Optimal Progressive Pension Systems and Heterogeneous Labor Market Risk*

Leanne Nam†

January 2024

Abstract

Heterogeneity in job stability is a salient feature of the labor market, which is a well-documented fact by empirical research. Job stability is key for welfare of individuals as interrupted work histories and the consequent earnings losses reduce pension entitlements and impair workers' ability to accumulate life-cycle savings. I study how a progressive pension system optimally considers heterogeneity in work histories and quantify welfare gains from implementing the optimal pension system. Pension progressivity provides insurance against the risk of interrupted work histories and mitigates lifetime earnings inequality caused by heterogeneity in job stability, but comes at the cost of distorting human capital investment and retirement decisions. Using a life-cycle model with heterogeneity in job stability, endogenous human capital accumulation, and retirement decision, I find that abolishing the Social Security cap and increasing pension progressivity relative to the current U.S. pension system is optimal. The optimal pension system leads to a welfare gain of 0.75% of lifetime consumption for labor market entrants. Following the observed macroeconomic shift in the job-stability distribution towards higher job stability since the 1990s, the optimal pension system becomes less progressive, but the welfare gain from implementing the optimal pension system remains sizeable.

JEL classification: E24, H21, H55, J64

Keywords: Pension System, Job Stability, Life-Cycle Earnings Risk

*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 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 Bonn, leanne.nam@uni-bonn.de, 53113 Bonn, Germany

1. Introduction

Heterogeneity in job stability is a salient feature of the labor market. While some workers have lifetime jobs, others have unstable work histories, and frequently experience unemployment. Job stability is key for welfare of individuals as interrupted work histories and earnings losses following job displacements lower pension entitlements and impair workers' ability to accumulate life-cycle savings. Heterogeneity in job stability has been already documented by Hall (1982) and is studied by a growing strand of the literature (Jung and Kuhn, 2019; Molloy et al., 2020; Kuhn and Ploj, 2020; Gregory et al., 2021; Jarosch, 2023; Ahn et al., 2023). Despite the well-documented empirical fact, we still know little about the consequences of such heterogeneity for optimal policy, in particular, the design of the pension scheme. This paper fills this gap.

In this paper, I apply a life-cycle model with heterogeneity in job stability, endogenous human capital accumulation, and consumption-saving and retirement decisions calibrated to the U.S. economy. I derive the optimal degree of progressivity for the pension system and study how a shift in the job-stability distribution towards more stable jobs affects the design of optimal pension systems. I find that abolishing the Social Security cap and increasing the degree of pension progressivity is optimal. The optimal pension system achieves a welfare gain of 0.75% of lifetime consumption for labor market entrants relative to the current U.S. pension system. Following the observed macroeconomic shift in the job stability distribution towards higher job stability since the 1990s in the United States, the optimal pension system becomes less progressive. However, the welfare gain from implementing the optimal pension system remains sizeable in the economy with higher job stability.

The life-cycle model features a frictional labor market where workers search on and off the job, make consumption-saving decisions in incomplete financial markets, and invest in risky human capital. When workers become eligible for pension benefits, they decide on their retirement age. Workers are heterogeneous with respect to their job-separation rates at each point in time of their life cycle, which builds on previous literature on heterogeneity in job stability (Pinheiro and Visschers, 2015; Kuhn and Ploj, 2020; Jarosch, 2023). The novel feature of this model is the incorporation of heterogeneity in job stability while treating human capital accumulation and retirement as endogenous. By allowing jobs to differ in their separation rates, I am able to analyze how the changing distribution of job stability shapes the optimal pension system. Endogenous human capital accumulation is an important determinant of earnings (Becker, 1964; Ben-Porath, 1967) and is key to explain increasing life-cycle earnings inequality (Hubmer, 2018; Jung and Kuhn, 2019). The interaction between human capital and job stability amplifies earnings losses upon job loss into long-run consequences (Jarosch, 2023). Workers get stuck in bad, unstable jobs as they fall down the job ladder and fail to accumulate human capital and life-cycle savings. These long-run consequences on earnings put forward the importance of job-stability heterogeneity in analyzing optimal pension systems. In the absence of progressivity, the pension system transmits interrupted work histories into the old age as pension entitlements depend on the pre-retirement history of earnings. I calibrate the model to the U.S. economy and the Social Security system, specifically, the empirical profiles of labor market transition rates, earnings profile (mean and variance), tenure distribution, and wealth-to-income ratios, and show that the model matches the data remarkably well.

In analyzing the optimal pension system, I focus on the classical trade-off between insurance and incentives for the policymaker. A progressive pension system assigns higher replacement rates for low-income workers compared to high-income workers and therefore, provides insurance against unstable employment histories and low lifetime earnings. However, this comes at the cost of reducing the return of human capital investment and thus, discouraging workers to invest in human capital. The progressivity level also affects retirement incentives of workers as the amount of pension wealth is one of the most important components in forming retirement decisions. Higher pension progressivity reduces labor supply at the extensive margin as low-income workers retire earlier. Another important margin in the analysis of the optimal pension system is the payroll tax finance of the pension system which plays a key role for welfare of workers in the presence of a borrowing constraint. Payroll taxes depress consumption of young low-income workers who are constrained in borrowing against their future income. Therefore, increasing the provision of pension coverage through higher payroll taxes may lower welfare of young workers even if they face a higher total amount of pension wealth.

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

The optimal pension system with higher progressivity level decreases consumption inequality over the life cycle by almost one third compared to the current U.S. Social Security system and it provides insurance to workers who suffer from frequent career interruptions and low lifetime earnings. On the other hand, the increase in pension progressivity distorts human capital investment and retirement decisions. The average stock of human capital decreases by 1.5% and the average retirement age decreases by 0.4% of total years worked in response to the policy change. However, these distortionary effects are not large enough to offset the insurance effect of the optimal policy for two reasons. First, young workers do not reduce their human capital investment. They strive to accumulate human capital to increase their prospective labor earnings. A lower return on human capital investment in the distant future has little impact on investment decisions for the young. Hence, young workers barely change their investment decision. The intuition is similar to Michelacci and Ruffo (2015) who show that an optimal age-dependent unemployment insurance system should provide higher insurance to young workers as the distortionary effect on human capital investment decision of young workers is small due to their high returns from human capital investment. Second, changes in retirement decisions depend on the accumulated stock of human capital. Less productive workers decide to retire earlier, whereas productive workers postpone their retirement in face of the increase in pension progressivity. The delay of retirement of productive workers partly offsets the distortion of retirement incentives on low-productivity workers such that the aggregate distortionary effect remains small.

I also analyze the transition from the initial economy to the economy with the optimal pension system assuming that the progressivity level is gradually increased over a time horizon of 40 years. While all labor market entrants benefit from the optimal progressive pension system as they are ex-ante identical, the welfare effect from increasing the progressivity level is heterogeneous for older workers depending on their labor market history. At age 50, workers with low job stability (bottom 25% of job-stability distribution) have a welfare gain of 0.6%, whereas workers with high job stability (top 25% of job-stability distribution) have a welfare gain of 0.3%. Importantly, the optimal policy induces a positive welfare gain for most of the workers except for the very top-income workers: the abolishment of the Social Security cap raises the tax revenue from high-income workers and reduces the payroll tax rate, which benefits the vast majority of workers in the economy, leading to a positive welfare gain.

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

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

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, followed by the baseline calibration in Section 4. Section 5 analyzes the effects of heterogeneity in job stability on inequality in labor market outcomes and its life-cycle consequences. I analyze the optimal pension system for the baseline economy in Section 6. Section 7 explores the consequences of a shift in job-stability distribution in the U.S. labor market on the optimal design of pension systems. Section 8 concludes.

2. Related literature

This paper contributes to three important strands of macroeconomic literature. First, by analyzing the policy implications of heterogeneity in job stability, it contributes to the growing strand of the literature that provides evidence for heterogeneity in job stability and labor market dynamics (Ahn et al., 2023; Gregory et al., 2021; Molloy et al., 2020). Among others, Jung and Kuhn (2019) find considerable heterogeneity in job 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 job stability, Jarosch (2023) points out that the interaction between human capital and job 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 and Ploj (2020) also provide evidence for job-stability heterogeneity in the U.S. labor market. Based on a life-cycle model, they show that job stability at labor market entry significantly affects income and consumption level over the whole life. I follow the strand of the literature that assumes that heterogeneity in employment stability is generated due to differences across jobs (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., 2021). The model implications in this paper are, however, also consistent with this strand of the literature. This model generates very stable and unstable career paths across workers and in fact, a worker with a very stable job is observationally indistinguishable from a very stable worker type within the model. Hence, differences in employment stability in this paper can be interpreted as worker types as in the above-mentioned empirical approaches.

Second, this paper relates to the literature on unemployment risk and inequality in lifetime earnings. Using the life-cycle model, I discuss the interaction between heterogeneity in job stability and how unstable career paths affect lifetime earnings of workers, which consequently affects their pension entitlements. These results complement the large strand of empirical literature on the relationship between inequality in lifetime earnings and labor market risks. For example, Bonhomme and Robin (2009) explore inequality in employment and lifetime earnings for France and find that unemployment is a key driver of lifetime earnings inequality. Similarly, Bowlus and Robin (2012) show that unemployment mobility has important consequences on lifetime earnings inequality. Bönke 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. (2019) find that heterogeneity in labor market outcome is a major source of inequality in lifetime earnings for workers with same abilities.

