Differential Mortality and the Progressivity of Social Security*

Shantanu Bagchi†

July 23, 2018

Abstract

There is a well-established negative correlation between earnings and mortality risk. Using a calibrated general-equilibrium macroeconomic model, this paper examines how this correlation interacts with Social Security's benefit-earnings rule. My findings suggest that the welfare ranking of alternative benefit-earnings rules is, in fact, sensitive to differential mortality risk. With a progressive benefit-earnings rule, Social Security's internal rate of return is lower for households with unfavorable earnings histories, and their relatively high mortality risk also heavily discounts their expected utility from old-age consumption. Because of these two effects, I find that welfare is maximized with a benefit-earnings rule *less* progressive than the current U.S. rule in the presence of differential mortality, but with a *more* progressive rule when differential mortality is eliminated.

JEL Classifications: E21, E62, H55

Keywords: differential mortality; Social Security; mortality risk; labor income risk; incomplete markets; social insurance

*This project received financial support from the W.E. Upjohn Foundation for Employment Research (ECRA Grant No. 15–150–01) and the College of Business & Economics Summer Support Grant Program. I would like to thank Frank Caliendo, Matt Chambers, Jim Feigenbaum, and Juergen Jung for useful comments and suggestions.

†Department of Economics, Towson University, 8000 York Road, Towson, MD 21252. Phone: +1-410-704-2191. Email: sbagchi@towson.edu

1 Introduction

Social Security is traditionally viewed as a vehicle that partially insures individuals against risks that markets do not insure well, such as the risk of an uncertain lifetime, and also the risk of old-age poverty caused by unfavorable labor-market outcomes. Social Security annuities are paid until death, so they are commonly believed to insure individuals against the risk of out-living their savings. Meanwhile, Social Security benefits are a concave, increasing function of work-life earnings. The curvature of this function determines how much insurance Social Security provides against unfavorable labor-market events, such as the inability to secure a high-paying job, or unemployment.[1]

While linking public pension benefits to work-life income is common within the industrialized world, the concavity of this relationship in the U.S. is unusual. Under current law, benefits replace 90% of the average work life earnings for an individual who is in the bottom 20% of the wage distribution, but only about 40-50% for individuals whose wages are higher than the average wage. Therefore, an individual at the bottom of the earnings distribution receives a higher return on every dollar of Social Security contributions paid, relative to an individual at the top of the earnings distribution. While this arrangement is intended to provide insurance against unfavorable labor income shocks, its effect on the distribution of lifetime utility across households depends on a multitude of economic and demographic factors.[2]

From the perspective of a household, an important determinant of the welfare gains or loses from Social Security is its life expectancy. Because Social Security is a retirement pension, households who expect to live longer are likely to experience a larger welfare gain, compared to households who do not. However, empirical evidence suggests that there is a significant positive correlation between income and life expectancy in the U.S. (Kitagawa and Hauser, 1987). This phenomenon, referred to as differential mortality, has important implications for Social Security, because the positive correlation between income and survivorship can potentially undo the redistribution from Social Security's current benefit-earnings rule (Coronado et al., 2002, 2011).

In this paper, I examine if differential mortality has any implications for the welfare effects of Social Security's benefit-earnings rule. To do this, I first set up a simple two-period overlapping generations model with incomplete markets, an unfunded public pension system that resembles Social Security, and rational households who experience mortality risks that are negatively correlated to their earnings. I use this simple model to illustrate two facts. First, I show that for households with lower earnings, Social Security's internal rate of return from a progressive benefit-earnings rule is actually lower in the presence of differential mortality, compared to when mortality risk is symmetric and uncorrelated to earnings. Second, I show that due to their higher mortality risk, these households also heavily discount the utility from higher old-age consumption. Together, these results suggest that the consumption-smoothing effects of Social Security may not be as useful when income and survivorship are positively correlated.[3]

Next, I evaluate the quantitative importance of these mechanisms using a rich multi-period

[1] In theory, Social Security also enables intergenerational risk sharing, but the model considered here does not have any aggregate risk to share across generations.

[2] It is well known that Social Security improves welfare in the aggregate only in the counterfactual case where the economy is dynamically inefficient (Samuelson, 1975).

[3] It is useful to note that Social Security does not really have a true longevity insurance role in the current model. Longevity risk, as it is usually understood, is the danger of outliving one's assets. This is impossible in a rational expectations framework, because households will always choose to deplete assets at the maximum possible age. The only time there is a true longevity insurance role for a public pension scheme in a rational life cycle framework is when the social planner has private information about individual mortality risks, and that they are actually lower than what the households believe them to be.

general-equilibrium model calibrated to the U.S. economy. In the model, households experience two types of risk: an age-dependent mortality risk that is negatively correlated to their earnings, and a labor income risk that consists of a permanent productivity fixed effect, and also an idiosyncratic shock with both persistent and transitory components. Factor markets in the model are competitive, firms maximize profit, and the government provides public goods and Social Security. I calibrate this model to match key institutional features of Social Security, and also U.S. macroeconomic aggregates, such as overall capital accumulation, the pattern of labor supply over the life cycle, earnings heterogeneity, and the share of government expenditures in GDP.

Then, I use this model to compute the welfare implications of alternative benefit-earnings rules, ranging from a fully proportional (or linear) function of past earnings (i.e. with zero implicit redistribution), to a fixed benefit that is completely unrelated to past earnings (i.e. with full redistribution). Finally, I examine if the welfare ranking of these rules is sensitive to the negative correlation between earnings and mortality risk. I also identify the macroeconomic effects of differential mortality on the labor market, capital accumulation, national income, and the government's budget.

My computations suggest that the welfare ranking of alternative benefit-earnings rules is, in fact, sensitive to differential mortality. I find that more progressive benefit-earnings rules actually provide worse work-retirement consumption smoothing for households with relatively unfavorable earnings histories in the presence of differential mortality. As a result, aggregate welfare is maximized under a benefit-earnings rule *less* progressive than the current U.S. rule. However, when differential mortality is eliminated, welfare is maximized with a benefit-earnings rule *more* progressive than the current U.S. rule. In other words, the lower internal rate of return and the lower expected utility of old-age consumption limits the value of improved work-retirement consumption smoothing when survivorship is positively correlated with earnings. I also find that modifying the shape of the benefit-earnings rule, given Social Security's current payroll tax rate and taxable maximum, has only a small effect on key macroeconomic aggregates: between the computational experiments, capital stock, labor, national income, and government expenditures, do not change by more than one percentage point, both with and without differential mortality.

My computations also predict that these results are fairly robust. Specifically, I examine how the welfare effects of the benefit-earnings rules, both with and without differential mortality, depend on particular modeling assumptions. First, I consider an alternative specification of the model, where the accidental bequests originating from the assets of the deceased, rather than being taxed away, are redistributed to the surviving population in a lump-sum fashion each period. Second, I examine how the welfare effects depend on the empirical estimates of differential mortality used to calibrate the initial model. In each case, I find that welfare is maximized with a benefit-earnings rule *less* progressive than the current U.S. rule with differential mortality, and with a benefit-earnings rule *more* progressive than the current U.S. rule when differential mortality is eliminated. Finally, I consider a "maximin" criterion to evaluate the welfare effects of the benefit-earnings rules, instead of the ex-ante expected lifetime utility criterion traditionally used in the literature. I find that in this case, welfare is maximized with a benefit-earnings rule *less* progressive than the current U.S. rule, both with and without differential mortality.

This paper contributes to three separate strands of the literature. First, it contributes to a large literature that characterizes the optimal redistribution scheme in a heterogeneous-agent economy, accounting for distortions to consumption, saving, and labor supply.[4] Two papers in this literature that highlight the importance of earnings history-dependent tax systems are Grochulski

[4]Notable studies in this literature include Saez (2002), Cremer et al. (2004), Sheshinski (2008), Golosov et al. (2011), and Farhi and Werning (2013), among many others.

and Kocherlakota (2010) and Michau (2014). Grochulski and Kocherlakota (2010) show that in an economy where agents have nonseparable preferences and private information about their skill levels, it is possible to implement a socially optimal allocation through a linear labor income tax during the working life, and constant payment during retirement that is conditioned on the agents' entire labor income history. Michau (2014) builds on their results and shows that an earnings history-dependent social security system can implement the optimal allocation that accounts for labor supply distortions both along the extensive and intensive margins. However, none of these studies account for differential mortality, thereby ignoring its potential redistributive consequences for such tax-and-transfer systems.

