Optimal Progressive Pension Systems in a Life-Cycle Model with Heterogeneity in Job Stability*

Leanne Nam†

April 2022

Abstract

Interrupted work histories and earnings losses following job displacements lower pension entitlements and impair workers' ability to accumulate life-cycle savings. Job stability is therefore key for inequality in lifetime earnings and for welfare of individuals both before and after retirement. In this paper, I study how a progressive pension system should take heterogeneity in job stability into account and quantify potential welfare gains from implementing the optimal pension system. Pension progressivity provides insurance against the risk of unstable employment histories and counteracts 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 welfare gains of 0.52% of lifetime consumption for labor market entrants from increasing the degree of pension progressivity relative to the existing U.S. pension system. Following a macroeconomic shift in the job-stability distribution towards higher job stability, the optimal pension system implies a higher degree of redistribution. In the presence of liquidity constraints, an increase in job stability makes the payroll tax finance of pension systems less costly and enhances the insurance effect of progressive pension systems.

JEL classification: E24, H21, H55, J64

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

*I thank Moritz Kuhn for invaluable advice, guidance, and encouragement throughout this project. I also thank Jonathan Heathcote, Marco Bassetto, Camille Landais, Thomas Hintermaier, and Eugenia Gonzalez-Aguado as well as seminar participants of the RTG 2281 Brown Bag seminar 2021 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

Interrupted work histories and earnings losses following job displacements lower pension entitlements and impair workers' ability to accumulate life-cycle savings. Job stability is therefore key for inequality in lifetime earnings and for welfare of individuals both before and after retirement. In light of the wide range of studies showing that heterogeneity in job stability is a salient feature of labor markets (see, for example, Hall, 1982, Kuhn and Ploj, 2020, and Morchio, 2020), it is important to study how heterogeneity in job stability affects the optimal design of progressive pension systems. Recent shifts in the distribution of job stability caused by the Covid-19 pandemic further enhance the importance of studying the policy implications of heterogeneity in job stability and its changing shape for the post-pandemic economy. Although the degree of job stability of a worker during working life remains unobservable to the government, the employment history at the entry to retirement allows to infer the worker's job stability throughout working life. The question is if and how pension systems should consider such observable inequality in employment history in determining individual pension entitlements.

In this paper, I show that progressive pension systems can effectively counteract consumption inequality produced by heterogeneity in job stability and quantify potential welfare gains from implementing the optimal pension system. Even though a worker's degree of job stability is not directly observable, pension systems redistribute earnings based on labor market histories and average lifetime earnings of workers, allowing for reallocation of resources to workers who suffer from frequent career interruptions. Using a calibrated life-cycle model with heterogeneity in job stability and consumption-saving decision, 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 increasing the degree of pension progressivity achieves welfare gains of 0.52% of lifetime consumption for labor market entrants relative to the existing U.S. pension system. Following a macroeconomic shift in the job stability distribution towards higher job stability, the optimal pension system implies a higher degree of redistribution.

The life-cycle model in this paper features a frictional labor market where agents search on and off the job, making consumption-saving decisions in incomplete financial markets and investing in risky human capital. The model incorporates endogenous retirement and a pension system calibrated to the U.S. social security system. Workers are heterogeneous with respect to their job-separation rates at each point in time of their life cycle. The novel feature of this model is the incorporation of heterogeneity in job stability while treating human capital accumulation and retirement as endogenous processes. Endogenous human capital accumulation is an important determinant of earnings (Becker, 1964 and Ben-

Porath, 1967) and is key to explain increasing life-cycle earnings inequality (Hubmer, 2018 and Jung and Kuhn, 2018). Therefore, along with endogenous retirement, endogenous human capital accumulation is an important dimension in analyzing optimal pension systems.

This study focuses on a classical trade-off between insurance and incentives for the policy maker. The redistributive effect of pension progressivity provides insurance against unstable employment histories and low lifetime earnings. However, this comes at the cost of reducing the returns on human capital investment and thus discouraging investment of workers. The progressivity level affects retirement incentives of workers as well, since it determines the amount of pension wealth of each worker. Moreover, the payroll tax finance plays a key role 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. Increasing the provision of pension coverage can therefore lead to utility losses for these workers, even though the policy change would increase their total amount of lifetime resources.

To develop an intuitive idea how pension systems provide insurance against the risk of unstable employment history produced by heterogeneity in job stability, I develop a two-stage model. In this model, workers are heterogeneous with respect to their job stability and make consumption-saving as well as human capital investment decisions. I compare and contrast a progressive pension system to annual earnings taxation. Progressive pension benefits take the full earnings history of an individual at the entry into retirement into account, whereas under annual earnings taxation, taxes and transfers are determined separately in each period. Analyzing the welfare effects for each of these policies, I find that in the presence of heterogeneity in job stability, a progressive pension system achieves higher ex-ante welfare. The reason is that annual earnings taxation does not capture persistent differences in the employment history and therefore fails to reduce inequality in lifetime earnings. On the contrary, a progressive pension system, which is a specific type of lifetime earnings taxation, allows to tackle the earnings inequality more directly. This result motivates the study of optimal pension systems to address earnings inequality produced by heterogeneity in job stability since pension systems are the most commonly applied form of lifetime earnings taxation in practice.

Using a life-cycle model with heterogeneity in job stability, human capital investment, and consumption-saving decision that is calibrated to the U.S. economy, I analyze the optimal pension system. 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 which is, in the absence of a progressive pension system, directly transmitted into the life after retirement. An increase in pension progressivity induces a welfare gain of 0.52% in terms of lifetime consumption. On the one hand, higher progressivity of the pension system decreases consumption inequality over the life cycle and offers insurance to workers who suffer from frequent career interruptions and low lifetime earnings. In particular, the optimal pension system reduces the increase in life-cycle consumption variance by almost one third of the initial increase. On the other hand, the increase in pension progressivity distorts human capital investment and retirement decisions. The average stock of human capital decreases by almost 2% and the average retirement age decreases by 0.5% 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 growth which is similar to the results of Michelacci and Ruffo (2015). 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.

In the last part of this paper, I analyze the consequences of a shift in the job-stability distribution of 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. Even though a shift towards higher job stability decreases inequality in labor market outcomes, I find that the optimal pension system becomes more redistributive in an economy with more stable jobs. This is because a high degree of redistribution through the provision of pension coverage requires a high payroll tax rate which is 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. Importantly, 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 weakens the negative impact of a borrowing constraint on consumption smoothing of workers. As a result, the optimal progressive pension system implies a higher degree of redistribution for the economy with higher job stability.

The remainder of this paper is structured as follows. Section 2 relates this paper to the existing literature. Section 3 studies the two-stage model. I present the life-cycle model in Section 4, followed by the baseline calibration in Section 5. Section 6 analyzes the effects of heterogeneity in job stability on inequality in labor market outcomes and its life-cycle consequences. I analyze the ex-ante optimal pension system for the baseline economy in Section 7. Section 8 explores the consequences of a shift in job-stability distribution in the U.S. labor market on the optimal design of pension systems. Section 9 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 several recent studies that document empirical evidence for such heterogeneity. The literature consistently finds that job-stability heterogeneity is a pronounced characteristic of labor markets. Among others, Jung and Kuhn (2018) find considerable heterogeneity in job stability for the United States and show that it is key to explain persistent earnings losses after job displacement. 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.

Second, this paper relates to the literature on inequality in lifetime earnings. In a simple theoretical framework, I show the consequences of heterogeneity in job stability on lifetime earnings. This model delivers an intuitive explanation why annual earnings taxation fails to address inequality in lifetime earnings. 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) point out that inequality in annual earnings may evolve differently from inequality in lifetime earnings. They 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) focus on the policy implications of inequality in lifetime earnings. They show that progressive annual earnings taxation is not able to counteract inequality in lifetime earnings produced by differences in employment history. Taxing lifetime earnings is more effective since it directly addresses the inequality. Additionally, the empirical findings of their paper indicate that heterogeneity in labor market outcome is a major source of inequality in lifetime earnings for workers with same abilities. Putting all these together, their result lends further support to the importance of lifetime earnings taxation in addressing earnings inequality produced by heterogeneity in job stability, which is an essential idea of this paper.

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. Some other prominent studies on pension systems are, for example, Krueger and Pischke (1992); Coile and Gruber (2001); Cremer et al. (2004); Gruber and Wise (2004); Bloom et al. (2007); and Haan and Prowse (2014). Past studies on optimal pension systems, however, neglect the fact that labor markets feature large heterogeneity in job stability. I add to this literature by considering heterogeneity in job stability as an important source of wealth and consumption inequality in analyzing the design of optimal pension systems.

3. Stylized two-stage model

Before setting up the life-cycle model, this section applies a stylized two-stage model with heterogeneity in job stability. This model illustrates that lifetime earnings taxation addresses earnings inequality produced by heterogeneity in job stability more effectively than annual earnings taxation, leading to higher ex-ante welfare. I show that annual earnings taxation is not able to capture persistent differences in the employment history and does not mitigate inequality in lifetime earnings.

Figure 1. Stylized two-stage model.

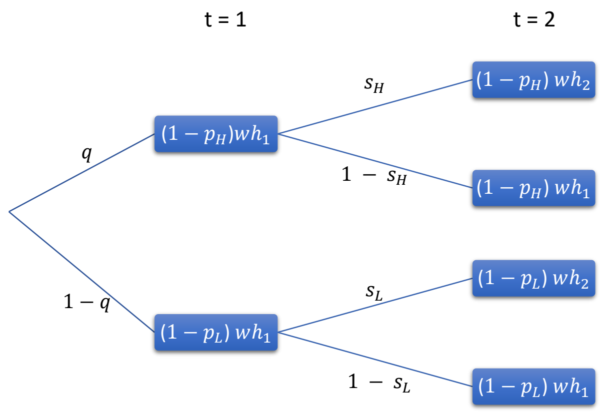

Notes: The graph shows the structure of the two-stage model. $t = 1$ represents the first and $t = 2$ the second half of working life.