Lastly, this study contributes to the literature on optimal design of pension systems by applying a life-cycle framework which explicitly takes heterogeneity in job stability into account. The application of a life-cycle model is a frequently used approach in the previous macroeconomic literature on optimal pension systems. Beginning with the work by Feldstein (1974), a large body of literature investigates the effects of pension systems in a life-cycle framework. For example, Hubbard and Judd (1987) consider a model with capital-market imperfections to show that the introduction of Social Security does not increase an economy's welfare as much as in an economy without borrowing constraints. A more recent work is the study by Fehr and Habermann (2008) who analyze the welfare effects of introducing a more redistributive pension system to Germany. They find that the positive liquidity and income insurance effects are large enough such that an increase in progressivity enhances welfare despite its distortive effects on labor supply. Fehr 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. (2022). This paper differs from the earlier work on pension reforms in that I study heterogeneity in job stability as a source of inequality in lifetime earnings. This also allows me to analyze how changes in the job-stability distribution affect the optimal pension system.

3. Life-cycle model

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

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

The period budget constraint of an agent is given by

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

where a and h denote the risk-free asset and the stock of human capital of the agent, respectively. Moreover, r denotes the risk-free rate in the economy, and y denotes the labor income of the current period including transfers. 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 $y(w_j, h_j, e) = w_j h_j$. When unemployed, a worker receives a transfer of $y(w_j, h_j, n) = b w_j h_j$, that is, an agent gets a benefit proportional to the labor earnings from the last job. The replacement rate of the unemployment insurance system is denoted by b. The model captures declining benefits in the spell of unemployment by assuming that the wage of the last job drops from w_k to $\max\{w_{k-1}, w_1\}$ if agents stay unemployed and therefore continue receiving unemployment benefits.

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

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

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

[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.3 for details.

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 of an employed agent is as follows:

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

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

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

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

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

$$\text{s.t.} \qquad c = (1+r)a + y(w,h,n) - a'.$$

Note that an unemployed worker receives unemployment benefits $y(w,h,n)$ which is reduced in the next period if the worker remains unemployed. At the search stage, an unemployed worker faces the value function

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

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

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

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

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

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

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

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

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

3.2. Payroll tax finance of the pension system

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

$$\mathbb{E}\left[\sum_{t=t^R}^{T}\frac{\omega(\bar{h})}{(1+r)^{t-1}}\right]=\mathbb{E}\left[\sum_{t=1}^{t^R-1}\frac{\tilde{y}(a,w,\lambda,h,t)}{(1+r)^{t-1}}\cdot\tau\right] \tag{2}$$

must be satisfied where

$$\tilde{y}(a,w,\lambda,h,t)=\begin{cases} y(a,w,\lambda,h,t) & \text{if } y(a,w,\lambda,h,t)<\text{cap}\\ \text{cap} & \text{if } y(a,w,\lambda,h,t)\geq\text{cap}\end{cases}$$

The cap denotes the maximum taxable earnings. The parameter τ in Equation (2) denotes the payroll tax rate. Taxes are only levied on labor earnings of employed agents in the working phase. Unemployed agents and agents in the retirement phase do not pay taxes. The tax rate is assumed to be constant over all periods and across all agents.

4. Calibration

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

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

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

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

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

retirement decisions. At age 67, all workers who decided to remain in the labor force up to this age are shifted to the retirement phase.

Wages and job-separation probabilities are discretized on grids with $[\underline{w}, \overline{w}]$ and $[\underline{\lambda}, \overline{\lambda}]$, respectively. The lowest quarterly separation probability is $\underline{\lambda} = 0.006$ which corresponds to lifetime jobs with an expected job duration of 42 years. The highest separation probability is set to $\overline{\lambda} = 0.35$. The job-offer distribution consists of a joint distribution of job-separation probability and wage. See Section A.4.1 in the Appendix for a detailed description of the job-offer distribution.

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

4.1. U.S. pension system and retirement elasticity

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

$$\omega(\bar{y}) = \begin{cases} \phi \cdot (\bar{y})^{1-\gamma} & \text{if } \bar{y} < \text{cap} \\ \overline{\omega} & \text{if } \bar{y} \geq \text{cap} \end{cases} \tag{3}$$

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

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

Figure 1. Pension System.

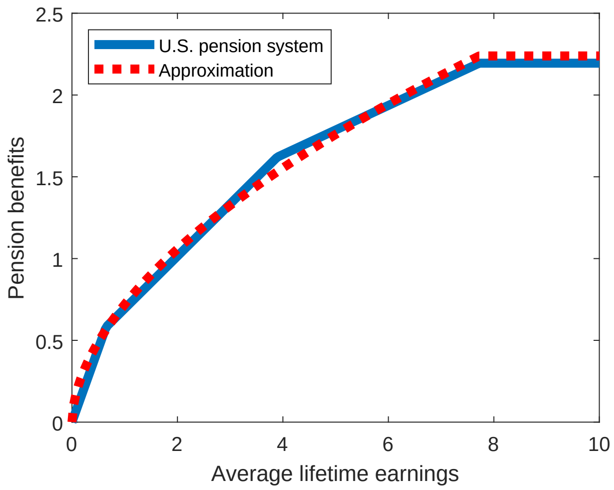

Notes: Comparison of the U.S. Social Security system (solid line) and the approximating function (dotted line).

In order to calibrate the pension system in the model to the U.S. pension system, I consider the U.S. Social Security legislation for 2019. Based on this, I set the Social Security cap to \$132,900 and the first and second bendpoints to \$926 and \$5,583, respectively. Then, I apply the retirement benefits formula as documented in http://www.ssa.gov/OACT/COLA/piaformula.html

$$\Omega(\bar{y}) = \begin{cases} 0.9\bar{y} & \text{if } \bar{y} < bp_1, \\ 0.9bp_1 + 0.32(\bar{y} - bp_1) & \text{if } bp_1 \leq \bar{y} < bp_2, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(\bar{y} - bp_2) & \text{if } bp_2 \leq \bar{y} < cap, \\ 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(cap - bp_2) & \text{if } \bar{y} > cap \end{cases} \tag{4}$$

where $\bar{y}$ denotes the average lifetime earnings, $\Omega(\bar{y})$ denotes the assigned benefit level, and bp_1 and bp_2 denote the two bendpoints. To match the pension system in the model to the U.S. Social Security, I target the net pension replacement rate for an average worker in terms of lifetime earnings in the model to the observed net pension replacement rate of 49.4% for an average worker in the United States reported by OECD (2019). The estimated parameters are $\phi = 0.72$ and $\gamma = 0.44$. Figure 1 shows the U.S. Social Security benefit function and the approximating function in the model.

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

Figure 2. Labor market transitions and tenure.

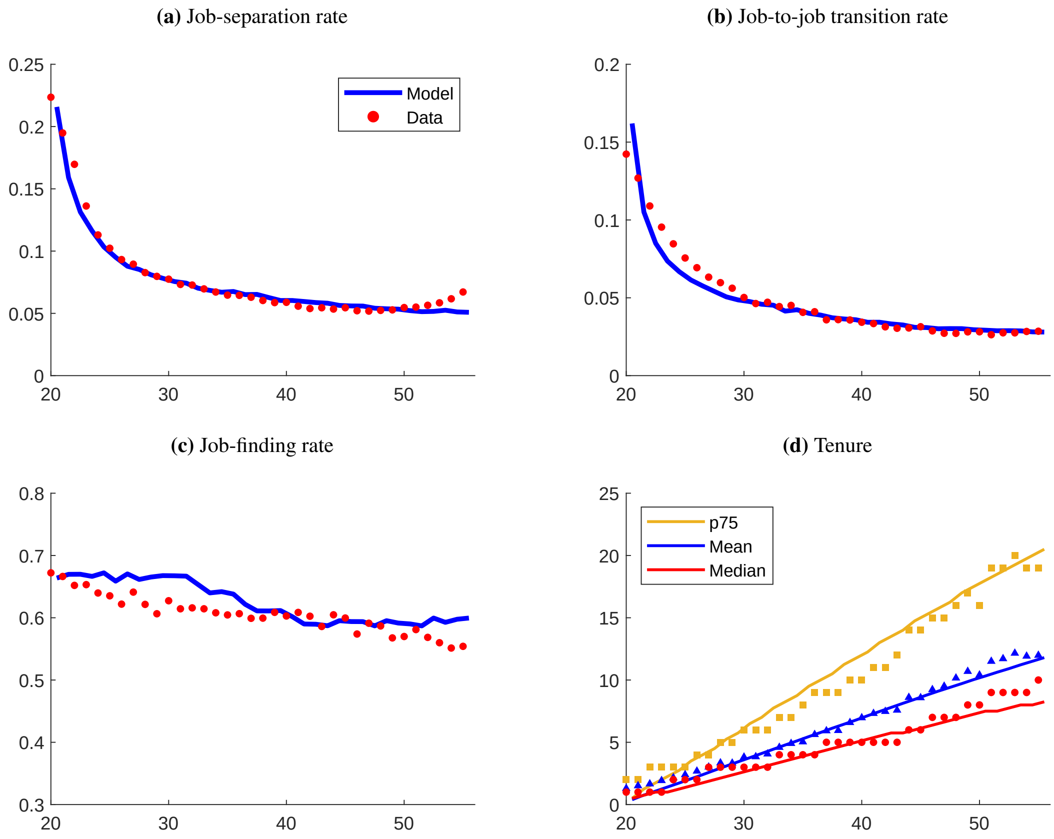

Notes: This figure shows quarterly life-cycle transition rates and tenure in years by age. Panel a shows the average separation rate, Panel b the job-to-job transition rate, and Panel c the job-finding rate. Panel d shows the tenure distribution for the mean, median, and the 75th percentile of the distribution. The dots display the empirical profiles and the solid lines show the model profiles. The empirical profiles are derived from the Current Population Survey.

regress the average lifetime earnings of workers on cubic polynomials of the final level of labor earnings of each worker which consist of human capital and wages. In fact, the final level of human capital and wages attained by workers until retirement explains more than 90% of variation in lifetime earnings. Section A.4.3 in the Appendix discusses the accuracy of this approximation of average lifetime earnings in detail.