Second, the current paper contributes to the quantitative-macro literature on the welfare consequences of alternative social security schemes in the context of the U.S. Three studies from this literature that are closest to the current paper are Huggett and Ventura (1999), Fuster et al. (2003), and Nishiyama and Smetters (2008). Huggett and Ventura (1999) examine the distributional consequences of replacing current U.S. Social Security with a two-tier pension system with a mandatory, defined-contribution first tier, and a guaranteed second tier with a minimum retirement income.[5] In general, they do not find substantial welfare improvements in switching from the current U.S. program to the two-tier structure. Fuster et al. (2003), on the other hand, examine the welfare effects of unfunded social security in a general equilibrium model with overlapping generations of altruistic individuals that differ in lifetime expectancy and earnings ability. They find that ability shocks and uncertain lifetimes generate significant heterogeneity among households to induce different preferences for Social Security. Finally, Nishiyama and Smetters (2008) find that while the progressive linking of earnings with retirement benefits in the U.S. has beneficial insurance effects, it also introduces various marginal tax rates that distort labor supply. In fact, Nishiyama and Smetters (2008) conclude that the optimal benefit structure in the U.S. is fairly proportional. However, both Huggett and Ventura (1999) and Nishiyama and Smetters (2008) ignore differential mortality, and the results in Fuster et al. (2003) are primarily driven by a two-sided altruism mechanism. Also, the computational experiments in these studies do not allow for a clear interpretation of the interaction between differential mortality and the welfare effects of alternative benefit-earnings rules.

Finally, this paper also contributes to an empirical literature that measures the effect of differential mortality on the progressivity of U.S. Social Security from survey data. Studies such as Coronado et al. (2002) and Coronado et al. (2011) conclude that once the positive correlation between wealth and survivorship is accounted for, Social Security is considerably less progressive than what is defined by the benefit-earnings rule. For example, Coronado et al. (2002) compute the "net tax rate" implicit in Social Security, given by the difference between the present value of the taxes paid and benefits received over the life cycle, expressed as a fraction of potential lifetime income. They find that Social Security becomes regressive after accounting for the mortality differentials between the different income groups, or in other words, it transfers resources from poorer households with shorter lives to wealthier households with longer lives. In a separate study, Meyerson and Sabelhaus (2006) compute the "money's worth" from Social Security, given by the ratio of the present value of benefits to that of the taxes paid over the life cycle. While they conclude that Social Security remains progressive even after accounting for the mortality differences, they find that the degree of progressivity is greatly reduced. However, these studies are purely actuarial in nature, as a result of which they do not account for how households, firms, and the overall economy respond to differential mortality and the modifications to Social Security. The current paper accounts for these effects.

[5] An example of such a reform proposal is the Boskin proposal (Boskin et al., 1987).

The rest of the paper is organized as follows: Section 2 illustrates the key mechanisms of the paper using a simple two-period model, Section 3 introduces the formal quantitative general-equilibrium model, and Sections 4 and 5 describe the baseline calibration and its results. I describe the computational experiments in Section 6, and I examine their results in Sections 7 and 8. Finally, I verify the robustness of the initial results in Sections 9, 10, and 11, and I conclude in Section 12.

2 A simple two-period model

In this section, I use a simple two-period overlapping generations model to illustrate how differential mortality interacts with Social Security's benefit-earnings rule. Lifetime consists of two periods: work (1), and retirement (2). A household earns income and pays taxes in both periods. There are two types of households in every cohort: type$-l$ with lower earnings, and type$-h$ with higher earnings. I assume that both types of households experience mortality risk, but a type$-h$ household has a higher likelihood (Q_h) of surviving into retirement, compared to a type$-l$ household (Q_l), i.e. $Q_l < Q_h < 1$. Fraction a of each cohort is born as type l, and fraction $(1-a)$ is born as type h, and cohort size grows at the rate of n over time. Social Security operates as follows: it collects taxes t_l and t_h from the respective types of households during the working period, and pays benefits b_l and b_h during retirement. Finally, Social Security taxes and benefits across the household types are related as $t_h = (1+x)t_l$ and $b_h = (1+y)b_l$, with $x, y > 0$. This specification allows us to consider Social Security programs with varying degrees of implicit redistribution.[6]

In this model environment, Social Security's budget balancing requires

$$aQ_l N_{t-1} b_l + (1-a)Q_h N_{t-1} b_h = aN_t t_l + (1-a)N_t t_h, \tag{1}$$

where N_t is the size of the cohort born on date t. Given that $b_h = (1+y)b_l$, $t_h = (1+x)t_l$, and $N_t = (1+n)N_{t-1}$, we have

$$\left[aQ_l + (1-a)Q_h(1+y)\right] b_l = \left[a + (1-a)(1+x)\right](1+n)t_l, \tag{2}$$

which gives the benefits for the type$-l$ household

$$b_l = \frac{(1+n)t_l\left[a + (1-a)(1+x)\right]}{\left[aQ_l + (1-a)Q_h(1+y)\right]}. \tag{3}$$

Therefore, the lifetime budget constraint for type$-l$ is

$$c_{l,1} + \frac{c_{l,2}}{1+r} = we_l, \tag{4}$$

where we_l is the present value of lifetime earnings, or

$$\begin{aligned} we_l &= y_{l,1} - t_{l,1} + \frac{y_{l,2} - t_{l,2}}{1+r} \\ &= y_{l,1} + \frac{y_{l,2}}{1+r} - t_{l,1} - \frac{t_{l,2}}{1+r}. \end{aligned} \tag{5}$$

[6] It is important to note here that in a framework with mortality risk and incomplete annuity markets, accidental bequests originate from the assets of the deceased. While I ignore these accidental bequests in the simple two-period model, this assumption is not completely innocuous. In fact, as Caliendo et al. (2014) have shown, if one accounts for how higher mandatory saving through Social Security crowds out these accidental bequests, then Social Security has zero effect on life-cycle wealth. However, this effect is more complicated in the current model, because the accidental bequests also redistribute from high- to low-income, and from high- to low-mortality households. To understand the important of the accidental bequests, I later examine the sensitivity of the baseline computational results with respect to how these bequests are treated within the model.

Substituting for $t_{l,2} = -b_l$, we get

$$\begin{aligned} we_l(SS=1) &= y_{l,1} + \frac{y_{l,2}}{1+r} - t_{l,1} + \frac{(1+n)t_{l,1}\left[a+(1-a)(1+x)\right]}{(1+r)\left[aQ_l+(1-a)Q_h(1+y)\right]} \\ &= we_l(SS=0) + t_{l,1}\left[\frac{(1+n)\left(a+(1-a)(1+x)\right)}{(1+r)\left(aQ_l+(1-a)Q_h(1+y)\right)} - 1\right], \end{aligned} \tag{6}$$

where $we_l(SS=1)$ and $we_l(SS=0)$ represent lifetime wealth with and without Social Security, respectively. Therefore, the necessary and sufficient condition for $we_l(SS=1) > we_l(SS=0)$ is

$$\gamma_l = (1+n)\left[\frac{a+(1-a)(1+x)}{aQ_l+(1-a)Q_h(1+y)}\right] - 1 > r. \tag{7}$$

The left-hand side of (7) is the net internal rate of return of Social Security for the type-l household in this model. Therefore, Social Security has a positive wealth effect if and only if its internal rate of return is higher than the market rate of return r.

Let us define the ratio $\frac{y}{x}$ as a measure of the progressivity of Social Security. That is,

$$\frac{y}{x}\begin{cases} < & 1 \quad \text{Progressive} \\ = & 1 \quad \text{Proportional} \\ > & 1 \quad \text{Regressive} \end{cases}.$$

Then, the following propositions can be stated:

Proposition 1 *When there is no mortality risk and Social Security is proportional, the internal rate of return of Social Security is equal to the population growth rate.*

Proof. Setting $\frac{y}{x} = 1$ and $Q_l = Q_h = 1$ in (7), we get

$$\gamma_l = (1+n) - 1 = n. \tag{8}$$

■

This is the well-known result that in a rational life-cycle framework without mortality risk and zero implicit redistribution, Social Security's internal rate of return is the population growth rate, and Social Security is welfare-improving if and only if that rate is larger than r (Samuelson, 1975).

