while workers with high human capital lose out. An increase in pension progressivity decreases the welfare gain of workers with high human capital in two respects. Firstly, higher pension progressivity reduces pension wealth and therefore the lifetime resources of workers with high human capital. Secondly, despite the decrease in the payroll tax rate due to the abolishment of the Social Security cap, contributions to the pension system increase relative to benefits. As a consequence, the increase in pension progressivity decreases welfare.

A.14 Model extensions

A.14.1 Bequest motive, mortality risk, and early application penalty

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

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

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

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

Here, d_{j_r+1} denotes the age-specific probability of death in the next period. The death probabilities are taken from the decennial life tables (1979-1981) published by the National Center for Health Statistics.[9]

For early claiming penalties in the Social Security system, I apply the reduction rates summarized in Table A.7.

Table A.7: Social Security benefit reduction by claiming age.

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

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

Table A.8 repeats the optimal pension system analysis from Section 6 for the model with mortality risk and a bequest motive to assess how these extensions affect the design and outcomes of the optimal policy. The welfare gain from implementing the optimal policy slightly increases and the optimal policy parameters in the extended model remain

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

Table A.8: Welfare change in the baseline economy under alternative pension systems.

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.7	7.48	-
No cap	0	0.38	–	−0.07
	0.2	0.54	–	0.40
	0.4	0.68	–	0.72
	0.6	0.84	–	0.96
	0.8	1.02	–	0.85
	1	1.24	–	0.58
	1.2	1.44	–	−1.10

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

almost unchanged.

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

A.14.2 Different levels of income tax progressivity

Next, I examine how changes in the progressivity of labor income taxation influence the design of the optimal policy. In the model, taxes are given by

$$T_y(y) = y - (1 - \nu)y^{1-\tau_y}.$$

As income tax progressivity increases, the welfare gains from introducing greater progressivity in the pension system are reduced. Nevertheless, the qualitative result remains robust: a more progressive pension system is optimal even in the presence of a highly progressive income tax. For instance, in a scenario where the income tax progressivity parameter is set to 0.45, a value higher than the baseline progressivity level of 0.181, the reform still yields welfare gains exceeding 0.3% of lifetime consumption.

Table A.9: Economy with income tax progressivity τ_y set to 0.15.

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.70	7.48	-
No cap	0	0.38	–	−0.73
	0.2	0.50	–	−0.16
	0.4	0.65	–	0.27
	0.6	0.82	–	0.62
	0.8	1.01	–	0.88
	1	1.22	–	−0.14
	1.2	1.39	–	−2.62

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

Table A.10: Economy with income tax progressivity τ_y set to 0.3.

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.70	7.48	-
No cap	0	0.38	–	−0.88
	0.2	0.51	–	−0.25
	0.4	0.65	–	0.24
	0.6	0.82	–	0.57
	0.8	1.00	–	0.81
	1	1.21	–	0.01
	1.2	1.35	–	−0.78

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

Table A.11: Economy with income tax progressivity τ_y set to 0.45.

Model	Parameters			Welfare change (%)
	γ	ϕ	cap	
Baseline	0.44	0.70	7.48	-
No cap	0	0.39	–	−0.89
	0.2	0.51	–	−0.29
	0.4	0.66	–	0.10
	0.6	0.82	–	0.36
	0.8	0.99	–	0.46
	1	1.17	–	0.50
	1.2	1.27	–	0.00

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

A.14.3 Different unemployment insurance systems

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

Table A.12 summarizes the results. A key finding is that the optimal pension policy parameter changes only marginally across alternative unemployment insurance systems. In all cases, it is optimal to remove the Social Security cap and the progressivity parameter γ of the pension benefit formula takes a value between 0.87 and 0.89. This is because a higher replacement rate helps workers smooth consumption during unemployment spells, but it does not compensate for potential losses in human capital. Consequently, the optimal degree of pension progressivity remains largely unaffected by the design of the unemployment insurance system. The induced welfare gains decrease with the replacement rate of the unemployment insurance system because workers with more frequent job interruptions receive higher unemployment benefits and the welfare gains from redistribution through the pension system decline.