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

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

where μ denotes the location parameter and σ denotes the scale parameter of the logistic distribution. I set on average $\mu = 13.68$ and $\sigma = 2.9$. The shock ε drawn from the logistic distribution can be interpreted as allowing for heterogeneity in preferences for leisure. It can be also interpreted as deriving an unexpected positive utility from retirement such as health problems, which are an important cause for retirement (see, for example, Dwyer

Figure 3. Earnings and wealth.

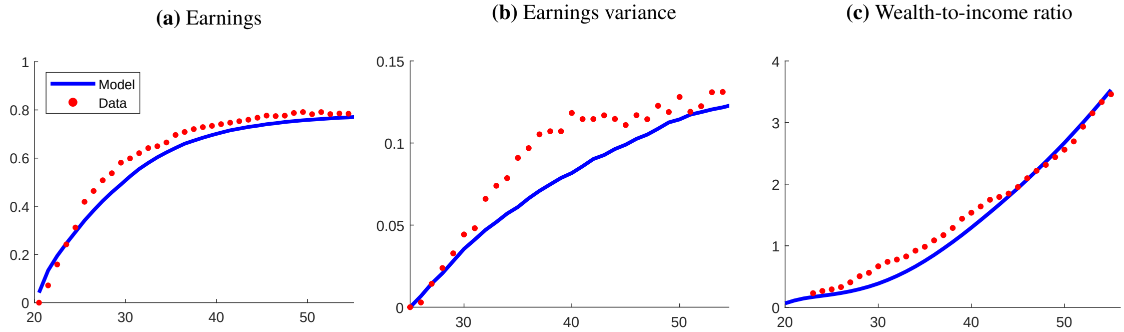

Notes: This figure shows quarterly life-cycle profiles of earnings and wealth-to-income ratio by age. Panel a shows the average labor earnings, Panel b the variance of labor earnings, and Panel c the wealth-to-income ratios. The dots display the empirical profiles and the solid lines show the model profiles.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

and Mitchell, 1999; Disney et al., 2006; Jones et al., 2010; Galama et al., 2013; Trevisan and Zantomio, 2016).

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

Figure 2 compares the life-cycle profiles of job-separation, job-to-job, job-finding rates, and the tenure distribution in the model and in the data. Both job-separation and job-to-job rates strongly decline between age 20 and 30 and decline until age 50. The model profiles fit the empirical profiles well with a steep decrease of job-separation rate and job-to-job transition rate at the beginning of the life cycle, followed by a flattening of the profiles after age 30. The empirical profile of job-finding rates also exhibit a decreasing trend over the life cycle which is closely matched by the model. In Figure 2d, I compare the empirical profiles of mean, median, and the 75th percentile of the tenure distribution with their model counterpart. The increasing gap over the life cycle between the 75th percentile, the mean, and the median of the tenure distribution indicates that the tenure distribution gets more dispersed over time. Heterogeneity in job stability generates differences in tenure across workers, and these differences grow over the life cycle both in the data and in the model. Figure 3 displays the profiles of mean log earnings, earnings variance, and wealth-to-income ratio. The model is able to match the strong increase in earnings at the beginning of the life cycle and the decreasing income growth over time when workers get closer to the end of prime age. The model also matches the steep increase of the earnings variance as well as the life-cycle profile of wealth-to-income ratio.

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

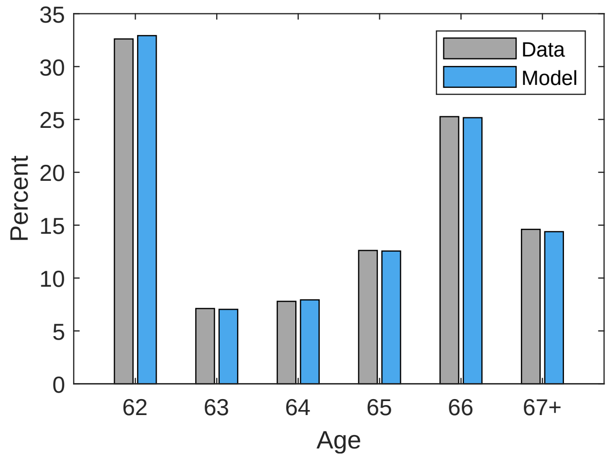

Notes: This figure shows the distribution of retirement age in the data for the United States in 2019 and in the baseline calibration of the model. The data source is Table 6.A4 of Social Security Administration (2020), Annual Statistical Supplement, available at https://www.ssa.gov/policy/docs/statcomps/supplement/.

4.2. Retirement age distribution and retirement elasticity

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

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

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

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

5. Heterogeneity in job stability and life-cycle dynamics

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

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

5.1. Inequality in labor market outcomes

Figure 5a illustrates how the two groups of agents differ in terms of their average job-separation rate over the life cycle. Comparing the two profiles in Figure 5a, 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 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 job 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 job stability between agents in the top and bottom quartile of average job stability over the life cycle remains sizeable. Compared to the top quartile, workers in the bottom quartile fail on average to find stable jobs over the life cycle. Overall, the difference in job stability is significant for the two groups of workers and remains persistent until the end of the working phase.

Figure 5. Inequality in labor market history.

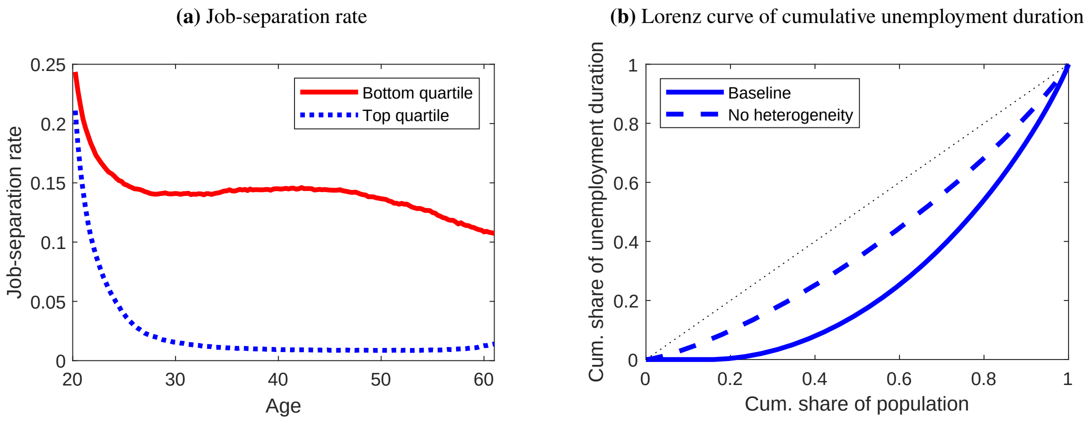

Notes: Panel a shows the average life-cycle profile of job-separation rate. The solid line displays the average life-cycle profile of job-separation rate of workers with most unstable jobs (bottom 25% of job-stability distribution). The dashed line represents workers with most stable jobs (top 25% of job-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 heterogeneity in job stability conditional on age.

This heterogeneity in job stability translates into a large inequality in employment history. Figure 5b shows the Lorenz curve for cumulative unemployment duration of prime-age workers between ages 25 and 55. The solid line displays the Lorenz curve for the baseline economy and the dashed line a counterfactual economy without heterogeneity in job 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 job 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 job stability conditional on age. Comparing the baseline economy and the counterfactual economy, the Lorenz curves indicate that heterogeneity in job stability is a key driver of inequality in employment history. In the baseline economy, about 40% of all workers account just for 8% of total cumulative unemployment duration of all workers, whereas the Lorenz curve of the counterfactual does not exhibit a conspicuous curvature.

5.2. Life-cycle consequences of job stability

Figure 6a shows the profiles of average human capital level for workers in the top and bottom quartile of the average life-cycle job-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 job 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

Figure 6. Average life-cycle profiles by job stability.

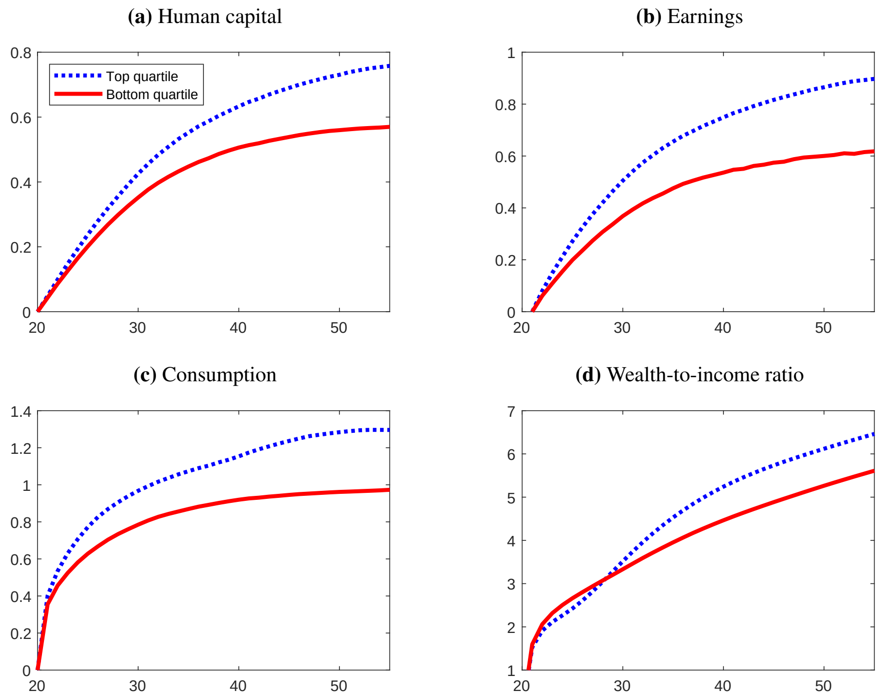

Notes: This figure shows the average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), and consumption (lower left panel), and wealth (lower right panel). The profiles are the log deviations from the initial values of the respective profiles. In all plots the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

capital between the two groups of agents rises over the life cycle. In particular, starting from the same initial level of human capital, the top quartile has on average 22% higher stock of human capital at the end of the working phase.