Proposition 2 *With earnings heterogeneity and symmetric mortality risk, a progressive benefit-earnings rule yields an internal rate of return for Social Security that is higher than the population growth rate.*

Proof. Setting $\frac{y}{x} < 1$ and $Q_l = Q_h = Q < 1$ in (7), we get

$$\gamma_l = \frac{(1+n)}{Q}\left[\frac{a+(1-a)(1+x)}{a+(1-a)(1+y)}\right] - 1. \tag{9}$$

Given that $\frac{y}{x} < 1 \implies \frac{x}{y} > 1$, and $Q < 1 \implies \frac{1}{Q} > 1$, we can conclude that $\gamma_l > n$. Therefore,

$$\gamma_l(\frac{y}{x} < 1, Q_l = Q_h = Q < 1) > \gamma_l(\frac{y}{x} = 1, Q_l = Q_h = 1) = n.$$

■

This shows that in a rational life-cycle framework with earnings heterogeneity and symmetric mortality risk, a progressive benefit-earnings rule yields an internal rate of return for Social Security that is higher than the population growth rate.

Proposition 3 *With earnings heterogeneity and differential mortality risk, Social Security's internal rate of return from a progressive benefit-earnings rule is lower, compared to when mortality risk is symmetric and uncorrelated to earnings.*

Proof. Setting $\frac{y}{x} < 1$ and $Q_h = (1+q)Q_l$, with $q > 0$ in (7), we get

$$\gamma_l = \frac{(1+n)}{Q_l}\left[\frac{a+(1-a)(1+x)}{a+(1-a)(1+q)(1+y)}\right] - 1. \tag{10}$$

To compare (10) with (9), note that if Q is the survival rate of the total population at the end of the first period, then $Q = aQ_l + (1-a)Q_h$. With this substitution, (9) can be rewritten as

$$\gamma_l = \frac{(1+n)}{(aQ_l + (1-a)Q_h)}\left[\frac{a+(1-a)(1+x)}{a+(1-a)(1+y)}\right] - 1. \tag{11}$$

Therefore, a sufficient condition for

$$\gamma_l(\frac{y}{x} < 1, Q_l < Q_h < 1) > \gamma_l(\frac{y}{x} < 1, Q_l = Q_h = Q < 1),$$

is

$$Q_l\left[a+(1-a)(1+q)(1+y)\right] < (aQ_l + (1-a)Q_h)\left[a+(1-a)(1+y)\right],$$

which after substituting $Q_h = (1+q)Q_l$ simplifies to

$$\left[a+(1-a)(1+q)(1+y)\right] < \left[a+(1-a)(1+q)\right]\left[a+(1-a)(1+y)\right]. \tag{12}$$

It is easy to show that inequality (12) does not hold for all non-negative q and y. ■

This shows that for a type$-l$ household, Social Security's internal rate of return from a progressive benefit-earnings rule is actually lower in the presence of differential mortality, compared to when mortality risk is symmetric and uncorrelated to earnings.

Proposition 4 *The higher mortality risk experienced by a type$-l$ household heavily discounts its marginal utility from higher old-age consumption.*

Proof. See Appendix A. ■

To summarize, the simple two-period model illustrates that a progressive social security program may actually deliver lower welfare for low-income households when mortality risk is negatively correlated with earnings. This is because of two effects: the internal rate of return from a progressive benefit-earnings rule is lower in the presence of differential mortality, and improved work-retirement consumption smoothing is not as useful for low-income households due to their higher mortality risk. Next, I turn to a quantitative examination of these effects in a rich multi-period general-equilibrium environment.

3 The quantitative general-equilibrium model

The unit of the quantitative general-equilibrium model is a permanent-income household that smooths consumption and labor supply over the life cycle by accumulating a risk-free asset: physical capital. Over the course of the life cycle, this household experiences two types of risk: labor income risk, which is exogenous, and mortality risk, which is endogenous and negatively correlated

to household earnings.[7] The household does not have access to markets where it can purchase insurance against these risks.

At each date, surviving households earn labor income if they work, and they also receive interest income from their asset holdings. Households receive Social Security benefits after the full retirement age, and they are also eligible for Supplemental Security Income (SSI) benefits. Firms operate competitively and produce output using capital, labor and a constant returns to scale technology. The government collects taxes on labor and capital income, and uses these revenues to provide public goods and the Supplemental Security Income benefits. Social Security, on the other hand, is funded through a payroll tax on labor income, up to a taxable maximum.

3.1 Preferences

Preferences are standard in the sense that households receive a utility flow every period from a chosen consumption-leisure bundle. The level of this period utility flow is typically irrelevant in most standard model environments. However, in a framework with endogenous mortality risk, the level of the period utility flow has important implications for life-cycle behavior. If period utility is negative, then rational agents will prefer a shorter life and will choose behavior that reduces the survival probabilities. Because of this reason, I follow Hall and Jones (2007) and Zhao (2014) and augment the standard CRRA utility function with a positive constant π_c so that the period utility flow is always positive. A household's period labor supply decision consists of two components: the extensive margin or the participation decision (P), and the intensive margin or the hours of work (h), conditional on participation. With these assumptions, the period utility function is given by

$$u(c, 1-h, P) = \begin{cases} \pi_c + \frac{\left(c^{\eta}(1-h-\theta_P \cdot P)^{1-\eta}\right)^{1-\sigma}}{1-\sigma} & \text{if} \quad \sigma \neq 1 \\ \pi_c + \ln\left(c^{\eta}(1-h-\theta_P \cdot P)^{1-\eta}\right) & \text{if} \quad \sigma = 1 \end{cases} \tag{13}$$

where η is the share of consumption, σ is the inverse of intertemporal elasticity of substitution (IES), θ_P is the age-dependent cost of labor force participation (measured in hours), and P is the labor force participation status: $P = 1$ if the household participates, and $P = 0$ otherwise. Also, since I normalize the period time endowment to unity, $0 \leq h \leq 1$.

Expected lifetime utility from the perspective of a household is given by

$$U = E\left[\sum_{s=0}^{T} \beta^s \ \left\{\Pi_{j=0}^{s-1} Q_j\left(x_j\right)\right\} \ u\left(c_s, 1-h_s, P_s\right)\right], \tag{14}$$

where β is the discount factor, and $Q_j\left(x_j\right)$ is the probability at age j of surviving to the next period, which depends on the age itself, as well as on the household's state vector that age, x_j.

3.2 Income

Conditional on labor force participation, a household earns before-tax wage income $y_s(\varphi_s) = h_s w_s e_s(\varphi_s)$ at age s, where w_s is the wage rate, and $e_s(\varphi_s)$ is a labor productivity endowment that depends on age and a stochastic productivity shock φ_s. This wage income is subject to two separate taxes: a progressive income tax $T_y(\cdot)$, and a payroll tax $T_{ss}(\cdot)$ for Social Security that is

[7] This approach of modeling differential mortality as a negative correlation between earnings and mortality risk follows Fuster et al. (2003) and Cristia (2009). However, some studies have also modeled health as a determinant of mortality risk (İmrohoroğlu and Kitao, 2012; Zhao, 2014). In these studies, unfavorable health shocks reduce life expectancy. In the current context, the exact source of differential mortality is less important; the relevant factor is how differential mortality interacts with the redistribution implicit in Social Security's benefit-earnings rule.

proportional up to the maximum taxable earnings of $\bar{y}$. Households also earn interest income on their asset holdings at the risk free rate r, and this interest income is subject to a capital income tax $T_k(\cdot)$. After-tax income at age s, therefore, is equal to before-tax wage plus interest income minus the income and Social Security taxes.

It is important to note here that because annuity markets are closed, deceased households at every age leave behind accidental bequests. I assume that the government imposes a confiscatory tax on these accidental bequests, which is equivalent to assuming that the government imposes an estate tax of 100%.[8]

3.3 Social Security

The government pays Social Security benefits to households after the full retirement age (T_r), and the amount of benefits paid to a particular household depends on its earnings history. For each household, the government calculates an average of past earnings (up to the maximum taxable earnings), referred to as the Average Indexed Monthly Earnings ($AIME$). The Social Security benefit amount, also called the Primary Insurance Amount (PIA), is calculated as

$$PIA = RR(AIME) \times AIME + \overline{PIA} \tag{15}$$

where RR is the replacement rate, which is a concave function of the $AIME$, and $\overline{PIA}$ is a fixed amount unrelated to the $AIME$.[9] Finally, the government scales the benefit amount up or down proportionally so that Social Security's budget is balanced.[10,11]

3.4 Public Goods and SSI

The government uses the tax revenues from wage and interest income, and also the accidental bequests, to provide public goods and the SSI benefits. The SSI program provides transfers to low-income households that are older than the full retirement age, so that old-age consumption is always above a consumption floor $\underline{c}$. Given its tax revenues and SSI program expenditures, the government adjusts its public goods expenditures so that its budget is balanced.