Table A.12: Optimal policy in economies with different unemployment insurance systems.

Replacement rate of UI	Optimal γ	Welfare change (%)
0.8	0.89	0.88
0.6	0.88	0.92
0.4	0.87	0.94
0.2	0.87	0.96

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

A.14.4 Human capital depreciation

In the baseline model, human capital does not depreciate when workers are unemployed. To analyze the potential consequences of human capital depreciation on the design of optimal pension systems, I assume that unemployed workers face a probability of 0.12 of human capital depreciation. This probability implies a yearly depreciation rate of approximately 12% which is within the range of human capital depreciation rates in the literature. For example, using the NLSY data, Keane and Wolpin (1997) find yearly depreciation rates of 9.6% and 36.5% for human capital. Jacobson et al. (1993) estimate depreciation rates between 10% and 25%.

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

Table A.13: Optimal pension systems in the economy with human capital depreciation.

Model	Optimal policy parameters			Progressivity (%)	Tax (%)	Welfare change (%)
	γ	ϕ	cap			
No cap	0.85	1.01	–	44.25	11.81	0.84

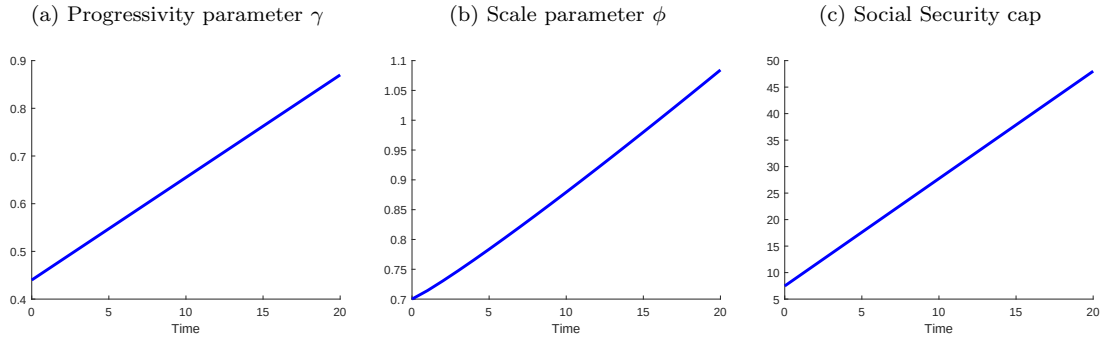
Notes: Column 1 specifies the model. Columns 2-4 show the optimal pension system parameters γ and ϕ, which determine the shape and the level in the pension benefit formula, and the cap. Column 5 presents the degree of progressivity achieved by the combination of policy parameters and Column 6 the tax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

Table A.13 summarizes the optimal policy parameters. Similar to the baseline model, removing the Social Security cap is optimal and the progressivity parameter γ increases

from 0.44 to 0.85 so that the progressivity level increases under the optimal policy. The inclusion of human capital depreciation does not change the previous results qualitatively.

A.15 Transition to the optimal policy

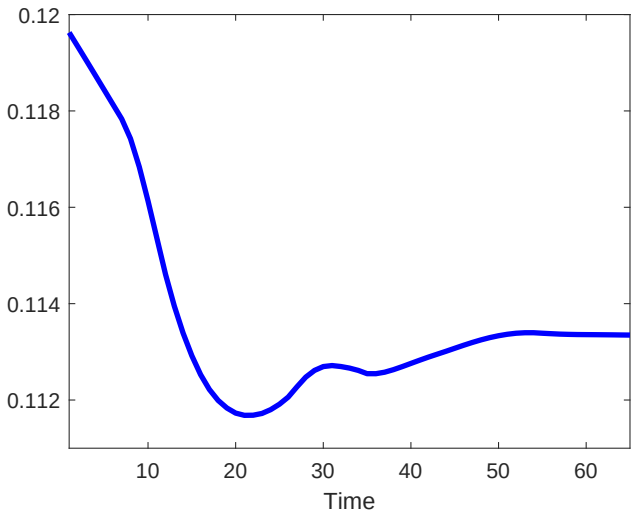
Figure A.25: Pension system parameters over the transition.