A similar pattern is observed in Figure 6b for average earnings and in Figure 6c 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 features a smaller growth over the life cycle.

Figure 6d shows the average life-cycle profiles of wealth. For workers in the top quartile, wealth starts to grow strongly 10 years after labor market entry. The profile of the bottom quartile also increases throughout life, but the growth in wealth is dampened by low earnings growth. Towards the end of the working phase, workers in the bottom quartile have significantly lower assets than agents in the top quartile. Since a high job-separation probability leads to fewer opportunities to invest in human capital and to more frequent job losses, workers in the bottom quartile have on average lower earnings which lead to lower savings. Moreover, workers dissave upon job loss in order to smooth consumption, which in turn decreases asset accumulation.

In sum, the risk of becoming unemployed constitutes a significant source of earnings uncertainty for individuals in the labor force and heterogeneity in job stability is a key driver of inequality in lifetime earnings. 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. Job 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 job stability affects the distribution of lifetime earnings and consumption. Therefore, the degree of heterogeneity in job stability has crucial implications on the desired level of redistribution in the economy and shapes the optimal design of pension systems. A progressive pension system implies decreasing replacement rates in average lifetime earnings and redistributes earnings from high-income workers to low-income workers, reducing consumption inequality across workers after retirement. Progressive pension benefits alleviate earnings shocks accumulated over the working life which would be otherwise fully carried over to the retirement.

6. Optimal progressive pension system

Starting from the baseline economy described in the previous section, I derive the optimal pension system. Holding the total expected pension benefits constant at the baseline level, I search for the optimal progressivity parameter γ and the optimal level of Social Security cap that lead to the highest welfare in the economy. I set up a grid for the parameter γ and the cap and compute the corresponding parameter ϕ and the payroll tax rate that achieve budget balance for the government while the total amount of benefits is equal to the total amount of benefits in the baseline economy. The next subsection explains the welfare measure which I apply for the welfare analysis.

6.1. Welfare measure

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

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

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

6.2. Measure of progressivity level

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

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

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

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

6.3. Optimal pension system

Figure 7 displays the welfare change induced by varying the shape parameter γ and the Social Security cap of the pension system in consumption-equivalent variations. Recall that for a given level of the Social Security cap, a pension system with $\gamma = 0$ implies that retirement benefits linearly increase in lifetime earnings up to the cap. An increase in γ to $\gamma > 0$ makes the pension system more progressive. The figure shows that welfare increases with the Social Security cap and a repeal of the cap is optimal. The concave shape of the welfare as a function of the shape parameter γ reveals the trade-off between redistribution and incentive distortions for the optimal design of the pension system. Progressivity offers insurance against unstable employment histories and low lifetime earnings, but it comes

Figure 7. Welfare and changes in progressivity level.

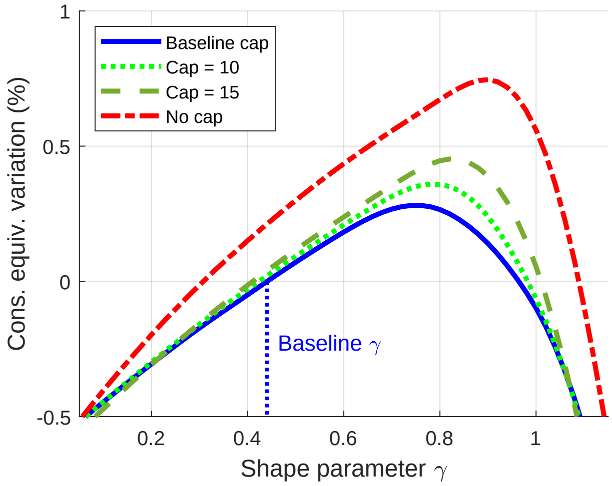

Notes: This figure shows the change in welfare in consumption-equivalent variation induced by changes in the shape parameter γ and the Social Security cap. The solid line displays the welfare change as a function of the shape parameter γ when the Social Security cap is at the baseline level (7.72). The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

at the cost of distorting human capital investment and retirement decisions.

Table 1 displays the policy parameters associated with each progressivity level γ. To give a manageable overview of the results, Table 1 displays the results only for selected values of parameters. The parameters in the baseline economy are displayed in the first line where $\gamma = 0.44$ and $\phi = 0.72$ with a payroll tax rate of 9.17%. Comparing the welfare change across all parameter specifications, the optimal policy is a pension system with no Social Security cap and the shape parameter set to $\gamma = 0.9$. The corresponding progressivity level is 42.23% and this parameter combination leads to a welfare gain of 0.75% in terms of consumption-equivalent variation for a labor market entrant. The payroll tax rate that achieves budget balance for the pension system decreases to 8.9% and the corresponding level of parameter ϕ increases to 1.12. To understand where this welfare gain comes from, in the next subsections I investigate how retirement decisions and life-cycle dynamics change under the optimal pension system compared to the baseline economy.

Note that the model assumes that all workers are ex-ante identical in their wealth level, human capital, and the initial labor market status. The optimal pension system therefore induces the same welfare change for all workers in the model. Ex-ante heterogeneity in the state variables, however, may have important welfare implications on the optimal pension system. In Appendix A.5, I discuss the welfare changes of the optimal pension system as functions of job-separation rate, wealth level, and human capital. I also show in Appendix A.6 that the optimal pension system remains unaffected if the unemployment insurance system changes. Finally, I consider a model with human capital depreciation in Appendix A.7. Overall, all qualitative results regarding the optimal pension system

Table 1. Welfare change in the baseline economy under alternative pension systems.

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

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

remains unchanged under the alternative model specifications.

How does heterogeneity in job stability affect the optimal level of progressivity of the pension system? In Section 5, I have shown that heterogeneity in job stability generates differences in employment histories and accounts for a large fraction of inequality in lifetime earnings. In order to analyze how heterogeneity in job stability affects the optimal pension system, I repeat the welfare analysis for the counterfactual economy without heterogeneity in job stability. In the counterfactual economy, there is no cross-sectional heterogeneity in job stability conditional on age, but the separation rates vary over the life cycle. Figure 8 displays the welfare changes induced by variations in the retirement benefit formula in the baseline economy and the economy without heterogeneity in job stability. In Figure 8a, the Social Security cap is kept constant at the baseline level. In the economy without heterogeneity in job stability and a fixed Social Security cap, the current U.S. pension system is close to optimal. Any deviation of the shape parameter γ from its baseline specification leads to welfare losses. Therefore, in the absence of heterogeneity in job stability, the optimal pension system implies a much lower level of progressivity.

The abolishment of the Social Security cap is, however, welfare increasing in both the counterfactual and the baseline economy. Figure 8b shows the welfare changes in the shape parameter in combination of the abolishment of the Social Security cap. In the absence of the cap, an increase in the shape parameter leads to welfare gains in both economies. Yet, the optimal level of progressivity remains lower in the counterfactual

Figure 8. Comparison of welfare changes in the baseline economy and the economy without heterogeneity in job stability.

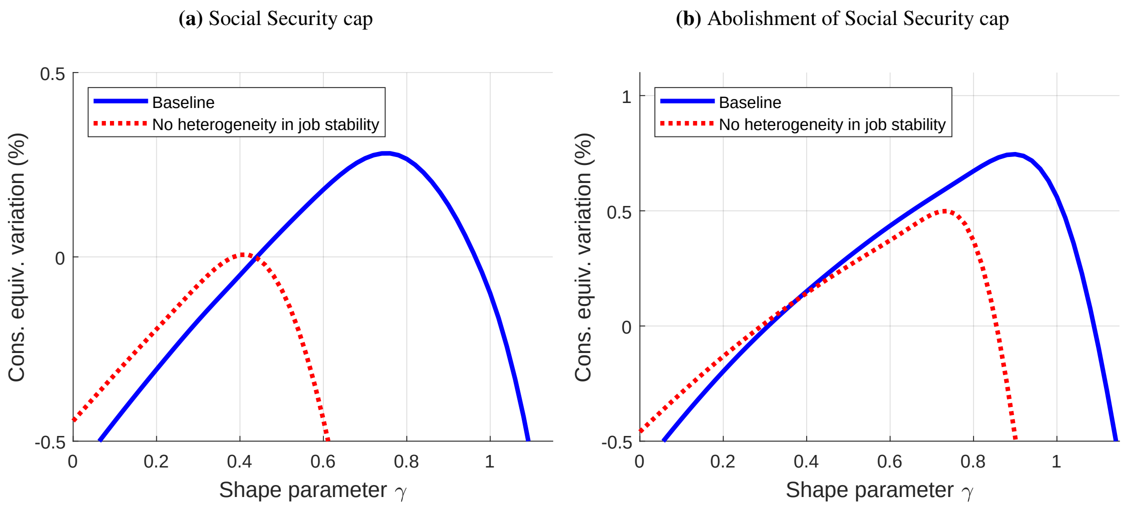

Notes: The graphs show the welfare changes in the baseline economy (solid lines) and the economy without heterogeneity in job stability (dotted lines) induced by variations in the shape parameter of the retirement benefit function. In the left panel, the Social Security cap is fixed to its initial level.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

without heterogeneity in job stability. In the baseline economy, the optimal progressivity level is 42.23%, whereas in the counterfactual economy, the optimal pension system implies a progressivity level of 32.45%. Hence, a disregard of heterogeneity in job stability decreases the optimal level of progressivity by a quarter.