3.1. Model

Figure 1 provides an overview on the structure of the two-stage model. The model is populated by a unit mass of agents consisting of two types. A fraction q of agents is denoted by type H and the remaining fraction $1 - q$ of agents is denoted by type L. An agent's life is divided into two stages where each stage comprises the first ($t = 1$) and the second half of working life after labor market entry ($t = 2$). In both stages, agents face a type-specific risk of unemployment p_H and p_L, where $p_H < p_L$. Assuming that time is continuous at each stage, agents spend a fixed fraction of their lifetime in unemployment depending on the type-specific risk of unemployment. On average, type H (type L) is unemployed for a fraction p_H (p_L) of lifetime. The average unemployment duration is larger for agents of type L since they face a permanently higher probability of job loss. Agents are indexed by $i \in [0, 1]$. Ex ante, agents are identical and they learn their types in the first period after labor market entry. In the first stage of life, both types of agents earn labor earnings wh_1 in each period when they are employed and decide at the end of the stage how much to invest in human capital. Human capital investment increases the probability of obtaining a higher level of human capital $h_2 > h_1$ in the second period from 0 to s_H (s_L), but the investment entails a disutility of $c_H \cdot s_H^3$ ($c_L \cdot s_L^3$). Human capital investment is more costly in terms of utility for agents L than for agents H such that $c_L > c_H$. This assumption of a higher utility cost for type L reflects that it is more difficult to make human capital investment for a worker with frequent career interruptions. Successful human capital accumulation increases labor earnings from wh_1 to wh_2 in employed years of the second stage. If human capital accumulation is unsuccessful, labor earnings remain unchanged at wh_1. Because agents get zero labor earnings during unemployment, the average labor earnings in each stage are $(1-p_H)wh$ for type H and $(1-p_L)wh$ for type L, where $h \in \{h_1, h_2\}$.

3.2. Tax policy implications of job-stability heterogeneity

In this model, there are two sources of earnings risk: first, the stochastic realization of human capital investment, and second, the risk of having unstable jobs over life which lead to lower lifetime earnings. The latter risk constitutes an exogenous shock to lifetime earnings and opens up scope for policy intervention. In particular, a welfare gain can be achieved in this economy by redistributing labor earnings from type H to type L. This provides insurance against the hazard of having a high employment risk.

Annual earnings taxation, however, does not achieve sufficient redistribution between different types of agents. Taxes and transfers cannot be conditioned on types explicitly as they are not directly observable by the tax authority. Moreover, the annual employment status and annual labor earnings provide little information on the underlying type of an agent. Note that agents with stable jobs do not have an incentive to report their type truthfully as a truthful report would result in lower net earnings for these agents.

This problem can be resolved by implementing lifetime earnings taxation in place of annual earnings taxation. Lifetime earnings are indicative of the employment history of agents, making types ex-post observable. In this way, the tax authority is able to capture inequality in lifetime earnings and redistribute resources from type H to type L. Note that, if there is no heterogeneity in job stability and all agents are confronted with the same level of employment risk, lifetime earnings taxation achieves no additional welfare gain in this model.

One important remark is that unemployment insurance is not an appropriate instrument to insure agents against the risk of low job stability. Recall that agents spend a fixed fraction of their lifetime in unemployment in this model. Within a stage, agents perfectly smooth out consumption which implies that they are already insured against the idiosyncratic risk of unemployment. This model thus allows to focus on the policy question of how to capture persistent differences in earnings risk and provide an optimal degree of redistribution.

3.3. Tax function and further model specifications

In this subsection, I establish a general framework for the tax system. Suppose that a utilitarian social planner levies taxes and pays out transfers so as to insure agents against bad earnings shocks. More specifically, assume that net earnings after taxes and transfers are described by the parametric specification

$$y^*_{i,t} = \phi y_{i,t}^{1-\gamma}$$

where $y_{i,t}$ denotes gross earnings and $y^*_{i,t}$ net earnings after taxes and transfers. The parameter ϕ scales the level of net earnings such that the resource constraint of the economy

$$\int_i \sum_t y^*_{i,t} di = \int_i \sum_t y_{i,t} di \tag{1}$$

is satisfied. The parameter γ determines the progressivity of the tax system. If $\gamma = 0$, net earnings linearly increase in gross earnings. The case of $\gamma > 0$ corresponds to a progressive tax system and if $\gamma = 1$, all agents receive the same level of net earnings. The social planner chooses the parameter γ to maximize ex-ante welfare of agents in the economy.

Agents derive logarithmic utility from consumption and can save and borrow freely in order to shift consumption across time provided that the following lifetime budget constraint holds:

$$\sum_t c_{i,t} = \sum_t y^*_{i,t} \quad \forall i \in [0, 1]$$

For the sake of simplicity, assume that the time preference parameter is set to $\beta = 1$ and the risk-free interest rate on assets to $r = 0$.[1] Moreover, for the numerical analysis hereinafter, I set $w = 1$, $h_1 = 1$, $h_2 = 1.5$, $c_H = 0.2$, $c_L = 0.8$, $p_L = 0.7$, and $p_H = 0$.

In the following, I consider two types of tax policies: In the first case, the social planner only observes labor earnings in each period and conditions taxes and transfers on these earnings. In the second case, the tax system takes into account the total labor earnings earned in both stages of life and redistributes earnings at the second stage. For each of these cases, I investigate the welfare-maximizing progressivity level of the tax system as a function of the parameter q, which denotes the share of agents with stable employment.

To solve the model and search for the optimal policies, I discretize the choice set for assets and human capital investment as well as the progressivity level of the tax system and the share of type H in the economy. The consumption-saving and the human capital investment decisions are solved by applying on-grid search. To search for the optimal policies, I proceed as follows. For all values of $\gamma \in [0, 1]$, I numerically search for the parameter value of ϕ that satisfies Equation (1). The optimal policy is given by the progressivity level γ that achieves the highest ex-ante expected lifetime utility.

3.4. Optimal tax systems

Figure 2 and Figure 3 display the optimal tax progressivity and the corresponding human capital investment by type as functions of the share of agents with stable jobs which

[1] The results can be generalized to $\beta(1+r) \neq 1$ and to assuming a general case of CRRA utility function.

Figure 2. Optimal progressivity level for annual earnings taxation and human capital investment.

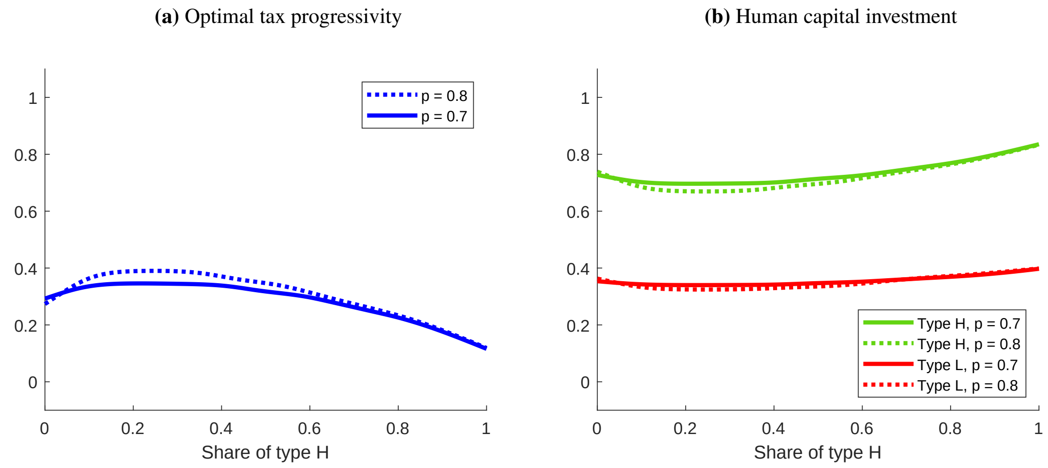

Notes: The left panel shows the optimal progressivity level of annual earnings taxation as a function of the share of agents with stable jobs in the economy. The right panel shows the associated human capital investment by type. In all plots the solid line represents the baseline economy and the dashed line represents the economy with higher job loss probability.

is represented by the parameter q. For annual earnings taxation, Figure 2a shows that the optimal progressivity level γ is strictly positive and slightly concave in q. The optimal progressivity level achieves its highest value at around $q = 0.25$ with $\gamma = 0.35$. Towards the extreme cases in which the economy is populated by a single type of agents, optimal progressivity decreases to 0.3 for $q = 0$ and to 0.13 for $q = 1$. Overall, Figure 2b shows that human capital investment follows the inverse pattern of the optimal tax progressivity as a function of q. Note that type L provides less effort for human capital accumulation for all values of q. This is because human capital investment is, by assumption, more costly for type L in terms of utility. Additionally, due to frequent career interruptions, the return on human capital investment is lower for type L than for type H with stable employment.

The result that the optimal progressivity is strictly positive and a fairly flat function of the share of agents with stable employment indicates that annual earnings taxation provides insurance against low earnings to some extent, but fails to fully account for the existing inequality in lifetime earnings and to redistribute earnings from type H to type L. More specifically, progressive annual earnings taxation reallocates labor earnings from agents whose stock of human capital has increased in the second stage to the first stage of life and to agents in the second stage with unsuccessful human capital investment. However, the tax authority is unable to sufficiently redistribute earnings from type H to type L since observing annual earnings at each point in time does not provide information on agents' types.