3.5 A household's optimization problem

A household's state vector is given by $x = \{k, \varphi, AIME\}$, where k denotes the beginning-of-period assets, φ the stochastic productivity shock, and $AIME$ the average past earnings that determine

[8] I later examine the sensitivity of the baseline results with an alternative specification where these bequests are redistributed back to the surviving population in a lump-sum fashion each period.

[9] For the current U.S. benefit-earnings rule, this fixed amount $\overline{PIA}$ is equal to zero.

[10] It is worth noting that because the unit of the current model is a household, I abstract from any intra-household risk-sharing effects of Social Security, which may be significant. For example, it is well known that female life expectancy is higher than male life expectancy, and also that Social Security benefits paid to a household are adjusted if there are surviving dependents. Additionally, Guner et al. (2012) find marked differences in how male and female labor supply responds to tax policies. Taken together, these factors may have important consequences for the insurance effects of Social Security's benefit-earnings rule.

[11] While in reality, Social Security has a trust fund and does not satisfy the definition of a Pay-As-You-Go program in the narrow sense, there is disagreement on whether or not the trust fund assets are "real", i.e. whether or not they have increased national saving. In fact, Smetters (2003) finds that the trust funds assets have actually increased the level of debt held by the public, or reduced national saving. Because of this reason, it is a common practice in the literature to ignore the trust fund and model Social Security's budget as balanced every period (See, for example, studies such as Huggett and Ventura (1999), Conesa and Krueger (1999), İmrohoroğlu et al. (2003), Jeske (2003), Conesa and Garriga (2009), and Zhao (2014), among others).

Social Security benefits. Conditional on a particular realization of the states, the household chooses consumption, assets holdings for the next period, and labor supply.

At a given age s, this optimization problem can be recursively expressed as

$$V_s(x) = \max_{c,k',P,h} \left\{ u(c, 1-h, P) + \beta \; Q_s(x) \; E\left[V_{s+1}(x')\right] \right\} \tag{16}$$

subject to

$$c_s + k' = (1+r)k + y_s^{at} + \Theta(s - T_r)\, b(AIME) \tag{17}$$

$$y_s^{at} = h_s w_s e_s(\varphi_s) + rk - T_y\left(h_s w_s e_s(\varphi_s)\right) - T_k\left(rk\right) - T_{ss}\left(h_s w_s e_s(\varphi_s); \bar{y}\right) \tag{18}$$

$$0 \leq h_s \leq 1, \tag{19}$$

$$k' \geq 0, \tag{20}$$

$$AIME' = \begin{cases} \left[AIME \times (s-1) + \min\left\{h_s w_s e_s(\varphi_s), \bar{y}\right\}\right]/s & s < T_r \\ AIME & s \geq T_r \end{cases}, \tag{21}$$

$$c_s \geq \underline{c} \quad \forall \quad s \geq T_r, \tag{22}$$

where

$$\Theta(s - T_r) = \begin{cases} 0 & s < T_r \\ 1 & s \geq T_r \end{cases}$$

is a step function. Households are born with and die with zero assets, i.e. $k(0) = k(T+1) = 0$, and prior to age T_r, the average earnings $AIME$ evolves based on the realized labor productivity shocks and the endogenous labor supply decisions.

3.6 Technology and factor prices

Output is produced using a Cobb-Douglas production function with inputs capital and labor

$$Y \;=\; K^{\alpha} L^{1-\alpha}, \tag{23}$$

where α is the share of capital in total income. Firms face perfectly competitive factor markets, which implies

$$r = MP_K - \delta = \alpha \left[\frac{K}{L}\right]^{\alpha-1} - \delta \tag{24}$$

$$w = MP_L = (1-\alpha)\left[\frac{K}{L}\right]^{\alpha} \tag{25}$$

where δ is the depreciation rate of physical capital and w is the wage rate. There is no aggregate uncertainty.

3.7 Aggregation

The population structure is as follows: at each instant a new cohort is born and the oldest cohort dies, and cohort size grows at the rate of n over time. Let us denote the measure of households at age s with state x as $\mu_s(x)$. Then, the aggregate capital stock and labor supply are given by

$$K = \sum_{s=0}^{T} N_s \sum_{x} \left\{\Pi_{j=0}^{s-1} Q_j(x)\right\} \; k_{s+1}(x)\mu_s(x) \tag{26}$$

$$L = \sum_{s=0}^{T} N_s \sum_{x} \left\{\Pi_{j=0}^{s-1} Q_j(x)\right\} \; h_s(x) e_s(x) \mu_s(x), \tag{27}$$

where N_s is the size of the age-s cohort. The budget constraints for Social Security and the general government budget are given by

$$\sum_{s=0}^{T} N_s \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q_j(x) \right\} \; T_{ss}\left(h_s(x) w_s e_s(x); \bar{y}\right) \mu_s(x)$$
$$= \sum_{s=0}^{T} N_s \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q_j(x) \right\} \; \Theta(s - T_r) b(x) \mu_s(x) \tag{28}$$

and

$$BEQ + \sum_{s=0}^{T} N_s \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q_j(x) \right\} \; \left(T_y\left(y_s(x)\right) + T_k\left(rk_s(x)\right)\right) \mu_s(x) = G + SSI, \tag{29}$$

where BEQ is the total value of the accidental bequests from deceased households, SSI is the total expenditure of the Supplemental Security Insurance program, and G is the endogenously determined level of public goods expenditures.

3.8 Competitive equilibrium

A competitive equilibrium in this environment is characterized by a collection of

1. cross-sectional consumption $\{c(s;x)\}_{s=0}^{T}$, participation $\{P(s;x)\}_{s=0}^{T}$, and labor hours allocations $\{h(s;x)\}_{s=0}^{T}$,
2. an aggregate capital stock K and labor L,
3. a rate of return r and a wage rate w,
4. Social Security benefits $b(x)$, SSI expenditures, and public-goods expenditures G, and
5. a measure of households $\mu_s(x)$ $\forall$ s,

that

1. solves the households' optimization problems,
2. maximizes the firms' profits,
3. equilibrates the factor markets,
4. balances the government's budgets, and
5. satisfies $\mu_{s+1}(x) = R_\mu\left[\mu_s(x)\right]$, where $R_\mu(\cdot)$ is a one-period transition operator on the measure distribution.

In equilibrium, total expenditures equals consumption plus net investment plus government expenditures, which is equal to the total income earned from capital and labor. Finally, I consider only steady-state equilibria, and I also normalize the initial newborn cohort size to $N(0) = 1$.

	Age groups		
Earnings quintile	35-49	50-64	65-75
Top	0.35	0.61	0.74
Fourth	0.56	0.68	0.94
Third	0.73	0.99	1.08
Second	1.13	1.10	1.14
Bottom	2.25	1.63	1.1

Table 1: Mortality ratios by lifetime earnings from Cristia (2009).

4 Calibration

4.1 Demographics

To set the demographic parameters, I first assume that households enter the model at the actual age of 25 (model age of zero), and they are alive for 75 periods (up to the actual age of 100). Second, I set the population growth rate to $n = 1\%$, which is consistent with the U.S. demographic history and also with the literature. Next, to calibrate the negative correlation between earnings and mortality risk, I use estimates of mortality ratios by earnings in Cristia (2009) to adjust the average age-specific death rates from the 2001 U.S. Life Tables (Arias, 2004).[12]

Cristia (2009) uses data from the 1984, 1993, 1996, and 2001 panels of the Survey of Income and Program Participation (SIPP), also matched with several files administered by the Social Security Administration (SSA), to compute the mortality ratios by income groups. I report these ratios in Table 1. Each entry in Table 1 gives the likelihood of death for the respective age-income group relative to the population average for that group. For example, an individual between ages 35-49 in the bottom earnings quintile is 2.25 times more likely to die, relative to the population average in that age range (conditional on surviving up to age 35).

To use this data, I first extrapolate the mortality ratios for ages 25-35 and ages 75-100, and then I calculate the age-specific death rates for each income quintile in the model by multiplying the mortality ratios by the average age-specific death rates in the population.