6.4. Welfare decomposition

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

Table 2. Welfare decomposition.

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

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

effect and the sum of all components.

The decrease in the payroll tax rate under the optimal pension system induces a welfare gain of 0.27%. The abolishment of the Social Security cap increases the contribution of high-income workers so that the tax rate declines. A lower tax rate is particularly beneficial for young workers as payroll taxes depress consumption when earnings are low and workers have a high motive for precautionary savings. The choice of the retirement age allows workers to adjust their labor supply at the extensive margin in response to the policy change and induces a welfare gain of 0.07%. The adjustment of labor supply allows to react to the changes in pension wealth. The insurance effect of the optimal pension system leads to a welfare gain of 1.78% as workers have a higher insurance against bad labor market outcomes. The distortinary effects arise as workers invest less into human capital and reduce on average their labor supply under the optimal pension system.

6.5. Retirement incentive distortion

The optimal pension system changes the retirement decision of workers and thus entails a distortionary effect. This section aims to examine this distortionary effect more closely. Table 3 compares the average retirement age in the baseline economy and the economy under the optimal policy. The first row shows the mean retirement age of all workers for each economy and indicates that the mean retirement age decreases from 64.33 to 64.16. This result implies that on average, workers choose to retire earlier in response to the increase in pension progressivity.

Yet, the change in the retirement age differs for workers depending on their stock of human capital. The second and third rows of Table 3 show the mean retirement age of two groups of workers: one group comprises workers who end up with a stock of human capital that is below the median of all workers; the other group refers to workers who achieve a human capital stock above the median of all workers. Workers below the median choose to retire earlier in the economy under the optimal policy than in the baseline economy, whereas workers above the median choose to stay longer in the labor force.

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

Table 3. The mean retirement age in the baseline economy under the ex-ante optimal pension system.

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

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

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

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

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

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 job stability over the life cycle. Taking agents up to the first quartile and above the third quartile of the distribution, I investigate the life-cycle consequences of the optimal progressive pension system for the two groups of agents.

Figure 9 shows the relative deviation from the baseline economy of human capital, labor earnings, consumption, and wealth in the economy under the optimal pension system. The first finding is that on average, the increase in pension progressivity discourages human capital investment over the life cycle. In Figure 9a, starting at age 30, the relative change

Figure 9. Relative change in average life-cycle profiles by job stability.

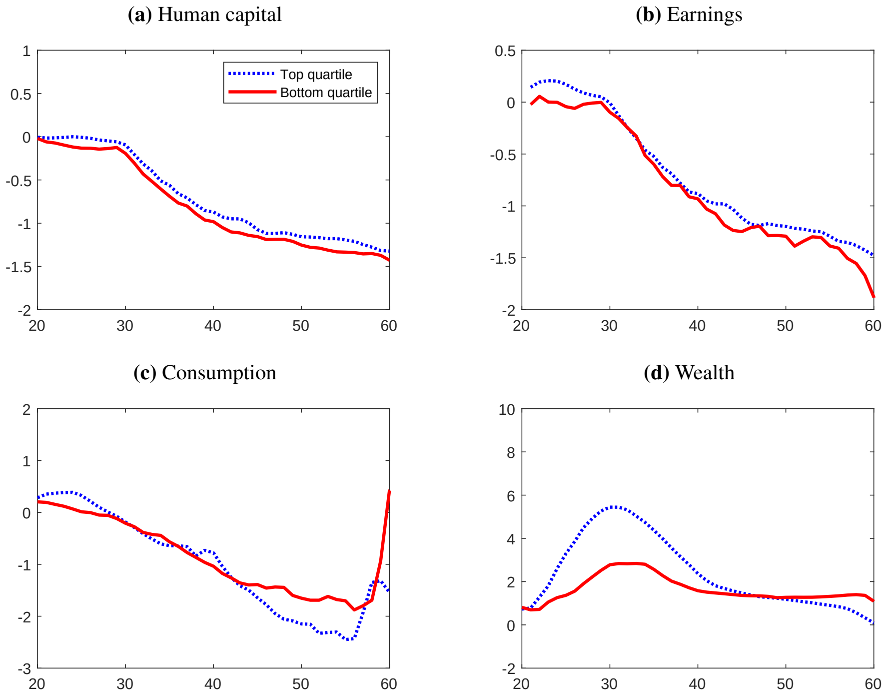

Notes: This figure shows the percent deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy under the optimal pension system compared to the baseline economy. In all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

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.[3] The gap in human capital level grows over life and amounts to -1.4% for agents in the top and -1.5% for agents in the bottom quartile at the early retirement age. Because achieving a high level of human capital is associated with a lower replacement rate in the retirement benefit formula than in the baseline economy, the return on human capital decreases for workers in the top quartile. Consequently, agents employed in stable jobs for whom human capital investment is large otherwise strongly decrease their effort provision for human capital accumulation. A similar reason accounts for lower human capital investment of workers in the bottom quartile.

Interestingly, during the first 10 years after labor market entry, workers do not reduce their investment in human capital. This is because human capital investment is highly productive in the early working phase. Young workers strive to accumulate human capital to increase prospective labor earnings growth and lifetime earnings such that human capital accumulation mostly happens at the beginning of the life cycle. The decrease in return on human capital investment, which realizes after retirement, has little impact on the incentives for human capital accumulation and labor market outcomes of young workers. As a consequence, the policy change does not impede human capital accumulation in

[3]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.

the early working phase. When the retirement age gets closer and the disincentives for human capital investment start to grow and become more relevant in the later stage of life, workers reduce their effort provision in any case because investment in human capital becomes unproductive. Hence, the gap in human capital between the baseline economy and the economy with higher pension progressivity does not grow further and remains almost unchanged close to retirement.

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

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 9b and Figure 9c. The first observation is that the relative changes in the profiles of average earnings closely follow the shape of the profiles in Figure 9a. For both groups of workers, the decrease in the average human capital implies lower average earnings compared to the baseline economy. The average consumption of workers in the top quartile decreases towards the later stage of the working phase. Also for the bottom quartile, the average consumption decreases because the average earnings become lower. However, the average consumption strongly increases 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 9d reflects this change in the consumption-saving behavior. For both groups of workers, the average asset level increases at the beginning of the life cycle, but declines and reaches the baseline level at the end of the working phase. Workers in the bottom quartile of the job-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 10 displays that the optimal pension system reduces consumption inequality over the life cycle. The profile for variance of log consumption in the baseline economy almost linearly increases over life by more than 50% reaching a value of 0.2 before retirement.

Figure 10. The life-cycle profiles for variance of log consumption and Gini coefficient of consumption.

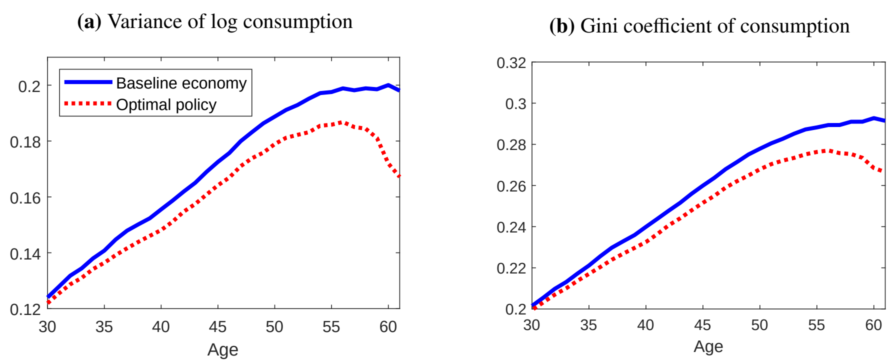

Notes: The left panel shows the variance of log consumption and the right panel the Gini coefficient of consumption over the life cycle. The solid lines display the inequality measure for the baseline economy and the dotted lines for the economy under the optimal pension system. The variances and Gini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

In the economy under the optimal pension system, the profile becomes flatter and has a concave shape. At the end of the working phase, the variance of log consumption reaches a value of 0.17. The optimal pension system therefore reduces the increase by almost one third of the increase observed in the baseline economy.

Comparing the profiles of the Gini coefficient of consumption in Figure 10b, a similar change as for the profile of variance of log consumption is observed. Whereas the Gini coefficient increases linearly in the baseline model and reaches a value of 0.29 before retirement, the optimal pension system dampens the increase in the Gini coefficient over life. In the economy under the optimal pension system, the Gini coefficient of consumption at the end of the working phase decreases to 0.27.

6.7. Transition to the optimal pension system

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

More specifically, the shape parameter γ and the Social Security cap are gradually raised by a constant percentage over a time horizon of 40 years. This is depicted in Figure 11. In year zero, the baseline pension system is implemented. Starting in year one, the government implements the new policies with higher parameter values of γ and the cap. Similar to Auerbach and Kotlikoff (1987), all workers have perfect foresight along the transition path regarding the pension system that will become effective when their cohort reaches the retirement age.