The reason why the optimal tax progressivity increases as the share of type H in the economy grows from zero to a quarter is that the aggregate unemployment duration of all agents declines such that the total amount of available resources in the economy

Figure 3. Optimal progressivity level for lifetime earnings taxation and human capital investment.

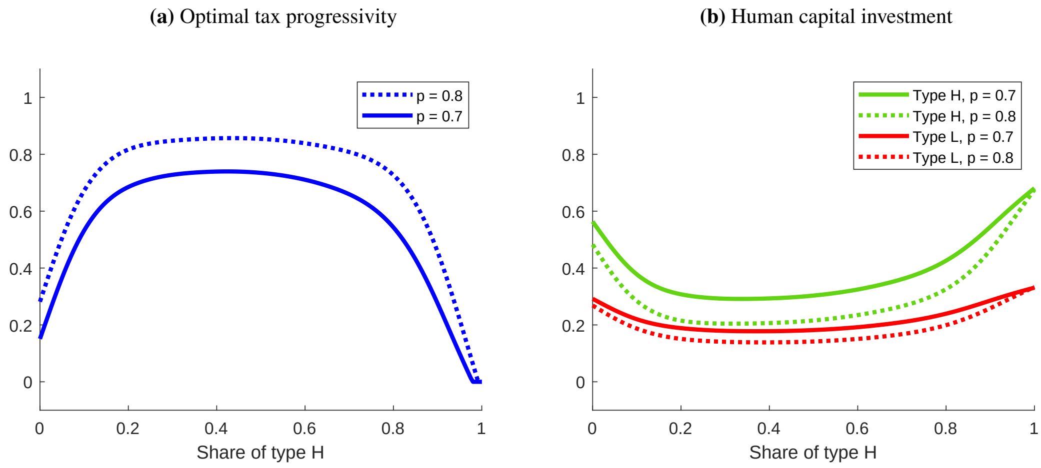

Notes: The left panel shows the optimal progressivity level of lifetime earnings taxation as a function of the share of agents with stable jobs in the economy. The right panel shows the associated human capital investment by type. In all plots the solid line represents the baseline economy and the dashed line represents the economy with higher job loss probability.

increases. Consequently, the tax system is able to allocate more resources to agents with low earnings through a raise in the progressivity level. This increases the expected lifetime utility of agents with low job stability whose marginal propensity to consume is particularly high in states of no additional human capital. In response to this policy change, human capital investment slightly decreases for both types in Figure 2b as the return on human capital investment declines in the progressivity level. The decrease in human capital investment partly offsets the insurance effect from higher progressivity. However, as long as the population share of type L constitutes a substantial part of the economy, the welfare weight on these agents remains large such that a higher progressivity level achieves the highest aggregate welfare. With a further increase in the share of type H, the optimal tax system implies a decline in the progressivity level. Since agents of type H have stable employment paths and thus higher labor earnings in both stages of life, their marginal propensity to consume is lower than for type L. Therefore, their utility gain from redistributing resources from high to low-income states is not sufficient enough to surpass the distortions on human capital investment. As a consequence, the optimal level of progressivity falls to 0.13 at $q = 1$.

In a next step, I investigate the optimal tax system for the case where taxes and transfers are conditional on lifetime earnings which correspond to the sum of labor earnings in both stages of life. The results are displayed in Figure 3. Figure 3a indicates that the optimal tax progressivity strongly varies in the composition of the economy by types. In particular, the optimal progressivity has a pronounced concave shape in q compared to the case of annual earnings taxation. If types are distributed almost equally, a high degree of tax progressivity is optimal. Human capital investment plotted in Figure 3b again follows the inverse pattern of the optimal progressivity level.

Optimal progressivity strongly depends on the distribution of types because lifetime earnings taxation allows to redistribute lifetime earnings from type H to type L, providing the latter group of agents insurance against the permanent risk of unemployment. The optimal tax progressivity declines in the share of type H due to the following reasons. Firstly, the Pareto-weight on type L decreases since its population share becomes smaller and thus the lifetime utility of type H gains in importance in the social welfare function. Because a high level of progressivity is welfare-detrimental for type H and the insurance effect of the tax progressivity becomes small, the optimal progressivity level decreases in q. Secondly, as the share of type H increases, the total return on human capital investment grows in absolute terms and this increases the total amount of resources in the economy. As a consequence, the welfare gain from redistributing earnings across the two types of agents decreases in the share of type H, and it becomes optimal to provide higher incentives for human capital investment by lowering the level of tax progressivity. Also, as the share of type L increases, the optimal progressivity level decreases and attains a value of 0.15 for $q = 0$. If the share of type L increases, the total available resources in the economy declines because human capital investments and the average level of lifetime earnings are lower for type L than for type H. Hence, there are little resources available to redistribute across agents. Moreover, redistribution among the same type does not lead to a large welfare gain since, as in the case of annual earnings taxation, agents are already insured against the idiosyncratic risk of unemployment. Providing higher incentives for human capital investment to type L becomes more efficient, and this is achieved by lowering the progressivity level of the tax policy.

In order to analyze the welfare consequences of the two different tax policies, I compute the ex-ante expected utility of agents before their labor market entry. Figure 4 shows that lifetime earnings taxation achieves a higher aggregate welfare than annual earnings taxation for $0 < q < 1$. In contrast to annual earnings taxation, taxing lifetime earnings captures persistent differences in job stability and the resulting inequality in lifetime earnings. Lifetime earnings taxation provides insurance against the exogenous shock of having an unstable employment path by redistributing earnings from type H to type L at the end of the second stage of life, whereas annual earnings taxation achieves little redistribution. This is because the insurance against low-earnings states under annual earnings taxation does not sufficiently capture the fact that type L has lower lifetime earnings due to higher risk of job loss. The concave shape of expected lifetime utility as a function of q for lifetime earnings taxation reflects this property. If the distribution of types in the economy is balanced and hence q is close to 0.5, the redistribution between the two groups of workers induces a large welfare gain for lifetime earnings taxation. Hence, the gap between the dotted and the solid line is substantial. Towards the extreme cases where the economy is

Figure 4. Ex-ante welfare under annual earnings taxation and lifetime earnings taxation.

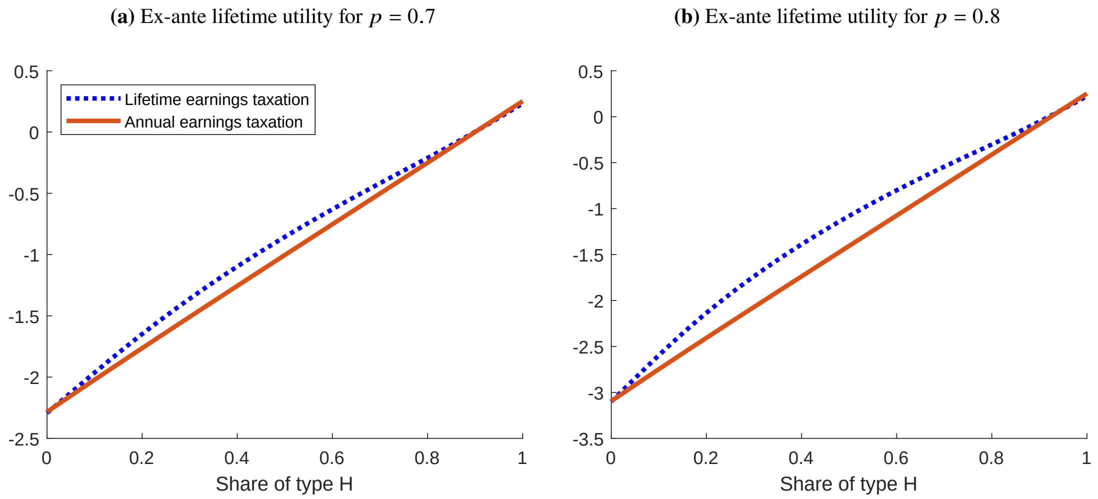

Notes: The left panel shows the ex-ante lifetime utility as a function of the share of agents with stable jobs in the economy under lifetime earnings taxation (dashed line) and under annual earnings taxation (solid line) for the baseline economy. The right panel considers the economy with higher job loss probability.

populated just by one type, there is no redistribution across different types of agents and taxing lifetime earnings does not lead to an additional welfare gain.

Next, I analyze the implications of a shift in the job-stability heterogeneity on the optimal tax policy. More precisely, I investigate the case where agents of type L face a higher employment risk and hence spend a larger fraction of lifetime in unemployment than in the initial economy. Assuming that all other economic conditions remain unchanged, I analyze the optimal policies and their welfare effects in response to an increase in the parameter p.

For annual earnings taxation, Figure 2a indicates that the optimal progressivity level increases slightly in the economy with higher employment risk. At $q = 0.3$ the optimal progressivity increases by approximately 0.45 percentage points. By contrast, the increase in progressivity is much stronger for lifetime earnings taxation. The dotted graph in Figure 3a shows that the optimal progressivity level increases everywhere as p rises from 0.7 to 0.8; for $q = 0.3$, the optimal level of progressivity increases by 12 percentage points and leads to a decrease in the optimal human capital investment for all agents. The increase in the employment risk of type L raises heterogeneity in lifetime earnings between agents of type L and type H. Since annual earnings taxation does not account for this heterogeneity between the two groups of agents, the adjustment of the optimal progressivity is not as strong as for lifetime earnings taxation. As already discussed, annual earnings taxation does not explicitly account for differences in the employment risk, whereas lifetime earnings taxation captures this heterogeneity. In Figure 4b, the additional welfare from implementing lifetime earnings taxation instead of annual earnings taxation becomes even larger in the new economy with larger employment risk compared to the additional welfare in Figure 4a. This result indicates that the gain from taxing lifetime earnings grows in heterogeneity in job stability because lifetime earnings taxation mitigates earnings inequality more effectively than annual earnings taxation. In particular, the insurance effect of lifetime earnings taxation is larger at a given level of tax progressivity due to its redistributive effect. Hence, the gain from providing higher insurance through higher tax progressivity is larger for lifetime earnings taxation, which produces the differences in the adjustment of the tax progressivity for both types of taxation.