4.2 Social Security

To calibrate Social Security in the model, I first set the payroll tax function to

$$T_{ss}(y;\bar{y}) = \begin{cases} \tau_{ss} y & y \leq \bar{y} \\ \tau_{ss}\bar{y} & y > \bar{y} \end{cases}$$

and then set the tax rate to $\tau_{ss} = 0.106$, which is the combined tax rate for the Old-Age and Survivors Insurance (OASI) component. The maximum taxable earnings ($\bar{y}$) is adjusted regularly relative to the average wage in the U.S. Huggett and Ventura (1999) calculate that the ratio of this taxable maximum to the average wage index has averaged at about 2.47 in the U.S., using which I set the maximum taxable earnings in the model to $\bar{y} = 2.47$.

Second, to compute the PIA amount, I incorporate the U.S. benefit-earnings rule into the model. This requires setting $\overline{PIA} = 0$ in equation (15), and then specifying the replacement rate (RR) as a piecewise linear function of past work-life income, the $AIME$. Huggett and Ventura (1999) estimate the "bend-points" of this piecewise linear function to be roughly 20%, 124%, and

[12] Later I examine the sensitivity of the baseline results using estimates of differential mortality from Bosworth et al. (2016).

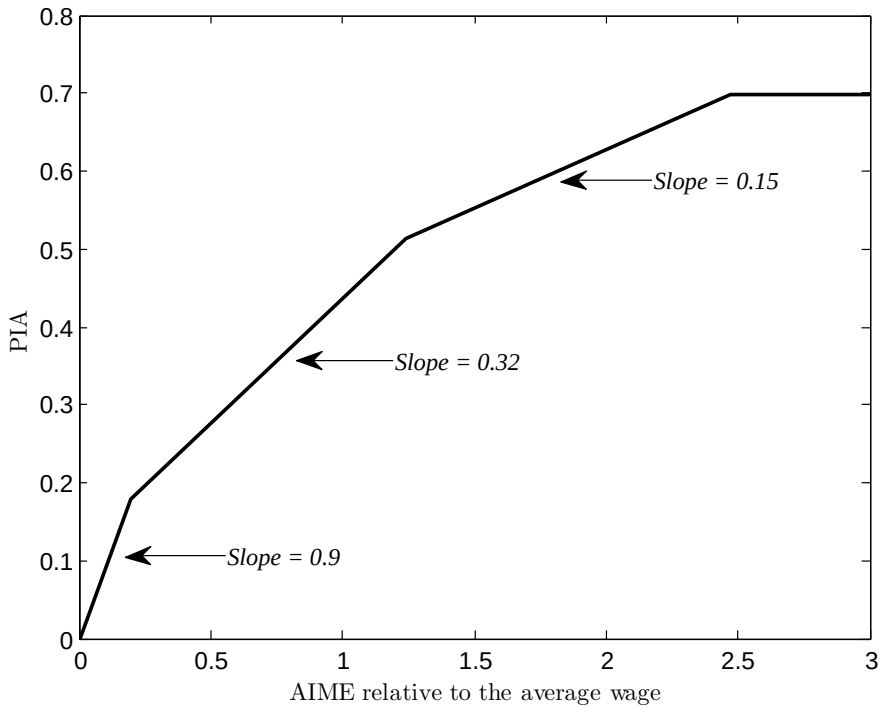

Figure 1: Benefit formula in the U.S.

247% of the average wage in the economy, and I take these values directly to the model. It is worth noting that the progressivity in the benefit-earnings rule is captured by the fact that the "replacement rate" is decreasing in the $AIME$ (Figure 1).

Finally, I assume that households receive Social Security benefits in the model after age $T_r = 41$, which corresponds to the current full retirement age of 66 in the U.S.

4.3 Labor productivity endowment

To calibrate the labor income process, I assume that the log of labor productivity at age s can be additively decomposed as

$$\log e_s(\varphi_s) = \epsilon_s + \varphi_s, \tag{30}$$

where ϵ_s is a deterministic age-dependent component, and φ_s is a stochastic component, given by

$$\varphi_s = p + z_s + \nu_s \tag{31}$$

$$z_s = \rho z_{s-1} + \upsilon_s, \tag{32}$$

where $p \sim N(0, \sigma_p^2)$ is a permanent productivity fixed effect realized at birth, $\nu_s \sim N(0, \sigma_\nu^2)$ is a transitory shock, and z_s is a persistent shock that follows a first-order autoregressive process with $z_0 = 0$, persistence ρ, and a white-noise disturbance $\upsilon_s \sim N(0, \sigma_\upsilon^2)$.

I parameterize ϵ_s using the estimates from Kitao (2014), who uses work hour and wage data from the 2007 PSID to derive this age-dependent component of productivity as a residual of wages, after accounting for hours worked and also the part-time wage penalty. To calibrate the stochastic component, I use estimates from Heathcote et al. (2010) and set the persistence parameter to $\rho = 0.973$, the variances of the permanent fixed effect and the transitory shock to $\sigma_p^2 = 0.124$ and

$\sigma_\nu^2 = 0.04$ respectively, and variance of the white-noise disturbance to $\sigma_v^2 = 0.018$. I use Gaussian quadrature to approximate the distribution of the permanent fixed effect using a three-point discrete distribution, and I approximate the joint distribution of the persistent and transitory shocks using a five-state first-order discrete Markov process following Tauchen and Hussey (1991).

4.4 Income tax

To calibrate the labor income tax function, I follow Storesletten et al. (2012) and Karabarbounis (2012) and assume that

$$T_y(y) = y - (1 - \tau_y)y^{1-\tau_1}, \tag{33}$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (33) reduces to a proportional tax function with a marginal rate of τ_y. With this income tax function, after-tax income is log-linear in before-tax income, and the parameter τ_1 controls the progressivity of the tax code. I set $\tau_1 = 0.151$ following the Storesletten et al. (2012) estimate of this parameter using data from the 2000, 2002, 2004, and 2006 waves of the PSID, and also NBER's TAXSIM program, accounting for federal and state income taxes plus public transfers.

Capital income tax in the U.S. consists of taxes on interest income, and also on capital gains. However, there are no capital gains in current model because it has only one asset and there is no aggregate uncertainty. Therefore, to parameterize the capital income tax function, I follow the literature and assume that $T_k(\cdot) = \tau_k \times rk$ (De Nardi et al., 1999; İmrohoroğlu and Kitao, 2012).

4.5 Technology

The historically observed value of capital's share in total income in U.S. ranges between 30-40%, so I set $\alpha = 0.35$. Finally, I calibrate the depreciation rate δ to get an investment-to-output ratio of 21% in equilibrium.

4.6 Unobservable parameters

Once all the observable parameters have been assigned empirically reasonable values, I jointly calibrate the remaining unobservable parameters of the model, i.e. the preference parameters π_c, σ, β, and η, the age-dependent labor force participation cost $\theta_P(s)$, and also the government parameters τ_y, τ_k, and $\underline{c}$ to match certain macroeconomic targets.

First, so that overall wealth accumulation in the model matches the U.S. economy, I fix the IES to $\sigma = 4$ and then calibrate the discount factor (β) to target an equilibrium capital-output ratio of 3.5. Second, two salient features of cross-sectional labor supply data in the U.S. are (i) a rapid decline in the labor force participation rate from about 90% to almost 30% between ages 55 to 70, and (ii) an average of 40 hours per week per worker spent on market work between ages 25 to 55 (Kitao, 2014). I adopt both of these empirical facts as targets.

Also, following Kitao (2014), I assume that the labor force participation cost increases with age based on the relationship

$$\theta_P(s) = \kappa_1 + \kappa_2 s^{\kappa_3},$$

where s is model age, and then parameterize κ_1, κ_2, and κ_3 to match the observed rapid decline in labor force participation after age 55. The consumption share parameter (η) controls the fraction of time a household spends on market work (conditional of participation), so I calibrate it to match the hours per week target. To calibrate the additive term in the utility function π_c, I follow Zhao

π_c	σ	β	η	κ_1	κ_2	κ_3	τ_y	τ_k	$\underline{c}$	δ
18.0	4	0.997	0.41	0.3186	1.788	5.48	0.126	0.1	\$28,600	0.053

Table 2: Unobservable parameter values under the baseline calibration.

	Target	Model
Capital-output ratio	3.5	3.34
Investment-output ratio	0.21	0.21
Avg. hours of market work per week per worker (25-55)	40	38
Share of govt. expenditures in GDP	0.2	0.21
Ratio of SSI expenditures to GDP	0.003	0.0036
Value of Statistical Life	\$4 m.–\$10 m.	\$9 m.
Interest rate	–	0.0517
Benefit adjustment factor	–	0.667

Table 3: Key macroeconomic variables under the baseline calibration.