Notes: Panel a shows the progressivity parameter γ, Panel b the scale parameter ϕ, and Panel c the Social Security cap along the transition path to the optimal pension system. Time indicates the year after the beginning of the transition.

This section describes how the pension system parameters are gradually changed during the transition until the optimal policy is reached. The progressivity parameter γ, the scale parameter ϕ as well as the Social Security cap are increased starting from the baseline economy at time zero to the optimal policy after 20 years. In each period, the tax rate is adjusted to achieve budget balance.

Figure A.26: Tax rate over the transition.



Notes: This figure shows the tax rate along the transition path. Time indicates the year after the beginning of the transition.

Given these exogenous adjustments to the pension parameters along the transition path, I compute the tax rate that adjusts endogenously over time to ensure budget balance

in every period. Although the pension system parameters are set to their optimal-policy values after 20 years, the economy requires more than 30 additional years to fully converge to the new steady state associated with the optimal pension system.

A.16 General equilibrium effects on interest rates and wages

The baseline model assumes that the interest rate and wages are fixed and therefore abstracts from general equilibrium feedback in the analysis of the optimal pension system. In this section, I extend the model to account for general equilibrium effects on both the interest rate and wages. Modeling a full general equilibrium labor market with wage bargaining is computationally infeasible in the current framework. Instead, I assume a standard Cobb–Douglas aggregate production function given by

$$Y = K^{\alpha} L^{1-\alpha}.$$

Under this specification, the interest rate and the wage are determined by

$$r = \alpha K^{\alpha-1} L^{1-\alpha} - \delta$$
$$\overline{w} = (1-\alpha) K^{\alpha} L^{-\alpha}.$$

Aggregate capital is given by the sum of asset holdings across households, while aggregate labor is defined as total effective labor input, equal to the sum of human capital supplied by employed workers in the economy. The capital share parameter α is set to 0.36, consistent with the average capital and labor shares observed in the data (Krusell and Smith, 1998), and the annual depreciation rate δ is set to 0.08. While holding wage dispersion across jobs fixed, I allow the average wage $\overline{w}$ to move endogenously, shifting the overall level of wages up or down.

Section 6 shows that when the pension system becomes more progressive, households accumulate more wealth while reducing labor supply at the extensive margin. I further find that during prime-age working life, workers remain unemployed for longer durations under higher pension progressivity, leading to a decline in aggregate labor supply. These behavioral responses imply that, in general equilibrium, higher pension progressivity puts downward pressure on the equilibrium interest rate, while increasing equilibrium wages to partially offset the reduction in labor supply. This introduces an additional trade-off: a more progressive pension system raises average wages but lowers the return on savings.

Table A.14 summarizes the welfare analysis under general equilibrium and compares the resulting optimal policy to that in the baseline (partial equilibrium) economy presented in Section 6. The optimal pension system parameters remain remarkably similar across the two environments. In the general equilibrium model, the interest rate declines modestly,

Table A.14: Optimal pension systems with general equilibrium effects.

Model	Optimal policy			Tax (%)	Interest rate (%)	Wage change (%)	Welfare change (%)
	γ	ϕ	cap				
Baseline	0.87	1.08	–	11.33	3.50	–	0.94
With GE	0.85	1.06	–	11.30	3.48	0.17	1.02

Notes: Column 1 specifies the model. Columns 2-4 show the optimal pension system parameters γ and ϕ, which determine the shape and the level in the pension benefit formula, and the Social Security cap. Column 5 presents the tax rate in the economy that achieves budget balance for the pension system. Columns 6 and 7 show the interest rate and the change in wage relative to the baseline economy. The interest rate and the wage level in the baseline economy are both fixed. Column 8 presents the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

while the equilibrium wage increases by 0.17%. The welfare gain is slightly larger under general equilibrium, as the increase in wages more than compensates for the decline in the interest rate. In sum, incorporating general equilibrium effects and the resulting trade-off between lower returns to savings and higher labor income leads to modest changes in the optimal pension system parameters relative to the partial-equilibrium benchmark. Overall these results indicate that the baseline policy conclusions are robust to general equilibrium effects on the capital and the labor market.