These results motivate the idea of studying optimal progressive pension systems in a richer model to address earnings inequality produced by heterogeneity in job stability. Pension systems are the most commonly applied form of lifetime earnings taxation in practice as pension systems redistribute earnings based on the labor market history and the average lifetime earnings of workers. The next section describes an extensive life-cycle model which I use to analyze optimal pension systems in the remainder of this paper.

4. Life-cycle model

The remainder of this paper applies a life-cycle model with heterogeneity in job stability which combines consumption-saving and labor market behavior. This framework builds on the life-cycle model of Kuhn and Ploj (2020). 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 last wage 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 $y_r(w_j, h_j, n) = \omega(\bar{h})$ where $\bar{h}$ is the final level of human capital agents achieve in the working phase before retirement. The function ω determines the level of benefits assigned to an agent with a final human capital level $\bar{h}$. The next subsection provides a detailed explanation for the mapping from the final human capital to retirement benefits and the shape of the function ω. Retirement benefits remain constant during the retirement phase and therefore, agents face no income risk during retirement. 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 shift to the investment stage at which they make their human capital investment decisions. In case of 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. For a given level of effort t, the realization of human capital investment is stochastic. More specifically, assuming that human capital levels are discrete and that h^{+} denotes the immediate successor and h^{-} the immediate predecessor of h, the law of motion for human capital is

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

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

4.1. 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

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

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

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

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 the lifetime utility from retirement is larger than the expected utility of remaining in the labor force, agents decide to retire. Hence, agents face the following discrete choice problem

$$V_{\max}(a, w, \lambda, h, j) = \max\{V(a, w, \lambda, h, j), V^r(a, w, h, j_r) + \varepsilon\}$$

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

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

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

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

4.2. 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

$$\mathbb{E}\left[\sum_{t=t^R}^{T} \frac{\omega(\bar{h})}{(1+r)^{t-1}}\right] = \mathbb{E}\left[\sum_{t=1}^{t^R} \frac{y(a,w,\lambda,h,t)}{(1+r)^{t-1}} \cdot \tau\right] \tag{3}$$

must be satisfied. The parameter τ in Equation (3) 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.

5. 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 66, workers make retirement decisions. At age 67, all workers who decided to remain in the labor force up to this age are shifted to the retirement phase.

Wages and job-separation probabilities are discretized on grids with $[\underline{w}, \overline{w}]$ and $[\underline{\lambda}, \bar{\lambda}]$, respectively. The lowest quarterly separation probability is $\underline{\lambda} = 0.006$ which corresponds to lifetime jobs with an expected job duration of 42 years. The highest separation probability is set to $\bar{\lambda} = 0.35$. The job-offer distribution consists of a joint distribution of job-separation probability and wage. See Section A.3.1 in the Appendix for a detailed description of the job-offer distribution.

5.1. U.S. pension system and retirement elasticity

In order to investigate the U.S. pension system, I consider the U.S. social security legislation for 2019. Based on this, I set the social-security income 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

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

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.

Then, I convert this cap and these bendpoints into model units. For this purpose, I match the gross pension replacement rates by individual earnings in the model to the observed net pension replacement rates in the United States reported by OECD (2019). I target a net replacement rate of 49.4% for a worker with average lifetime earnings in the economy.

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

To study the optimal level of progressivity of the retirement benefit system, I restrict attention to the parametric class of retirement systems given by

$$\omega(\bar{y}) = \phi \cdot (\bar{y})^{1-\gamma}. \tag{5}$$

Here, the parameter γ governs the progressivity level of the benefit system, with $\gamma = 0$ representing a system with flat replacement rate (benefits linearly increase in lifetime earnings) and $\gamma > 0$ corresponding to a progressive pension system. This formulation of the pension system implies that an increase in the parameter γ makes the pension system more redistributive.

I approximate the retirement benefits formula in Equation (4) by this two-parameter function which implies that retirement benefits are a (weakly) concave function of average lifetime earnings. 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). I estimate the parameters in the benefit function by fitting the function

$$\log \omega(\bar{y}) = \log \phi + (1 - \gamma) \log \bar{y}$$

to the U.S. social security system specified in Equation (4). The estimated parameters are $\phi = 0.627$ and $\gamma = 0.316$. If the average lifetime earnings exceeds the social-security income cap, I assume that retirement benefits are fixed to

$$\omega(\bar{y}) = 0.9bp_1 + 0.32(bp_2 - bp_1) + 0.15(cap - bp_2)$$

for $\bar{y} >$ cap as in Equation (4). In the Appendix, Figure A.1 compares the U.S. pension system and the approximating function.

In the model, I approximate the average lifetime earnings of workers by the level of human capital workers achieve before they are shifted to the retirement phase. This final level of human capital $\bar{h}$ constitutes a good proxy for an individual's earnings history because human capital accumulation is only possible in employment. Thus, the final level of human capital is informative about a worker's average labor earnings, the past duration of employment, and the degree of job stability the worker faces before retirement. I regress the average lifetime earnings of workers on cubic polynomials of the final level of human capital of each worker. Section A.3.4 in the Appendix discusses the accuracy of this approximation of average lifetime earnings.

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 $\mu = 13.5$ 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 also be interpreted as deriving an unexpected positive utility from retirement such as health problems, which are an important cause for retirement (see, for example, Dwyer and Mitchell, 1999; Disney 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.3 in the Appendix provides further details on the estimation procedure and Table A.1 summarizes the estimated parameters.

5.2. Retirement age distribution and retirement elasticity

Figure 5 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 5 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 -0.07 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. For the case that the policy change already takes place at labor market entry, the model implies an elasticity of retirement relative to a 1% reduction in retirement benefits of -0.0424.

These changes in retirement decisions fall into the range of retirement elasticities

Figure 5. Retirement age distribution in the data and in the baseline model.

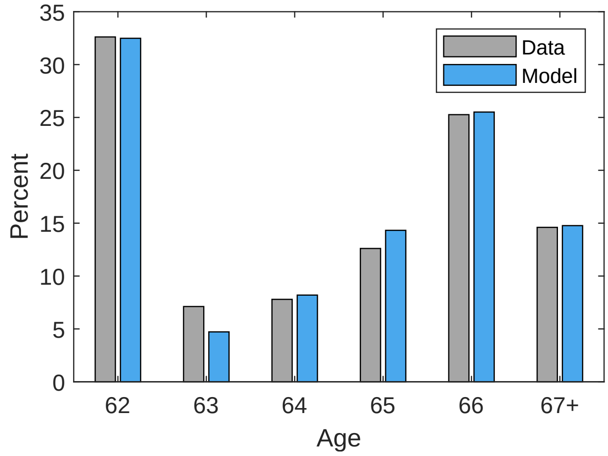

Notes: The data source is Table 6.A4 of Social Security Administration (2020), Annual Statistical Supplement, available at https://www.ssa.gov/policy/docs/statcomps/supplement/.

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. The response of retirement age to changes in the generosity of retirement benefits is also close to the results of French (2005). Based on a structural model, French (2005) conducts several policy experiments and shows that a reduction in benefits by 20% leads to an increase in years worked from 32.6 to 33, which implies a retirement elasticity of -0.06.

6. Heterogeneity in job stability and life-cycle dynamics

The calibrated model is used as the baseline for the analysis in the remainder of this paper. 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

Figure 6. Average life-cycle profile of job-separation rate.

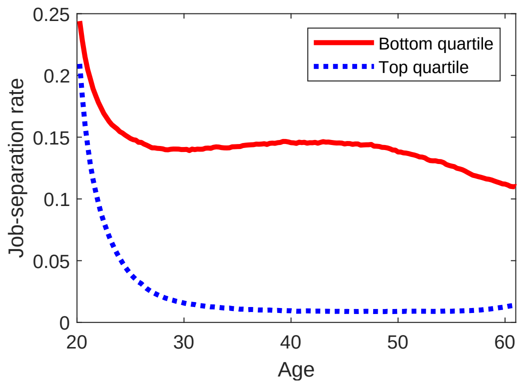

Notes: The solid line displays the average life-cycle profile of job-separation rate of workers with most unstable jobs. The dashed line represents workers with most stable jobs throughout working life.

the baseline calibration of the model. Because an optimal progressive pension system takes the entire earnings history of an agent into account to efficiently redistribute resources, it is important to learn how job stability over the whole life cycle affects consumption and earnings dynamics of workers.

For this purpose, 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.

6.1. Inequality in labor market outcomes and lifetime earnings

Figure 6 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 6, on average, the top quartile already finds more stable jobs at labor market entry. Moreover, over the life cycle, the average separation rate of the top quartile remains close to the most stable job which represent 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

Figure 7. Lorenz curve of cumulative unemployment duration.

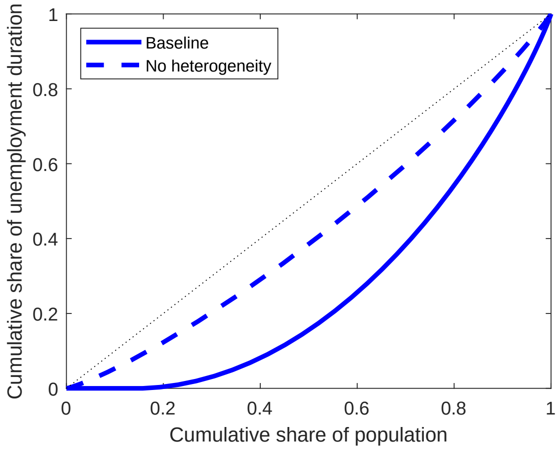

Notes: The solid line displays the Lorenz curve of cumulative unemployment duration for the baseline economy. The dashed line represents the Lorenz curve for a counterfactual economy in which there is no heterogeneity in job stability conditional on age.