(2014) and adopt a Value of Statistical Life (VSL) target of \$4 to \$10 million to be consistent with empirical estimates (Viscusi and Aldy, 2003).[13]

I calibrate the tax parameters τ_y and τ_k such that the model generates a ratio of government expenditures to GDP of 20% in equilibrium. This step ensures that the scale of tax revenues relative to GDP in the model is consistent with that in the U.S. economy and also the literature. Finally, I calibrate the old-age consumption floor $\underline{c}$ so that total SSI expenditures are 0.3% of GDP.

5 The baseline economy

The unobservable parameter values under which the baseline equilibrium reasonably matches the above targets are reported in Table 2. Note that with leisure in period utility, the relevant inverse elasticity for consumption is $\sigma^c = 1 + \eta(\sigma - 1) = 2.23$, which lies within the range frequently encountered in the literature. Also, with the above values of κ_1, κ_2, and κ_3, the labor force participation cost increases at a faster rate with age.

The model-generated values for key macroeconomic variables under the baseline calibration are reported in Table 3 along with their targets, and the cross-sectional labor force participation and labor hours data (conditional on participation) are reported in Figures 2 and 3. Note that the benefit adjustment factor in the baseline calibration is smaller than unity, which implies that the PIAs obtained from the U.S. benefit-earnings rule are adjusted downward to balance Social Security's budget in the baseline equilibrium. It is clear from the figures that the baseline calibration does a good job of matching the rapid decline in labor force participation at older ages quite well, and it also reasonably matches the general declining trend of weekly hours over the life cycle. However, the model fails to replicate the mild-hump shape in the hours profile observed in the data. One way to potentially improve the model's fit along this dimensions is to treat the age-dependent component of labor productivity ϵ_s as unobservable. Treating ϵ_s as an unobservable parameter would potentially eliminate any selection bias arising from measuring it as residual wages (Bullard and Feigenbaum, 2007; Bagchi and Feigenbaum, 2014).

Finally, it is worthwhile examining the correlation between earnings and mortality risk in the

[13] I measure VSL in the current model as the average willingness to pay to survive an additional year between ages 30-55.

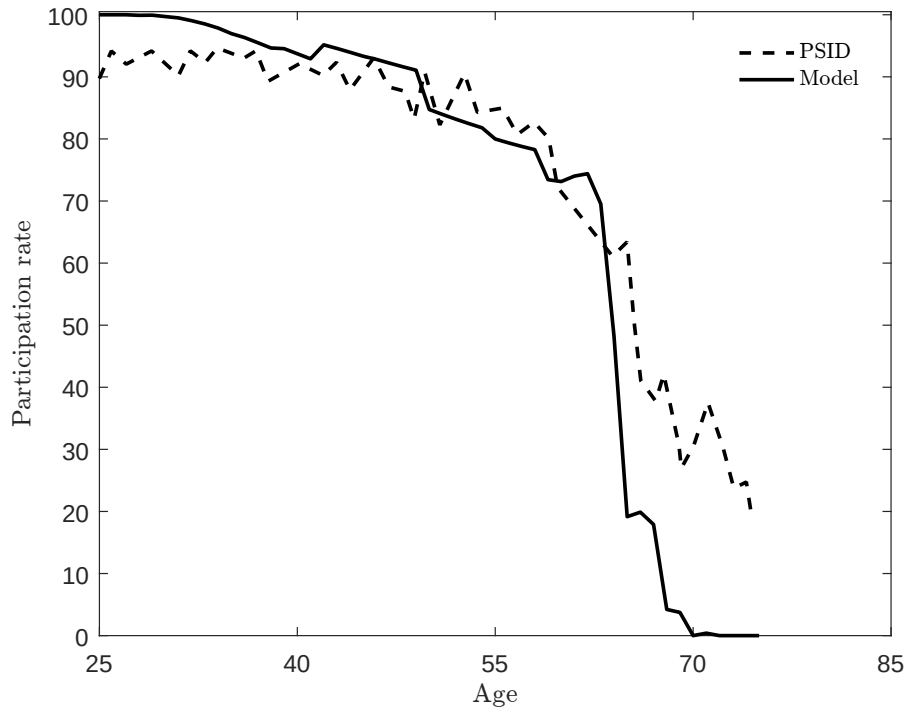

Figure 2: Cross-sectional labor force participation rates under the baseline calibration.

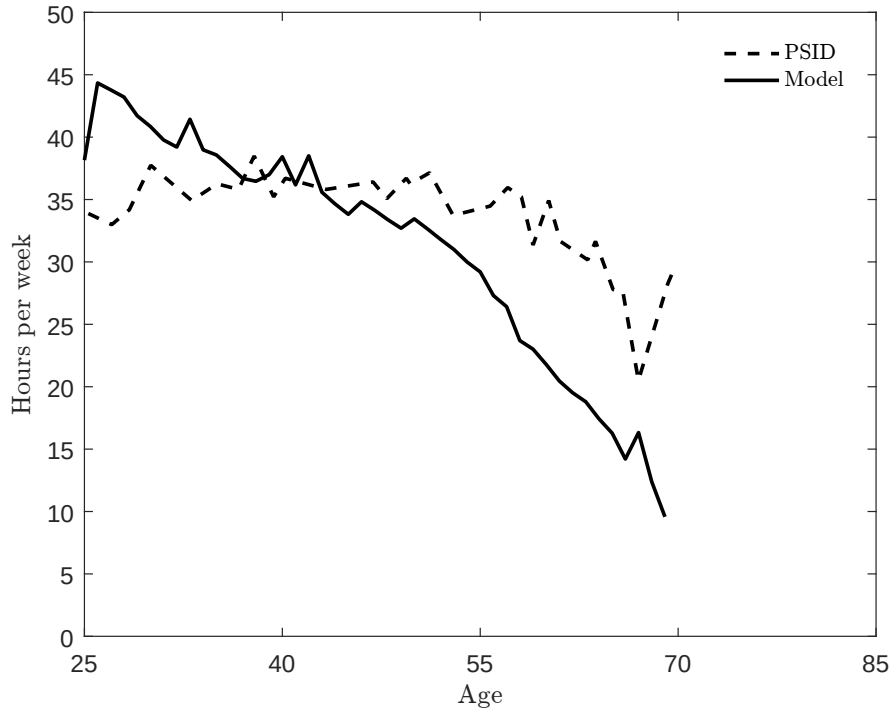

Figure 3: Cross-sectional mean of labor hours per week (conditional on participation) under the baseline calibration.

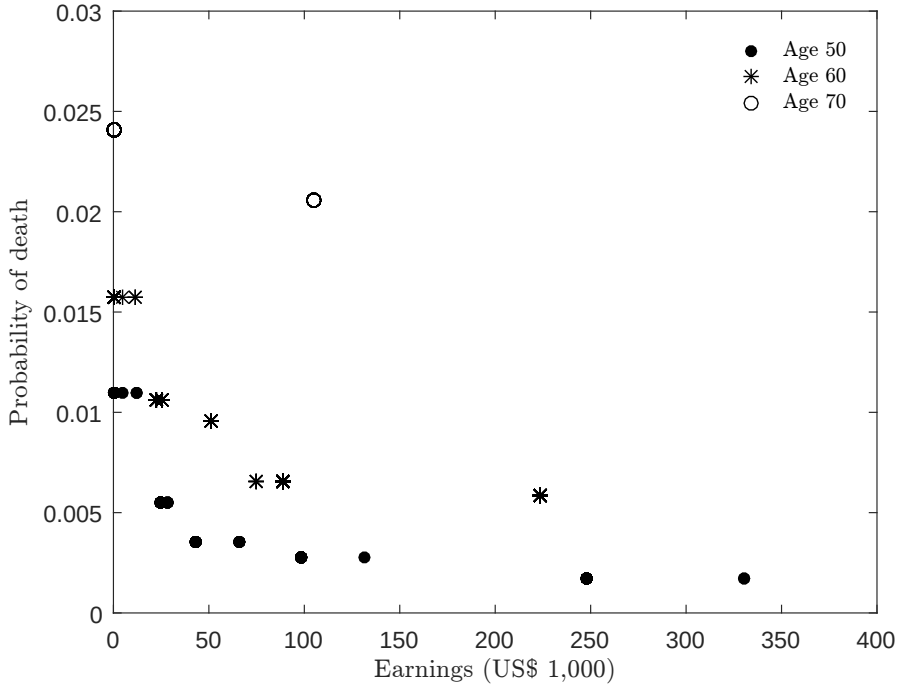

Figure 4: Age-specific death rates as a function of earnings in the baseline calibration.

baseline calibration. In Figure 4, I report the death rates at ages 50, 60, and 70 as a function of earnings. Three facts are clear from the figure. First, as one would expect, mortality risk is positively related to age: death rates experienced by 70-year-olds are larger than those experienced by 60-year olds, which in turn, are almost always larger than those for the 50-year olds. Second, at a given age, households at the bottom of the earnings distribution experience higher death rates. For example, a 60-year-old household with earnings less that $50,000 is almost twice as likely to die, compared to a 60-year-old household with earnings between $50,000–$100,000. This pattern remains unchanged for younger as well as older households. Finally, the negative correlation between earnings and mortality risk in the baseline calibration is strong enough to cause some cross-over in survivorship, in the sense that death rates of 60-year olds earning $50,000–$100,000 are lower than those experienced by 50-year olds at the bottom of the earnings distribution.