The welfare changes for the different cohorts are shown in Figure 12. These welfare

Figure 11. Transition to the optimal policy.

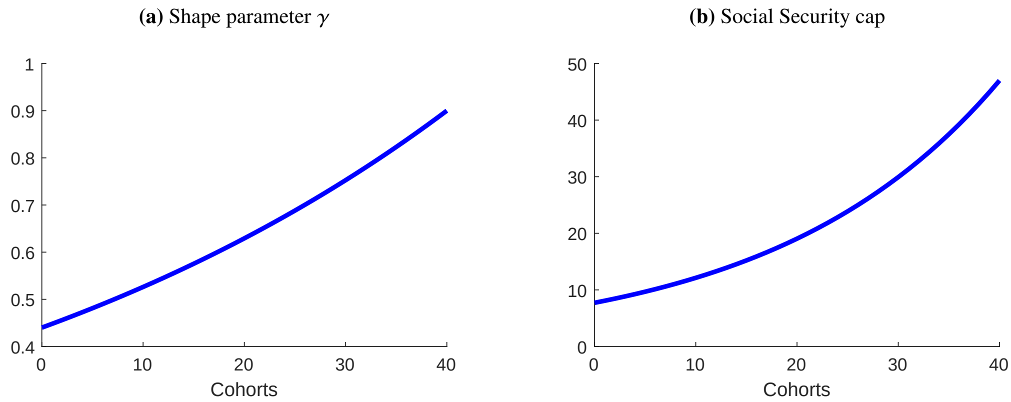

Notes: The left panel shows the shape parameter γ and the right panel the Social Security cap along the transition path to the optimal pension system. Cohort denotes the sequential order of cohorts that face the new policy parameters for the pension system.

changes are computed in the period the government credibly announces the entire path to the optimal pension system. As the policy parameters are changed gradually for each cohort and the remaining lifetime differs for different cohorts until the retirement age, the transition to the optimal policy induces differential welfare effects across cohorts. The transition to the optimal policy induces for the majority of workers a positive welfare effect which is increasing from zero to 0.75% when reaching the steady state with the optimal pension system.

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

Note that older workers close to the retirement also have positive welfare gains from increasing the progressivity level. At age 50, workers with low job stability (bottom 25% of job-stability distribution) have a welfare gain of 0.6% in terms of life-time consumption. The welfare gain for high job-stability workers (top 25% of the job-stability distribution) is at 0.3% which is much lower compared to low-job stability workers, but still positive. The reason why even older workers benefit from the implementation of the

Figure 12. Welfare along the transition path.

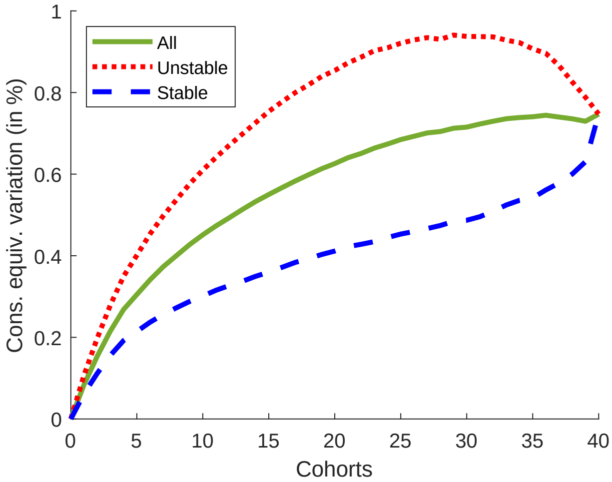

Notes: This figure shows the change in welfare in consumption-equivalent variation along the transition path to the optimal pension system. The solid line displays the welfare change for all workers. The dotted line and the dashed line show the welfare change for workers with unstable (bottom 25% of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

higher progressivity level is due to the abolishment of the Social Security cap. The tax revenue from high-income workers increases as the Social Security cap is abolished so that the vast majority of workers in the economy face a lower payroll tax rate. This induces a positive welfare effect even for most of the older workers close to the retirement.

7. Consequences of a shift in the job-stability distribution in the U.S. labor market

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

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

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

Given these empirical observations, I study the consequences of a shift in the job-stability distribution on the optimal design of pension systems. To this end, I adjust the job-stability distribution of the baseline economy in order to capture the empirical finding that the change in job-separation rates is largely driven by the extent of short-duration jobs. While holding the job-separation rate of the most stable job constant at the baseline value, I reduce the separation rate of the most unstable job. As a result, the support of the job-stability distribution becomes narrower such that the degree of job-stability heterogeneity declines. The average job-separation rate decreases by 6% in the economy with higher job stability.[4]

7.1. Life-cycle consequences of a shift in the job-stability distribution

How does a lower job-separation rate affect life-cycle dynamics? To address this question, I analyze the relative deviations from the baseline economy of human capital, earnings, consumption, and wealth for the economy with higher job stability in Figure 13. Figure 13a shows that in the economy with higher job stability, workers in both the top and bottom quartile achieve on average larger stocks of human capital at the early retirement age. Because of the increase in average job stability, workers are less affected by career interruptions and thus make more human capital investments than in the baseline economy. The increase in human capital is larger for workers with most unstable jobs. Whereas human capital increases by 0.5% at the early retirement age for workers with most stable jobs, workers with unstable jobs achieve on average a human capital level that is 1.8% higher than the baseline level. Since job stability increases for jobs with the largest job-separation rates, the effect of higher job stability is strongest for workers who hold on average the most unstable jobs over the life cycle.

In Figure 13b, the average life-cycle profiles of labor earnings reflect these changes in human capital. Earnings become around 2.7% higher for workers with the most unstable

[4]In Appendix A.8, I show how the average separation rate changes over the life cycle in the economy with higher job stability.

Figure 13. Relative change in average life-cycle profiles by job stability.

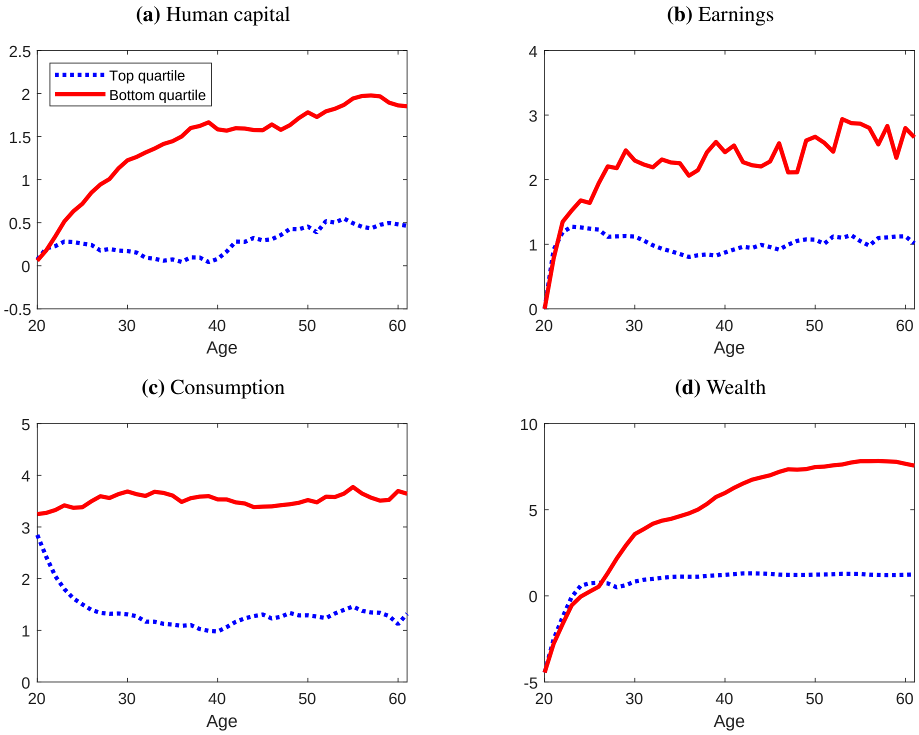

Notes: This figure shows the percent deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy with higher job stability compared to the baseline economy. In all plots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

jobs, and around 1% for workers with most stable jobs over life. The increase in average earnings is larger than the relative increases in human capital since higher job stability not only leads to larger human capital stocks, but also facilitates moving up the job ladder such that workers find jobs with higher wages. Moreover, both consumption and wealth increase in the economy with higher job stability. More stable career paths and lower earnings losses due to higher job stability enable workers to accumulate larger wealth. Before the early retirement age, workers with most unstable jobs have 7.5% higher wealth and workers with most stable jobs 1.2% higher wealth than in the baseline economy.

7.2. Distribution of job stability and optimal pension systems

In order to analyze how these changes in life-cycle dynamics affect the optimal pension system in the economy with higher job stability, I search for the optimal parameter combination of the parameter γ and the cap in the retirement benefit function. As in the previous section, policy changes are accompanied by a constant payroll tax rate on labor earnings and the parameter ϕ which scales the level of benefits in order to achieve budget balance for the pension system. The results are summarized in Table 4.

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

Table 4. Optimal pension systems in the baseline economy and in the economy with higher job stability.

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

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

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

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

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

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

8. Conclusion

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

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

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

References

Ahn, Hie Joo, Bart Hobijn, and Ayşegül Şahin (2023). *The Dual U.S. Labor Market Uncovered*. Working Paper 31241. National Bureau of Economic Research.

Auerbach, Alan J and Laurence J Kotlikoff (1987). "Evaluating fiscal policy with a dynamic simulation model." In: *The American Economic Review* 77.2, pp. 49–55.