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.

This heterogeneity in job stability translates into a large inequality in employment history. Figure 7 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.

Table 1 reports the Gini coefficient of cumulative unemployment duration as well as the variance and the quantile ratio of lifetime earnings for the baseline economy and for the counterfactual economy. For each measure, I use the relative difference between the two

Table 1. Inequality of cumulative unemployment duration and lifetime earnings in the baseline economy.

Model	Gini of unemployment	Lifetime earnings	
		Variance	Quantile ratio
Full model	0.36	4.74	1.66
No heterogeneity	0.13	2.69	1.55
Relative difference (%)	176.9	76.30	6.49

Notes: The full model refers to the baseline economy. The model without heterogeneity refers to the counterfactual economy in which there is no heterogeneity in job stability conditional on age.

economies to quantify the effect of heterogeneity in job stability on inequality of employment and lifetime earnings. The associated Gini coefficients of cumulative unemployment duration are 0.36 and 0.13 for the baseline and the counterfactual economy, respectively. With heterogeneity in job stability, the Gini coefficient of cumulative unemployment duration more than doubles compared to the counterfactual without heterogeneity. Moreover, heterogeneity in job stability increases the variance and the quantile ratio of lifetime earnings by 76.3% and 6.5%, respectively. This finding implies that heterogeneity in job stability is also an important driver of inequality in lifetime earnings. In the remaining part of this section, I analyze the mechanism through which heterogeneity in job stability generates inequality in lifetime earnings.

6.2. Life-cycle consequences of job stability

Figure 8a shows the profiles of average human capital level for workers in the top and bottom quartile of the average life-cycle job-stability distribution. 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 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 22.5% higher stock of human capital at the end of the working phase.

A similar pattern is observed in Figure 8b for average earnings and in Figure 8c 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. Because labor earnings depend on human capital, the earnings profile of the bottom quartile features a smaller growth over the life cycle. The profiles of the bottom quartile becomes flat after growing in the first 10 years, whereas the average earnings profile of the top quartile exhibits strong concavity and increases over

Figure 8. Average life-cycle profiles by job stability.

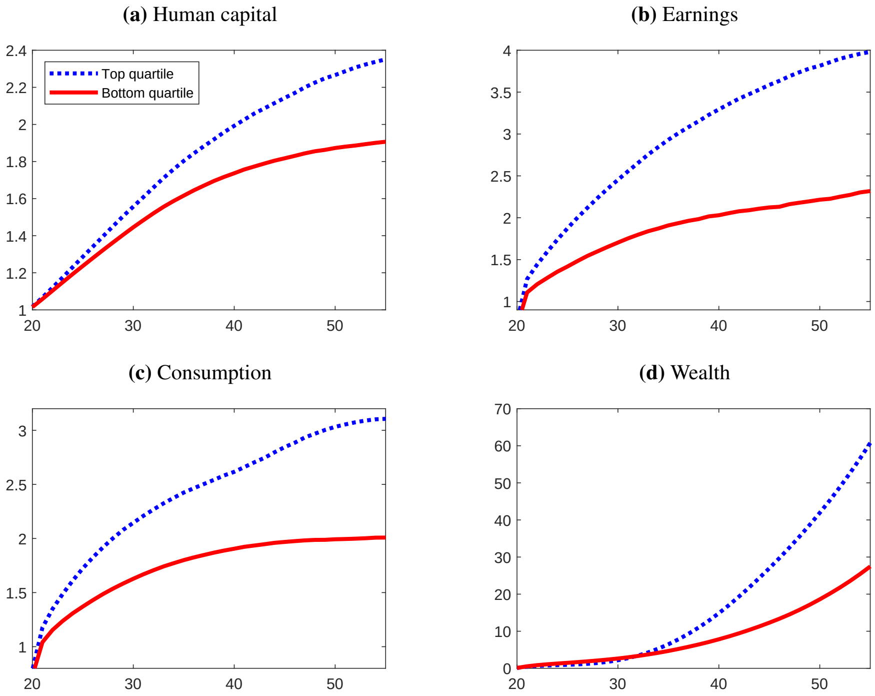

Notes: This figure shows the average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel). In all plots the solid line represents workers with most unstable jobs and the dashed line represents workers with most stable jobs throughout working life.

the life cycle.

The consumption profiles in Figure 8c have a similar shape. The borrowing constraint impedes the desire for consumption smoothing of young workers and the consumption profile strongly tracks the earnings profile at the beginning of the life cycle. Towards the end of the working phase, the average consumption profile of the top quartile becomes attenuated and there is no co-movement with the earnings profile. The reason is that workers strongly increase their wealth accumulation for retirement purposes around the middle of the working phase which dampens their consumption growth. For workers in the bottom quartile, the average consumption profile does not exhibit an upward movement after age 40 as a result of the weak earnings growth for the remaining lifetime.

Figure 8d shows the average life-cycle profiles of wealth. For workers in the top quartile, the average asset level is close to zero at the beginning of the working phase, but starts increasing in a strictly convex manner 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-earners to low-earners, reducing consumption inequality across workers after retirement. Consumption inequality after retirement arises from the difference between workers with small pre-retirement earnings who can save little for retirement purposes and workers with large pre-retirement earnings and wealth. Progressive pension benefits alleviate earnings shocks accumulated over the working life which would be otherwise fully carried over to the retirement. Pension progressivity also reduces consumption inequality of younger workers before retirement. Anticipating redistribution through pension benefits after retirement, workers can increase their consumption during working life. Hence, the insurance effect of a redistributive pension system against low lifetime earnings provides workers a better capability to smooth consumption over the life cycle.

7. 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 that leads to the highest welfare in the economy. I set up a grid for the parameter γ 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.

Figure 9. Welfare and changes in progressivity level.

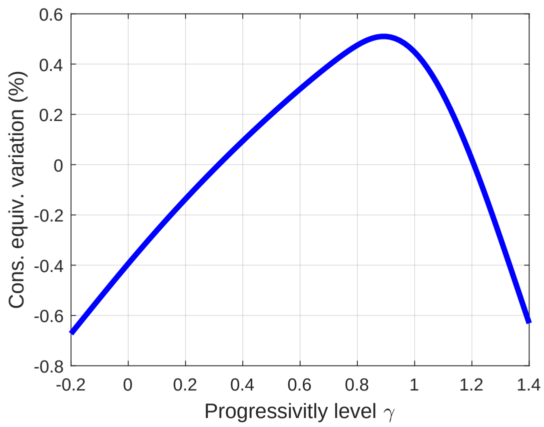

7.1. Welfare measure

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.

7.2. Ex-ante optimal pension system

Figure 9 displays the welfare change induced by varying the progressivity level of the pension system as consumption-equivalent variations. Recall that a pension system with $\gamma = 0$ implies that retirement benefits linearly increase in lifetime earnings, whereas a pension system is progressive if $\gamma > 0$. The figure shows that welfare exhibits a strictly concave shape in the progressivity level of the pension system. This result 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

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

Parameters		Tax (%)	Welfare change (%)
γ	ϕ		
−0.2	0.350	8.552	−0.676
0	0.442	8.580	−0.394
0.2	0.553	8.609	−0.137
0.316	0.627	8.626	-
0.4	0.685	8.638	0.094
0.6	0.841	8.667	0.298
0.8	1.022	8.695	0.473
0.894	1.118	8.708	0.518
1	1.231	8.723	0.450
1.2	1.468	8.742	0.028
1.4	1.745	8.753	−0.633

Notes: The results are obtained by varying the parameters γ 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 and 2 show the parameters of the pension system. Column 3 presents the payroll tax rate in the economy that achieves budget balance for the pension system and column 4 the welfare change relative to the baseline economy as equivalent variation in consumption in percentages.

low lifetime earnings, but it comes at the cost of distorting human capital investment and retirement decisions.

Table 2 displays the policy parameters associated with each progressivity level γ. To give a manageable overview of the results, Table 2 displays the results only for selected values of the progressivity level that were taken into account for the welfare analysis. The parameters in the baseline economy are $\gamma = 0.316$ and $\phi = 0.627$ with a payroll tax rate of 8.63%. Comparing the welfare change across all parameter specifications, the optimal policy is a pension system with $\gamma = 0.894$ which leads to a welfare gain of 0.52% in terms of consumption-equivalent variation. The payroll tax rate that achieves budget balance for the pension system increases to 8.71% and the corresponding level of parameter ϕ 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 the Appendix, I discuss the welfare changes of the optimal pension system as functions of job-separation rate, wealth level, and human capital.

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	64.40	64.18
Human capital below median	64.11	63.63
Human capital above median	64.92	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.

7.3. Retirement incentive distortion

The ex-ante optimal pension system changes the retirement decision of workers. 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 64.4 to 64.18. 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 accumulated 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 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

Figure 10. Relative change in average life-cycle profiles by job stability.