6 The experiments

This paper examines if differential mortality has any implications for the welfare effects of Social Security's benefit-earnings rule. To do this, one must compute new equilibria of the baseline model with alternative benefit-earnings rules, while holding all the other institutional features of Social Security fixed at their baseline level, under two scenarios: with and without differential mortality risk. At this point, two important choices must be made. First, what kind of alternative benefit-earnings rules should be considered in the computational experiments? Second, what type of welfare measures ought to be used in evaluating these alternative benefit-earnings rules?

First, while it is certainly possible to investigate the globally optimal structure of Social Security benefits in a model such as this, I focus only on benefit-earnings rules that are structurally similar

to the current U.S. rule.[14] As discussed earlier, Social Security's current benefit-earnings rule is piecewise linear, with a replacement rate of 90% of the AIME for the first 20% of the average wage, 32% of the AIME for the next 104% of the average wage, and 15% of the AIME for the remaining, up to the taxable maximum of 247% of the average wage in the economy (Figure 1). I henceforth refer to these replacement rates as the primary, secondary, and the tertiary replacement rates, respectively.

Given this structure, a convenient way to vary the progressivity of the benefit-earnings rule is to vary only the primary replacement rate in the computational experiments, while holding the secondary and the tertiary rates constant at their current values. This way, reducing the primary replacement rate makes the benefit-earnings rule more proportional (or linear), which has the effect of reducing Social Security's implicit insurance. On the other hand, increasing the primary replacement rate makes the benefit-earnings rule more progressive (or concave), which has the effect of increasing Social Security's implicit insurance. I consider benefit-earnings rules ranging from a linear function of past earnings (i.e. with zero redistribution), to a fixed benefit that is uniform across retirees and completely unrelated to past earnings (i.e. with full redistribution).[15]

It is important to note here that even though I focus only on the primary replacement rate in the computational experiments, keeping Social Security's budget balanced with the current payroll tax rate and the taxable maximum requires the secondary and tertiary replacement rates to be adjusted as well. For example, increasing the primary replacement rate from 90%, while keeping the secondary and tertiary rates fixed at 32% and 15% respectively, leads to an overall increase in Social Security benefits. Therefore, so that Social Security can achieve Pay-As-You-Go balance with the current tax rate and taxable maximum, the secondary and tertiary replacement rates must be reduced. This is accomplished automatically in the model through the benefit adjustment factor, which adjusts the PIA obtained from every alternative benefit-earnings rule to ensure that Social Security's budget is balanced in equilibrium. In fact, as we will see, this adjustment factor even offsets some of the direct change in the primary replacement rate, in addition to the secondary and tertiary rates. This approach has the merit of allowing for the cleanest interpretation of the results, especially from a policy-making perspective, because the benefit-earnings rule fundamentally alters the progressivity in Social Security without altering the overall size of the program.

Second, to evaluate the welfare implications of the alternative benefit-earnings rules, I define the following two measures. To understand the overall welfare consequences, I follow the literature and define

$$W = \sum_{s=0}^{T} \beta^s \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q_j(x) \right\} \; u\left(c_s(x), 1 - h_s(x), P_s(x)\right) \mu_s(x) \tag{34}$$

which is the ex-ante expected lifetime utility of a newborn household. Then, to understand the distributional consequences of these benefit-earnings rules, I define a consumption equivalence ψ

[14] Note that in general, the current model is not suitable for studying the globally optimal Social Security tax and benefit structure. This is because in a calibrated rational-agent model such as this, the distortions from Social Security are typically larger than its insurance effects (Hubbard and Judd, 1987; İmrohoroğlu et al., 1995; Nishiyama and Smetters, 2008; Bagchi, 2015). As a result, the globally optimal Social Security tax and benefit structure in a model such as the current one would trivially warrant zero Social Security.

[15] In terms of equation (15), a fixed benefit that is completely unrelated to past earnings is obtained by setting the primary replacement rate equal to zero and allowing $\overline{PIA} > 0$.

for each realization of the permanent productivity fixed effect (p) that solves

$$E\left[\sum_{s=0}^{T}\beta^{s}\ \left\{\Pi_{j=0}^{s-1}Q_{j}(x^{C})\right\}\ u\left((1+\psi)\,c_{s}^{C}(x^{C}),1-h_{s}^{C}(x^{C}),P_{s}^{C}(x^{C})\right)\right]=$$
$$E\left[\sum_{s=0}^{T}\beta^{s}\ \left\{\Pi_{j=0}^{s-1}Q_{j}(x^{H})\right\}\ u\left(c_{s}^{H}(x^{H}),1-h_{s}^{H}(x^{H}),P_{s}^{H}(x^{H})\right)\right], \tag{35}$$

where C denotes current Social Security law, and H denotes a hypothetical Social Security law with the alternative benefit-earnings rule. Intuitively, this consumption equivalence captures the welfare gains (or losses) in units of consumption, as a function of the permanent productivity fixed effect, under each one of the computations. Taken together, these two measures provide an overall, as well as a disaggregated picture of the welfare implications of the alternative benefit-earnings rules.[16]

7 Welfare consequences with differential mortality

A good benchmark for examining the welfare consequences of alternative Social Security benefits-earnings rules in a general-equilibrium life-cycle economy is Nishiyama and Smetters (2008). In this study, the authors examine the optimal Social Security benefit structure in an overlapping-generations macroeconomic model with labor income and mortality risk, missing annuity markets, and borrowing constraints. Calibrating the model to match some key features of the U.S. economy, Nishiyama and Smetters (2008) find that the optimal benefit-earnings rule is fairly proportional (or linear), with a strong link between benefits and past work-life income. They argue that Social Security's relatively long averaging period already provides some insurance against negative labor income shocks, but in a manner that is more efficient than explicit redistribution through the progressive benefit-earnings rule. This is because while the progressivity in the benefit structure provides insurance against labor income risks that are difficult to insure privately, it also introduces implicit tax rates that distort labor supply. Nishiyama and Smetters (2008) find that the welfare losses from these distortions outweigh the welfare gains from the increased insurance.

I report the welfare consequences of several alternative benefit-earnings rules from the baseline model with differential mortality risk in Table 4. In the first column, I report the primary, secondary, and the tertiary replacement rates of the benefit-earnings rule being examined, and in the second column, I report the corresponding adjustment factor needed to balance Social Security's budget, given the current payroll tax rate and taxable maximum. I combine these two statistics to calculate the "effective" replacement rates in the third column, and in the last column I report overall welfare.[17]

The following two facts are clear from Table 4. First, the welfare effects of varying the progressivity of Social Security's benefit-earnings rule are quite small in equilibrium. This is primarily because modifying the primary replacement rate of the benefit-earnings rule requires adjusting the secondary and tertiary replacement rates, so that Social Security's budget is balanced with the current payroll tax rate and taxable maximum. The benefit-earnings rule adjustment factor declines when the primary replacement rate is increased, and increases when it is reduced from its baseline level. In addition to offsetting some of the direct change in primary replacement rate, this leads to

[16] I later examine how differential mortality affects the welfare ranking of the benefit-earnings rules under a Rawlsian "maximin" welfare criterion.

[17] For the fully proportional (or linear) benefit-earnings rule, I set the primary, secondary, and the tertiary replacement rates equal to each other.