Becker, Gary S. (1964). "Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education." University of Chicago Press.

Ben-Porath, Yoram (1967). "The production of human capital and the life cycle of earnings." In: *Journal of Political Economy* 75.4, pp. 352–365.

Bénabou, Roland (2000). "Unequal Societies: Income Distribution and the Social Contract." In: *American Economic Review* 90.1, pp. 96–129.

— (2002). "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" In: *Econometrica* 70.2, pp. 481–517.

Bilal, Adrien (2023). "The Geography of Unemployment." In: *The Quarterly Journal of Economics* 138.3, pp. 1507–1576.

Bloom, David E., David Canning, Richard K. Mansfield, and Michael Moore (2007). "Demographic change, social security systems, and savings." In: *Journal of Monetary Economics* 54.1, pp. 92–114.

Bonhomme, Stéphane and Jean-Marc Robin (2009). "Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990–2000." In: *The Review of Economic Studies* 76.1, pp. 63–92.

Bönke, Timm, Giacomo Corneo, and Holger Lüthen (2015). "Lifetime Earnings Inequality in Germany." In: *Journal of Labor Economics* 33.1, pp. 171–208.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang (Mar. 2022). "Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labour Supply?" In: *The Review of Economic Studies* 90.1, pp. 102–131.

Bowlus, Audra J. and Jean-Marc Robin (2012). "An International Comparison of Lifetime Inequality: How Continental Europe Resembles North America." In: *Journal of the European Economic Association* 10.6, pp. 1236–1262.

Braxton, J. Carter, Kyle F Herkenhoff, Jonathan L Rothbaum, and Lawrence Schmidt (2021). *Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter*. Working Paper 29567. National Bureau of Economic Research.

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

Burtless, Gary (1986). “Social Security, Unanticipated Benefit Increases, and the Timing of Retirement.” In: *The Review of Economic Studies* 53.5, pp. 781–805.

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

— (2007). “Future Social Security Entitlements and the Retirement Decision.” In: *The Review of Economics and Statistics* 89.2, pp. 234–246.

Cremer, Helmuth, Jean-Marie Lozachmeur, and Pierre Pestieau (2004). “Social security, retirement age and optimal income taxation.” In: *Journal of Public Economics* 88.11, pp. 2259–2281.

De Nardi, M, S Imrohoroglu, and TJ Sargent (Mar. 2001). “Saving and pension reform in general equilibrium models.” In: *Oxford Review of Economic Policy* 17.1, pp. 20–39.

Disney, Richard, Carl Emmerson, and Matthew Wakefield (2006). “Ill health and retirement in Britain: a panel data-based analysis.” In: *Journal of Health Economics* 25.4, pp. 621–649.

Dwyer, Debra S. and Olivia S. Mitchell (1999). “Health problems as determinants of retirement: are self-rated measures endogenous?” In: *Journal of Health Economics* 18.2, pp. 173–193.

Fehr, Hans and Christian Habermann (2008). “Risk Sharing and Efficiency Implications of Progressive Pension Arrangements.” In: *The Scandinavian Journal of Economics* 110.2, pp. 419–443.

Fehr, Hans, Manuel Kallweit, and Fabian Kindermann (2013). “Should pensions be progressive?” In: *European Economic Review* 63, pp. 94–116.

Feldstein, Martin (1974). “Social Security, Induced Retirement, and Aggregate Capital Accumulation.” In: *Journal of Political Economy* 82.5, pp. 905–926.

Galama, Titus, Arie Kapteyn, Raquel Fonseca, and Pierre- Carl Michaud (2013). “A health production model with endogenous retirement.” In: *Health economics* 22.8, pp. 883–902.

Golosov, Mikhail, Ali Shourideh, Maxim Troshkin, and Aleh Tsyvinski (2013). “Optimal Pension Systems with Simple Instruments.” In: *American Economic Review* 103.3, pp. 502–507.

Gregory, Victoria, Guido Menzio, and David G Wiczer (2021). *The Alpha Beta Gamma of the Labor Market*. Working Paper 28663. National Bureau of Economic Research.

Gruber, Jonathan and David A. Wise (2004). “Social Security Programs and Retirement around the World: Micro-Estimation.” University of Chicago Press.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner (2021). “Lifetime Earnings in the United States over Six Decades”. University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2021-60.

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

Haan, Peter and Victoria Prowse (2014). "Longevity, life-cycle behavior and pension reform." In: *Journal of Econometrics* 178.3, pp. 582–601.

Hall, Robert (1982). "The Importance of Lifetime Jobs in the U.S. Economy." In: *American Economic Review* 72.4, pp. 716–724.

Hall, Robert and Marianna 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, Jonathan, Fabrizio Perri, Giovanni L Violante, and Lichen Zhang (2023). *More Unequal We Stand? Inequality Dynamics in the United States 1967–2021*. Working Paper 31486. National Bureau of Economic Research.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017). "Optimal Tax Progressivity: An Analytical Framework." In: *The Quarterly Journal of Economics* 132.4, pp. 1693–1754.

Hubbard, R. Glenn and Kenneth L. Judd (1987). "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax." In: *The American Economic Review* 77.4, pp. 630–646.

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

Hyatt, Henry R. and James R. Spletzer (2013). "The recent decline in employment dynamics." In: *Journal of Labor Economics* 2.5.

— (2016). "The shifting job tenure distribution." In: *Labour Economics* 41, pp. 363–377.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993). "Earnings Losses of Displaced Workers." In: *The American Economic Review* 83.4, pp. 685–709.

Jarosch, Gregor (2023). "Searching for Job Security and the Consequences of Job Loss." In: *Econometrica* 91.3, pp. 903–942.

Jones, Andrew M., Nigel Rice, and Jennifer Roberts (2010). "Sick of work or too sick to work? Evidence on self-reported health shocks and early retirement from the BHPS." In: *Economic Modelling* 27.4, pp. 866–880.

Jung, Philip and Moritz Kuhn (2019). "Earnings Losses and Labor Mobility Over the Life Cycle." In: *Journal of the European Economic Association* 17.3, pp. 678–724.

Keane, Michael P. and Kenneth I. Wolpin (1997). "The Career Decisions of Young Men." In: *Journal of Political Economy* 105.3, pp. 473–522.

Kimball, Miles S. and Mattew D. Shapiro (2003). "The Retirement Elasticity: Theory and Implications." Working paper.

Krueger, Alan B. and Jörn-Steffen Pischke (1992). “The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation.” In: *Journal of Labor Economics* 10.4, pp. 412–437.

Kuhn, Moritz and Gašper Ploj (2020). “Job Stability, Earnings Dynamics, and Life-Cycle Savings.” CESifo Working Paper Series 8710.

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

Michelacci, Claudio and Hernán Ruffo (2015). “Optimal Life Cycle Unemployment Insurance.” In: *American Economic Review* 105.2, pp. 816–859.

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

Morchio, Iacopo (2020). “Work Histories and Lifetime Unemployment.” In: *International Economic Review* 61.1, pp. 321–350.

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

O’Dea, Cormac (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, Roberto and Ludo Visschers (2015). “Unemployment risk and wage differentials.” In: *Journal of Economic Theory* 157, pp. 397–424.

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

Shimer, Robert (2005). “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” In: *American Economic Review* 95.1, pp. 25–49.

Trevisan, Elisabetta and Francesca Zantomio (2016). “The impact of acute health shocks on the labour supply of older workers: Evidence from sixteen European countries.” In: *Labour Economics* 43, pp. 171–185.

9. Appendix

A.1. Logistic distribution and expected utility

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

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

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

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

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

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

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

$$\mathbb{E}V_{\max} = pV_1 + (1 - p)V_2 + \int_{\delta}^{\infty} \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon \tag{5}$$

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

$$\int_{\delta}^{\infty} \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon = \delta \cdot \frac{\exp\left(-\frac{\delta - \mu}{\sigma}\right)}{1 + \exp\left(-\frac{\delta - \mu}{\sigma}\right)} + \sigma \cdot \log\left(1 + \exp\left(-\frac{\delta - \mu}{\sigma}\right)\right). \tag{6}$$

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

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

I use these expressions to simplify Equation (6) 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 (5) 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 (1) in Section 3.1.

A.2. Model solution

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

A.3. Human capital

In this section, I discuss an alternative specification for the human capital accumulation process. In particular, I consider a human capital accumulation process that corresponds to the Ben-Porath (1967) model. While I leave all other parts of the model unchanged,

Figure A.1. Earnings variance.

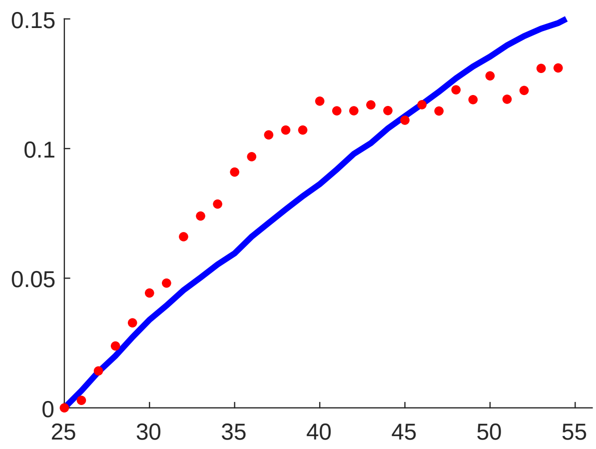

Notes: This figure shows the earnings variance in a model with Ben-Porath (1967) human capital accumulation.