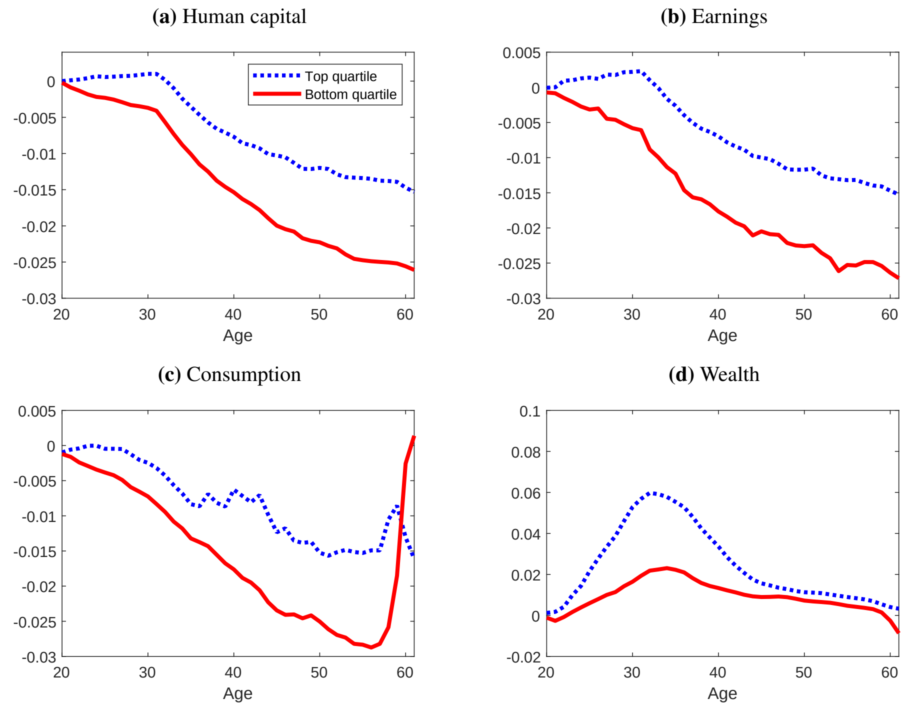

Notes: This figure shows the relative deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy under the optimal pension system compared to the baseline economy. In all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

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.

7.4. 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 ex-ante 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 10 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 10a, starting at age 30, the relative change in the profiles of average human capital declines steadily until the early retirement age, implying that workers accumulate less human capital under the optimal policy than under the baseline specification.[3] The gap in human capital level grows over life and amounts to -1.5% for agents in the top and -2.6% 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. The relative change in human capital level is, however, larger for the bottom quartile for two reasons. First, a higher level of porgressivity reduces the returns on human capital investment and second, workers with unstable jobs anticipate higher retirement benefits compared to the baseline case. This increase in retirement benefits leads to higher lifetime resources of workers with unstable jobs and this wealth effect generates less need to acquire human capital. Therefore, the discouragement of human capital investment is stronger for workers in the bottom quartile.

Interestingly, in the first 10 years after labor market entry, the relative change of average human capital is positive for the top quartile and slightly negative but close to zero for agents in the bottom quartile. Moreover, from age 45 onwards, the divergence in human capital from the baseline levels decelerates and the profiles in Figure 10a become flatter over time. The reason why workers do not reduce their investment in human capital at the beginning of the life cycle is that human capital investment is highly productive in the early working phase. Young workers strive to accumulate human capital to increase prospective labor earnings growth and lifetime earnings such that human capital accumulation mostly happens at the beginning of the life cycle. The decrease in return on human capital investment, which realizes after retirement, has little impact on the incentives for human capital accumulation and labor market outcomes of young workers. As a consequence, the policy change does not impede human capital accumulation in the early working phase. When the retirement age gets closer and the disincentives for human capital investment start to grow and become more relevant in the later stage of life, workers reduce their effort provision in any case because investment in human capital becomes unproductive. Hence, the gap in human capital between the baseline economy and the economy with higher pension progressivity does not grow further and remains almost unchanged close to retirement.

[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 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 shown above, 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 10b and Figure 10c. The first observation is that the relative changes in the profiles of average earnings closely follow the shape of the profiles in Figure 10a. For both groups of agents, 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 policy change 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.

The profiles of average assets in Figure 10d reflect this change in the consumption-saving behavior. For both groups of agents, 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. Agents in the top quartile increase their savings up to age 30 because their average earnings are higher under the optimal pension system. As human capital and earnings decrease over the life cycle, the average level of assets gets close to the average assets in the baseline economy. Agents in the bottom quartile reduce their asset accumulation close before 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 11 displays that the additional progressivity of the 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 67% reaching a value of 0.23

Figure 11. The life-cycle profiles for variance of log consumption and Gini coefficient of consumption.

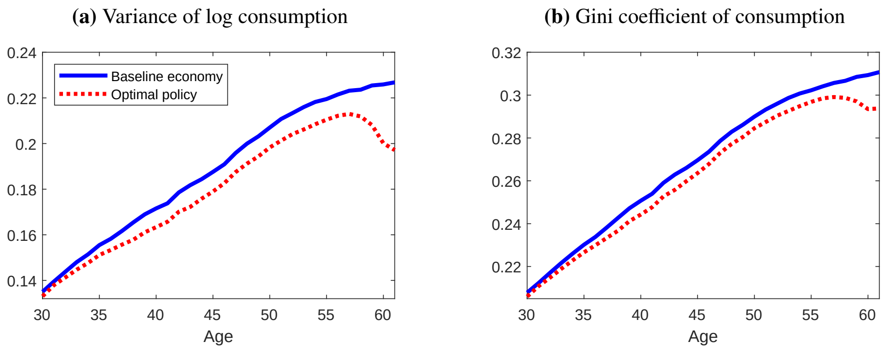

Notes: The left panel shows the variance of log consumption and the right panel the Gini coefficient of consumption over the life cycle. The variances and Gini coefficients are computed using the whole sample.

before retirement. In the economy under the optimal pension system, the profile becomes flatter and has a concave shape. At the end of the working phase, the variance of log consumption reaches a value of 0.2. 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 11b, 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 0.31 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 0.29.

To summarize, the welfare analysis shows that starting from the U.S. pension system, an increase in pension progressivity leads to a welfare gain of 0.52% in terms of lifetime consumption. The increase in progressivity distorts human capital investment and the retirement decision of workers. However, the insurance effect of the increase in pension progressivity exceeds the distortionary effects. It provides redistribution to workers who suffer from unstable jobs and reduces consumption inequality over the life cycle.

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

This section analyzes the consequences of a shift in the job-stability distribution on inequality in labor market outcomes and lifetime earnings, and the implications thereof 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 of an economy. In particular, the Covid-19 pandemic severely affected the dynamism of the U.S. labor market (see, for example, Coibion et al., 2020, Bartik et al., 2020, and Albanesi and Kim, 2021) which has led to macroeconomic shifts in the distribution of job stability. Such recent development in labor markets puts forward the importance of studying the implications of a shift in the job-stability distribution on the optimal design of pension systems for the post-pandemic economy.[4]

Intuitively, a decrease in job-separation rate reduces the incidence of job separations and 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 optimal pension design should take account of such 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 inequality in labor market outcomes and lifetime earnings as well as the implications thereof 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,

[4]While this section focuses on analyzing the consequences of a shift in the job-stability distribution towards more stable jobs, it should be noted that following a decline in job stability, all results obtained in the following go in the opposite direction.

Figure 12. Average life-cycle profile of separation rate for the baseline economy and for the economy with higher job stability.

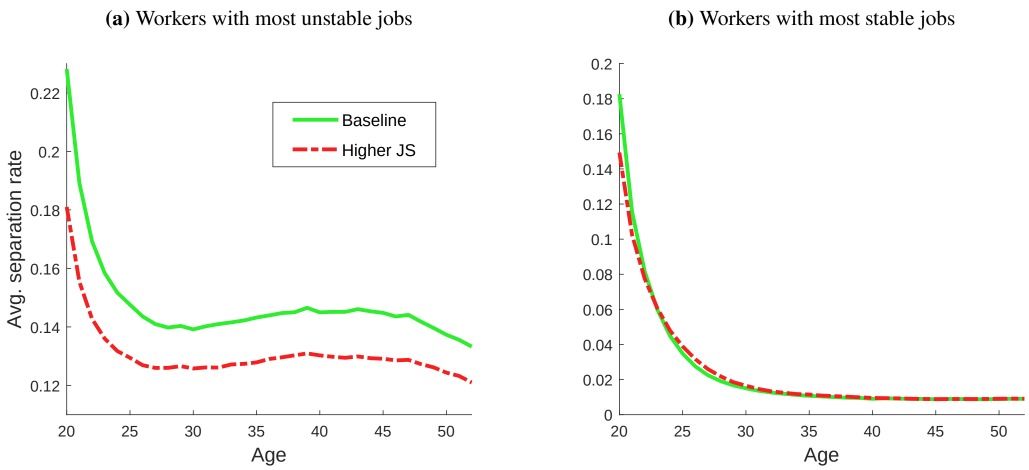

Notes: The left panel shows the average separation rate of workers with most unstable jobs and the right panel the average separation rate of workers with most stable jobs.

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 5% in the economy with higher job stability.

8.1. Inequality in labor market outcomes and lifetime earnings

Figure 12 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 5 percentage points for workers with most unstable jobs. This gap 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.

Table 4 reports the Gini coefficient of cumulative unemployment duration and the variance and the quantile ratio of lifetime earnings for each economy. As in the previous

Table 4. Inequality of cumulative unemployment duration and lifetime earnings in the baseline economy and in the economy with higher job stability.

Model	Gini of unemployment	Lifetime earnings	
		Variance	Quantile ratio
Baseline			
Full model	0.36	4.74	1.66
No heterogeneity	0.13	2.69	1.55
Relative difference (%)	176.9	76.30	6.49
Higher job stability			
Full model	0.35	5.10	1.64
No heterogeneity	0.13	3.45	1.59
Relative difference (%)	176.23	47.91	3.59

Notes: The full model refers to the baseline economy with heterogeneity in job stability. The model without heterogeneity refers to the counterfactual economy in which there is no heterogeneity in job stability conditional on age.

section, I simulate a counterfactual economy in which all workers have the same job-separation rate conditional on age. The age-dependent separation rate corresponds to the average age-dependent separation rate of all workers in the full model with heterogeneity in job stability. Compared to the baseline economy, higher job stability leads to lower Gini coefficient of unemployment. This result indicates that higher job stability does not only lead to smaller inequality of unemployment, but also decreases the relative difference between the full model and a counterfactual without heterogeneity in job stability. A smaller proportion of unemployment inequality is therefore attributable to heterogeneity in job stability if the distribution of job stability shifts towards more stable jobs. Moreover, heterogeneity in job stability accounts for a smaller proportion of lifetime earnings inequality for the economy with higher job stability as indicated by the relative difference between the full model and the counterfactual in Table 4.

One surprising result is that the Gini coefficient of unemployment duration in the counterfactual does not decrease in the economy with higher job stability. The reason is the following. By construction, a lower separation rate of the most unstable job decreases the average separation rate in the economy and induces a lower job-stability heterogeneity. This leads to two opposing effects. On the one hand, lower average separation rate raises unemployment inequality, while on the other hand, smaller job-stability heterogeneity decreases the inequality. In the full model, the effect of smaller heterogeneity in job stability outweighs the decrease in the average separation rate such that the Gini coefficient declines. In the counterfactual model, there is no heterogeneity in the job-separation rate conditional on age such that only the effect of lower average separation rate exists, and this leaves the unemployment inequality unchanged. Nevertheless, the relative difference between the full model and the counterfactual decreases in response to the increase in job stability implying that a smaller proportion of the existing unemployment inequality is due to heterogeneity in job stability.

Concerning the inequality measures for lifetime earnings, Table 4 indicates that the variance of lifetime earnings increases in the economy with higher job stability. As explained above, the average separation rate moves in response to the adjustments of the separation rate of the most unstable job, which affects the human capital accumulation process of workers. A lower average separation rate offers more opportunity to invest in human capital. The stochastic process of human capital accumulation becomes more relevant for the earnings growth over the life cycle and increases the variance of lifetime earnings. Hence, higher job stability leads to a higher variance of lifetime earnings for both the full model and the counterfactual without heterogeneity in job stability.

Moreover, higher job stability leads to a decrease in the quantile ratio of lifetime earnings in the full model. The increase in the quantile ratio of the counterfactual model implies that, as for the variance of lifetime earnings, a lower average separation rate leads to more heterogeneously distributed lifetime earnings. However, job-stability heterogeneity decreases in the model with higher job stability, and the total effect implies a decrease in the ratio between the 25th and 75th percentile of the lifetime-earnings distribution.

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

How does a shift in the job-stability distribution 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 1% at the early retirement age for workers with most stable jobs, workers with unstable jobs achieve on average a human capital level that is 2.5% 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. On average, earnings increase for workers in the economy with higher job stability, and workers in the bottom quartile benefit more from the shift in the job-stability distribution. More specifically, earnings become around 6% higher for workers with the most unstable jobs, and around 2% for workers with most stable jobs over life. The increase

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

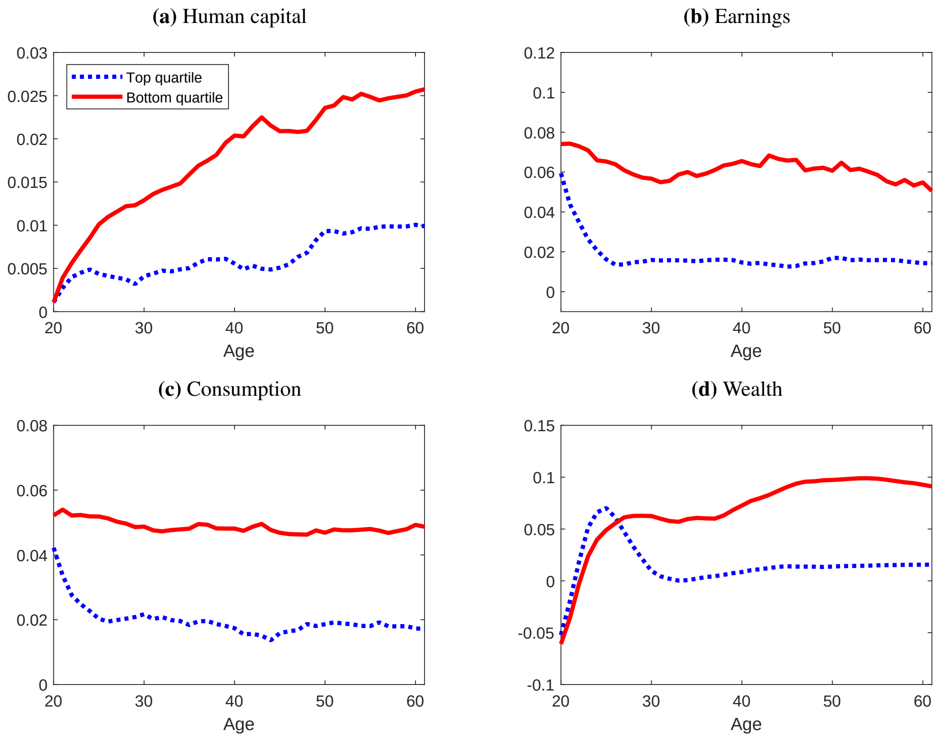

Notes: This figure shows the relative deviation in average life-cycle profiles of human capital (upper left panel), labor earnings (upper right panel), consumption (lower left panel), and wealth (lower right panel) in the economy with higher job stability compared to the baseline economy. In all plots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

in average earnings are 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, workers get less frequently unemployed which leads to higher average labor earnings. Note that the profile of the top quartile moves towards zero immediately after labor market entry because these workers find on average very stable jobs at the beginning of their career path as shown in Section 6.

A similar result is obtained for consumption in Figure 13c. As workers are identical in their state variables ex ante and consumption is a forward-looking variable, its relative deviation from the profiles in the baseline economy exhibits a jump in the first period after labor market entry in response to the shift in the job-stability distribution. On average, workers in the bottom and top quartile consume around 5% and 2% more in the economy with higher job stability than in the baseline economy.

Finally, Figure 13d displays the profiles of relative deviation from the baseline economy for wealth. One striking observation is that the profiles do not imply a persistent increase in wealth over the life cycle for the two groups of workers. The relative change in wealth strongly varies during the first 15 years after labor market entry. At the beginning of life, workers decrease their precautionary savings relative to the baseline economy. The average

Table 5. Ex-ante optimal pension systems in the baseline economy and in the economy with higher job stability.

Model	Optimal policy parameters		Tax (%)	Welfare change (%)
	γ	ϕ		
Baseline	0.89	1.12	8.71	0.52
Higher job stability	0.93	1.16	8.47	0.53

Notes: Column 1 specifies the model. Column 2 shows the optimal progressivity level of the pension system. Column 3 and 4 present the associated parameter ϕ that determines the level of pension benefits and the tax rate in the economy that achieve budget balance for the pension system. Column 5 shows the welfare change in percentages relative to the baseline economy as equivalent variation in consumption.

level of wealth is thus lower by more than 5% compared to the baseline economy. Over time, this relative change in wealth turns positive. 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 9% higher wealth and workers with most stable jobs 1.6% higher wealth than in the baseline economy.

8.3. Distribution of job stability and ex-ante 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 progressivity level γ 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 5. The optimal progressivity parameter γ increases to 0.93. At this optimal degree of progressivity, the parameter value of ϕ increases from 1.12 to 1.16 in the economy with higher job stability. Interestingly, the payroll tax rate decreases by 24 basis points despite the increase in the parameter ϕ. The reason for this decrease in the tax rate is the shift in the job-stability distribution. Recall that the average job stability and 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.

These changes in the optimal benefit level indicate that the optimal pension system becomes more redistributive in an economy with more stable jobs. This may be surprising given the results in Section 8.1 that higher job stability decreases inequality in unemployment as well as the gap in earnings and consumption between workers with the most stable and the most unstable jobs over the life cycle. One would expect that the pension system should become less redistributive in an economy with higher job stability because the risk of unstable employment decreases and therefore, it might be efficient to provide incentives for human capital investment.

To understand this result, recall from Section 8.1 that a shift in the lower tail of the job-stability distribution primarily affects the average job-separation rate of young workers. Moreover, since these young workers face on average a lower unemployment risk, they accumulate smaller buffer stocks as shown in Section 8.2. Importantly, in the presence of a borrowing constraint, a low degree of job stability and payroll taxes restrict the ability to smooth consumption over the life cycle, and these effects mutually affect each other: Low job stability results in low labor earnings as well as a large precautionary motive for saving, pushing workers towards higher marginal propensities to consume. This makes the payroll tax finance of the pension system more costly in terms of utility for young workers in an economy with lower job stability. At the same time, payroll taxes depress consumption and precautionary savings. Payroll taxes thus constrain workers from adjusting their consumption-saving decision in an economy with lower job stability and amplify the detrimental effects of job instability.

Moreover, increasing the degree of pension progressivity requires an increase in the payroll tax rate. This is welfare-detrimental for young low-income workers and the negative effect of payroll taxes is stronger in the baseline economy compared to the economy with higher job stability. Because workers face a borrowing limit of zero and labor earnings are relatively low at the beginning of the working phase, additional tax paid for every unit of labor earnings depresses consumption and precautionary savings of young workers. Although a more redistributive pension system increases the present value of additional pension wealth relative to the additional taxes paid for workers with unstable jobs, the negative effect of a higher tax rate dominates the gain from obtaining higher insurance against low lifetime earnings. Therefore, a higher benefit level leads to large utility losses for young workers in the economy with lower job stability and it is ex-ante optimal to implement a less redistributive pension system in the baseline economy.