I assume that human capital investment entails a monetary cost as a function of the investment time instead of a utility cost. The budget constraint of a worker is then

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

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 A.1 shows that compared to the baseline model, the increase in earnings variance over the life cycle is slightly overstated as it increases almost linearly in age. The interaction between job 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.

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

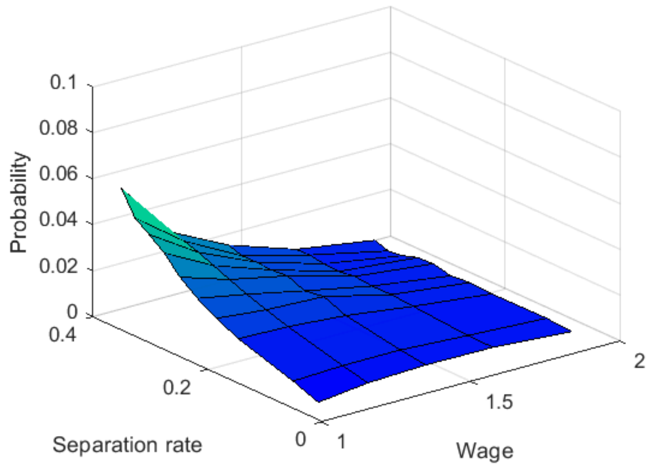

Notes: This figure shows the joint distribution of wages and separation rates.

A.4. Calibration

A.4.1. Job-offer distribution

The number of gridpoints for wages is set to $N_w = 5$ where $\underline{w} = 1$ and $\overline{w} = 1.85$. All gridpoints are equidistant in logs. Concerning the job-separation probabilities, it is assumed that $N_\lambda = 10$ with $\underline{\lambda} = 0.006$ and $\overline{\lambda} = 0.35$. All remaining gridpoints are located non-linearly between these two values. The marginal distributions for wages and job 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 job 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 job stability is pinned down by the parameter θ. Table A.1 presents the estimated parameters.

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

A.4.2. Parameters

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

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

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

Table A.1. Estimated parameters

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

A.4.3. Approximation of average lifetime earnings

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

A.5. 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.3.

Figure A.3a 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

Figure A.3. Welfare change as a function of the initial state variables of separation rate, wealth, and human capital.

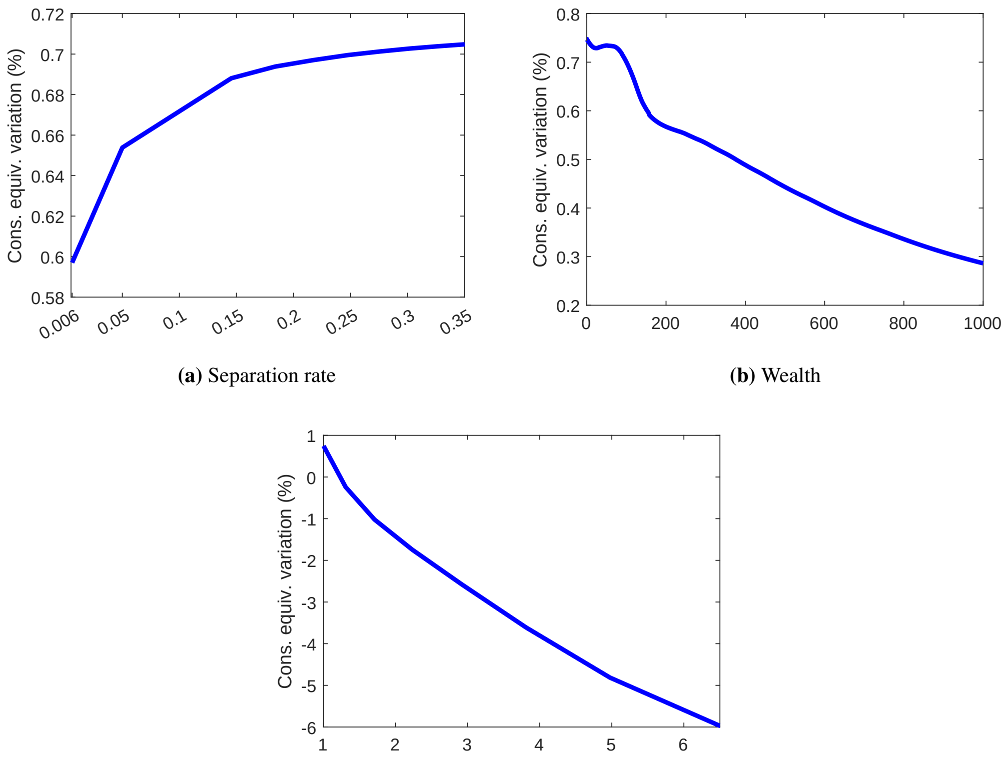

(a) Separation rate

(b) Wealth

(c) Human capital

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 job stability.

Figure A.3b presents the welfare change as a function of the wealth level. Interestingly, the function is not monotonic: the welfare change exhibits a hump close to zero wealth, but then, decreases steadily in the asset level. Starting from zero wealth, higher wealth allows for a better ability of consumption smoothing in the presence of the payroll tax rate. Because young workers have relatively low labor earnings and they save a large proportion of their earnings for precautionary reasons, an increase in wealth leads to higher welfare gains. However, as the amount of wealth increases further, the total amount of lifetime resources available increases in a way that the insurance effect of higher pension progressivity becomes less important. Still, the welfare change from the optimal progressive pension system remains positive since prospective labor market outcome remains uncertain and pension progressivity provides insurance against this risk.

Finally, Figure A.3c 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

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

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

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

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, contribution to the pension system increases relative to benefits. As a consequence, the increase in pension progressivity decreases welfare.

A.6. 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.2 summarizes the results. One striking result is that the optimal degree of progressivity does not change for the pension system under alternative unemployment insurance systems. In all cases, it is optimal to remove the Social Security cap and the shape parameter γ of the pension benefit formula takes a value of 0.9. Note, however, that the induced welfare gains are decreasing in the replacement rate of the unemployment insurance system. As workers with frequent job interruptions get a higher replacement rate, the welfare gain from redistribution through the pension system decreases. However, whereas a higher replacement rate of the unemployment insurance system helps workers to smooth consumption upon job interruptions, it does not compensate for the potential loss of human capital. Hence, the optimal pension progressivity remains unaffected by the unemployment insurance system.

Table A.3. Optimal pension systems in the economy with human capital depreciation.

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

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

A.7. Human capital depreciation

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

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

Table A.3 summarizes the optimal policy parameters. Similar to the baseline model, removing the Social Security cap is optimal and the shape parameter γ increases from 0.44 to 0.64 (with cap) and 0.86 (no cap) so that the progressivity level increases under the optimal policy. The inclusion of human capital depreciation does not change the previous results qualitatively.

A.8. Shift in job-stability distribution

Figure A.4 displays the average separation rate by age in the baseline economy and the economy with higher job stability for workers in the bottom and top quartile of the average job-stability distribution over the life cycle. Relative to the average separation rate in the baseline economy (solid line), the dashed-dotted line shows that the life-cycle profile of average job-separation rate of workers with most unstable jobs shifts downwards in the economy with higher job stability. At labor market entry, the average separation rate decreases by 3 percentage points for workers with most unstable jobs. This gap

Figure A.4. Average life-cycle profile of separation rate for the baseline economy and for the economy with higher job stability.

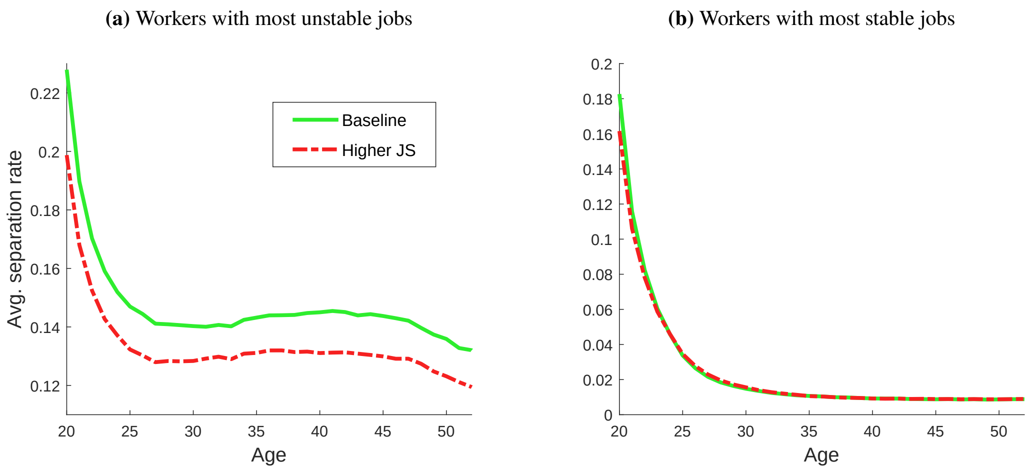

Notes: The left panel shows the average separation rate of workers with most unstable jobs (bottom 25% of job-stability distribution) and the right panel the average separation rate of workers with most stable jobs (top 25% of job-stability distribution).

declines over the life cycle. A shift in the lower tail of the job-stability distribution primarily affects the average separation rate of young workers because young workers also accept unstable jobs: Employment offers the opportunity to invest in human capital and to increase prospective earnings growth, and is therefore highly valuable for young workers. Over time, workers climb the job ladder and an average worker finds more stable jobs as the worker spends more time in the labor market. Note that the separation rates of the most stable jobs remain unchanged in the recalibrated economy. By construction, workers with most stable jobs are little affected by the shift in the job-stability distribution and their profile of average separation rate remains almost unchanged.