This result is in contrast to the implication of the stylized two-stage model discussed in Section 3. In the stylized model, which abstracts from life-cycle components and borrowing constraints, the increase in job-stability heterogeneity implies an increase in the progressivity level for the optimal tax system. As agents can perfectly smooth consumption over life, the increase in lifetime earnings inequality makes a higher degree of redistribution 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). If these components are neglected, one may draw wrong conclusions and inaccurate policy prescriptions from analyzing the design of optimal pension systems.

9. Conclusion

This paper studies how a progressive pension system should take heterogeneity in job stability into account and quantifies potential welfare gains from implementing the optimal pension system. Using a life-cycle model with heterogeneity in job stability, human capital investment, and retirement decision, I show that heterogeneity in job stability is a key driver of inequality in earnings and consumption over the life cycle. Progressive pension systems provide redistribution and 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 redistributive and 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.

I motivate this study by developing a stylized two-stage model which shows that heterogeneity in job stability generates inequality in lifetime earnings that cannot be sufficiently addressed by taxing annual earnings. This is because annual earnings taxation fails to capture persistent differences in the employment history and does not achieve sufficient redistribution. Lifetime earnings taxation leads to a large welfare gain by tackling the inequality in lifetime earnings more directly. This result motivates the study of optimal pension systems in the presence of heterogeneity in job stability. Pension systems redistribute based on lifetime averages of earnings and therefore, pension systems are the most commonly applied form of lifetime earnings taxation in practice.

In a realistically calibrated life-cycle model with job-stability heterogeneity, I analyze the optimal pension system. 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 0.52% 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 in the last few decades, I study the consequences of a shift in the job-stability distribution on the optimal design of progressive pension systems. Even though a shift towards higher job stability decreases inequality in labor market outcomes, I find that the optimal pension system becomes more redistributive in an economy with more stable jobs. The reason is that a high degree of redistribution through the provision of pension coverage requires a high payroll tax rate which is welfare-detrimental for young workers who have to cope with a high risk of job loss. Payroll taxes impede the ability to smooth consumption over the life cycle, and the increase in job stability makes payroll taxes less costly due to a decrease in career interruptions and earnings losses.

The results of this paper demonstrate that heterogeneity in job stability has crucial implications on the optimal tax policy. In future research it would be interesting to investigate multiple tax systems simultaneously: The interaction between progressive pension systems and progressive annual earnings taxation is potentially important in studying optimal tax policies in the presence of job-stability heterogeneity. As pointed out by Diamond (1977), the two tax instruments address different types of earnings inequality and may complement each other.

References

Albanesi, Stefania and Jiyeon Kim (2021). “Effects of the COVID-19 Recession on the US Labor Market: Occupation, Family, and Gender.” In: *Journal of Economic Perspectives* 35.3, pp. 3–24.

Bartik, Alexander, Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matt Unrath (2020). “Measuring the labor market at the onset of the COVID-19 crisis.” Working Paper 27613. National Bureau of Economic Research.

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.

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.

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.

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.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2020). "Labor Markets During the COVID-19 Crisis: A Preliminary View." Working Paper 27017. National Bureau of Economic Research.

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.

Diamond, Peter (1977). "A framework for social security analysis." In: *Journal of Public Economics* 8.3, pp. 275–298.

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.

French, Eric (2005). "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." In: *The Review of Economic Studies* 72.2, pp. 395–427.

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.

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

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.

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.

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 (2018). "Earnings Losses and Labor Mobility Over the Life Cycle." In: *Journal of the European Economic Association* 17.3, pp. 678–724.

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.

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.

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

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.

10. 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 $V_2 := V^r(a, w, h, j_r)$. A worker's decision problem is described by

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

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

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

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

$$\mathbb{E}V_{\max} = pV_1 + (1 - p)V_2 + \int_\delta^\infty \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon \tag{6}$$

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

$$\int_\delta^\infty \varepsilon f(\varepsilon, \mu, \sigma) d\varepsilon = \delta \cdot \frac{\exp\left(-\frac{\delta-\mu}{\sigma}\right)}{1 + \exp\left(-\frac{\delta-\mu}{\sigma}\right)} + \sigma \cdot \log\left(1 + \exp\left(-\frac{\delta - \mu}{\sigma}\right)\right). \tag{7}$$

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 (7) 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 (6) can be simplified to

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

which yields Equation (2) in Section 4.1.

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. Calibration

A.3.1. 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 $\bar{\lambda} = 0.35$. 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 $[1-\bar{\lambda}, 1-\underline{\lambda}]$, 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.

A.3.2. 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

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

$\pi_s(a,\theta)$ 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	0.992	Quarterly discount factor
κ	0.325	Utility cost of effort
π_e	0.429	Probability of a job offer when employed
π_u	0.867	Probability of a job offer when unemployed
ψ_w	0.539	Marginal distribution of w^*
ψ_λ	0.466	Marginal distribution of $1-\lambda^*$
θ	0.435	Joint distribution of w^* and $1-\lambda^*$
$\bar{p}_H$	0.051	Skill upgrading probability
ρ	0.985	Persistence of skill upgrading probability
p_H^*	0.057	Probability to move to h^*

A.3.3. Approximation of the U.S. pension system

Figure A.1. U.S. pension system and the approximating function.

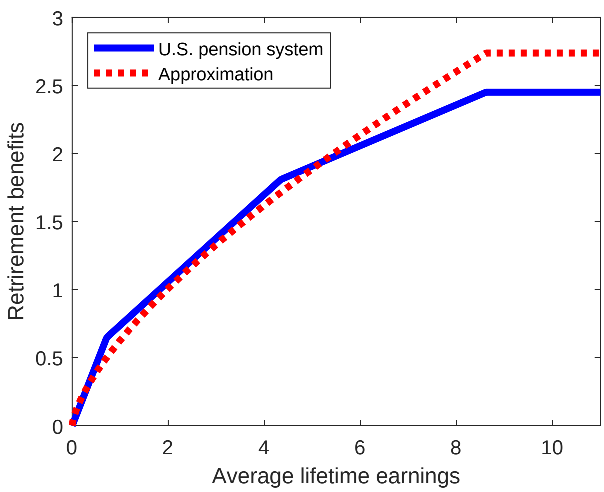

A.3.4. Approximation of average lifetime earnings

To evaluate the accuracy of the approximation of average lifetime earnings by the final level of human capital that workers attain 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 final level of human capital of each worker. This yields an R^2 statistic of 0.885. The final level of human capital of workers thus approximates the average lifetime earnings well, explaining 88.5% of the variance in the average lifetime earnings.

A.4. 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.2.

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

Figure A.2b presents the welfare change as a function of the wealth level. Interestingly, the function is not monotonic: the welfare change first increases in the level of wealth, but then, it begins to decrease as the asset level increases further. Recall that the ex-ante optimal policy is associated with a higher payroll tax rate. Starting from zero wealth, higher wealth allows for a better ability of consumption smoothing in face of the increase in 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.

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

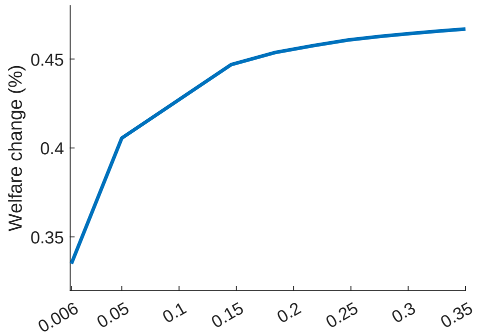

(a) Separation rate

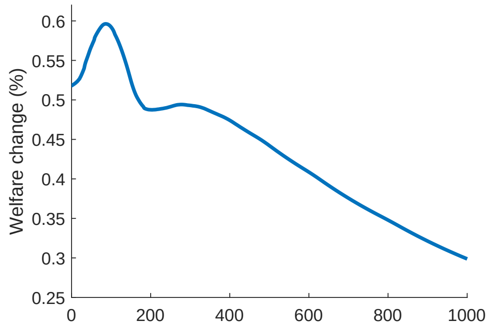

(b) Wealth

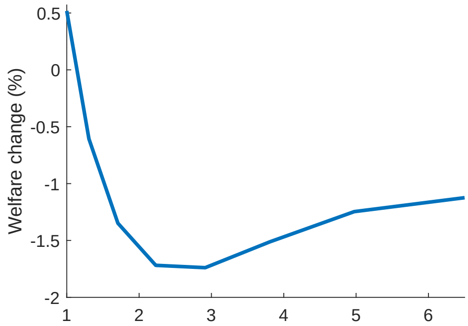

(c) Human capital

Finally, Figure A.2c shows the welfare change as a function of human capital. Starting from the initial level, a larger stock of human capital implies a decrease in welfare. This points out that workers with low human capital gain from an increase in pension progressivity, while workers with high human capital lose out. An increase in pension progressivity decreases the welfare gain of workers with high human capital in two respects. Firstly, higher pension progressivity reduces pension wealth and therefore the lifetime resources of workers with high human capital. Secondly, the increase in the payroll tax rate entailed by the increase in pension progressivity also decreases welfare. One interesting finding is that the graph exhibits a bend point. Above a certain level of human capital, the latter effect diminishes because labor earnings become high enough such that the effect of the increase in payroll tax rate does not affect consumption as strongly as for lower labor earnings. As a consequence, the welfare change slightly increases for higher human capital levels, but remains negative implying that the increase in pension progressivity reduces expected lifetime utility for workers with high levels of human capital.