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	0.339	0.2/0.2/0.2	596.729
0.45/0.32/0.15	0.767	0.35/0.25/0.12	597.435
Baseline (0.9/0.32/0.15)	0.667	0.6/0.21/0.1	597.160
1.8/0.32/0.15	0.525	0.94/0.17/0.08	596.976
Uniform benefits	–	–	596.152

Table 4: Replacement rates and aggregate welfare under alternative benefit-earnings rules with differential mortality risk.

Permanent productivity fixed effect (p)	0.54	1.00	1.84
Proportional/linear	−0.261	−0.064	−0.072
0.45/0.32/0.15	0.079	0.092	0.092
1.8/0.32/0.15	−0.027	−0.054	0
Uniform benefits	−0.039	−0.48	−0.045

Table 5: Consumption equivalences (ψ%) under alternative benefit-earnings rules with differential mortality risk.

a consistent flattening of the benefit-earnings rule. Second, with differential mortality risk, ex-ante expected utility is maximized with the replacement rates 0.45/0.32/0.15, which is less progressive than the current U.S. rule. Increasing the primary replacement rate from its baseline value of 90% to 180% consistently reduces overall welfare. Therefore, similar to Nishiyama and Smetters (2008), I find that the insurance effects of more progressive benefit-earnings rules are weak, even in the presence of differential mortality.

To assess the distribution of welfare gains and losses under the alternative benefit-earnings rules, I report in Table 5 the consumption equivalence (ψ) for each realization of the permanent productivity fixed effect (p) under each computation. Although these values are small, it is clear from the table that increasing the degree of progressivity from the baseline hurts almost all households, but decreasing it slightly leads to welfare gains for everyone. Adopting the replacement rates 0.45/0.32/0.15 leads to a welfare gain equivalent to an increase of 0.08% in period consumption for households with $p = 0.54$, and of almost 0.1% for households with $p = 1.0$ and 1.84. Not surprisingly, the welfare gains from this policy are the smallest for households at the bottom of the income distribution, because this is still a slight reduction in Social Security's implicit redistribution. For the same reason, households with $p = 0.54$ suffer the largest welfare losses with the proportional (linear) benefit rule.

To summarize, consistent with Nishiyama and Smetters (2008), I find that in a general-equilibrium environment with uninsurable labor income and differential mortality risk, the insurance effects of Social Security's benefit-earnings rule are not strong enough: a benefit-earnings rule less progressive than the current U.S. rule maximizes aggregate welfare.

8 Eliminating differential mortality

Based on our simple two-period model outlined earlier, a progressive benefit-earnings rule reduces Social Security's rate of return for households with relatively low earnings in the presence of differential mortality. Moreover, due to their higher mortality risk, these households also heavily discount the marginal utility of better work-retirement consumption smoothing. The results in the previous section demonstrate the combined effect of these two mechanisms: with Social Security's

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	0.352	0.21/0.21/0.21	609.339
0.45/0.32/0.15	0.755	0.34/0.24/0.11	609.488
Baseline (0.9/0.32/0.15)	0.654	0.59/0.21/0.1	609.510
1.8/0.32/0.15	0.507	0.91/0.16/0.08	609.543
Uniform benefits	–	–	609.438

Table 6: Replacement rates and aggregate welfare under alternative benefit-earnings rules without differential mortality risk.

current payroll tax rate and taxable maximum, a slightly less progressive benefit-earnings rule delivers higher welfare than Social Security's current benefit formula. In this section, I examine the key question of this paper: how important is the role of differential mortality in the welfare ranking of these benefit-earnings rules?

To do this, I compute a hypothetical version of the baseline model with all the observable and structural parameters held fixed at their initial values, but without differential mortality risk. Specifically, I adopt the average age-specific death rates from the 2001 U.S. Life Tables in Arias (2004) to generate the survivor functions for all households under this experiment. However, two adjustments are needed in the model when differential mortality is eliminated. First, adopting the average age-specific death rates for all households leads to an increase in the overall life expectancy. Given that the period utility flow is positive, this has a spurious welfare effect. To control for this effect, I rescale the death rates from the 2001 U.S. Life Tables so that life expectancy without differential mortality is identical to that in the baseline model with differential mortality. Second, I adjust the consumption floor for the SSI program so that equilibrium SSI expenditures as a percentage of GDP without differential mortality is identical to that in the baseline model with differential mortality. This ensures that the overall magnitude of redistribution *outside* Social Security is constant across the two models.

It is worthwhile at this point to note an important effect of the elimination of differential mortality from the baseline model. With this modification, households at the bottom of the earnings distribution experience a longevity improvement, whereas households at the top experience a longevity reduction. These longevity effects induce households with unfavorable productivity shocks to increase their life-cycle labor supply relative to the baseline model, and households with favorable productivity shocks to reduce their labor supply. As a result, in this counterfactual model, households with relatively unfavorable productivity shocks actually end up with slightly better earnings histories on the average. Because of these slightly better earnings histories, households with $p = 0.54$, on the average, receive higher benefits than those with $p = 1.0$ and 1.84 in this counterfactual model. As we will see, this has important distributional consequences for the insurance effects of Social Security's benefit-earnings rule.

The next step is to compute the welfare implications of varying Social Security'e benefit-earnings rule with this counterfactual model (Table 6). It is clear from the table that modifying Social Security's benefit-earnings rule has vastly different welfare implications in the absence of differential mortality. In this case, overall welfare is highest under the replacement rates 1.8/0.32/0.15, which is considerably more progressive than the current U.S. rule. Moreover, welfare decreases monotonically when benefit-earnings rules less progressive than the current U.S. rule are adopted. In other words, elimination of differential mortality leads to a significant reordering of the welfare ranking of the benefit-earnings rules.

The distributional consequences of the alternative benefit-earnings rules in the absence of differential mortality are reported in Table 7. The table shows that similar to the baseline model, the

Permanent productivity fixed effect (p)	0.54	1.00	1.84
Proportional/linear	0.022	−0.03	−0.032
0.45/0.32/0.15	0	−0.011	−0.011
1.8/0.32/0.15	0.007	0.012	0.014
Uniform benefits	−0.062	−0.007	0

Table 7: Consumption equivalences ($\psi\%$) under alternative benefit-earnings rules without differential mortality risk.

	K	L	GDP	r	G/Y
Proportional/linear (0.6/0.6/0.6)	0.993	1.001	0.998	1.013	0.999
0.45/0.32/0.15	1.010	1.004	1.006	0.996	1.002
1.8/0.32/0.15	1.005	1.003	1.004	1.001	1.004
Uniform benefits	1.009	1.002	1.004	0.998	1.006

Table 8: Select macroeconomic variables (relative to the baseline) under alternative benefit-earnings rules with differential mortality risk.

$\psi-$values are quite small, but as expected, benefit-earnings rules more progressive than the current U.S. rule usually generate welfare gains for households with less favorable earnings histories, i.e. those with $p = 1.0$ and 1.84 in this counterfactual model. Similarly, benefit-earnings rules less progressive than the current U.S. rule generate welfare losses for these households.

In the context of our simple two-period model, the absence of differential mortality improves Social Security's internal rate of return for households with lower earnings, and also increases their expected utility from improved work-retirement consumption smoothing. The above results suggest that their combined effect is quantitatively important: eliminating differential mortality from the baseline model causes a benefit-earnings rule more progressive than the current U.S. rule to maximize welfare.

So far, we have focused only on the welfare effects of the alternative benefit-earnings rules. I now turn to the macroeconomic effects of these experiments. In Table 8, I report key macroeconomic variables such as aggregate capital, labor, national income, the interest rate, and the share of government expenditures in GDP, relative to the baseline with differential mortality risk. The table shows that overall, reducing the progressivity of the benefit-earnings rule does not appear to have significant general equilibrium effects: the changes in relevant macroeconomic variables are within a single percentage point, and this result continues to hold even without differential mortality. These results, however, are in stark contrast with Nishiyama and Smetters (2008), who find significant macroeconomic effects of changes in the benefit-earnings rule.

A potential reason behind this discrepancy is the size of labor supply distortions in their experiments. While considering alternative benefit-earnings rules, Nishiyama and Smetters (2008) also adjust the labor income tax rate proportionally to maintain a given level of government expenditures outside Social Security in their model. However, the labor supply distortions from such an experiment are a combined effect of a changing benefit-earnings rule, as well as a changing labor income tax rate. Because Nishiyama and Smetters (2008) require higher labor income tax rates to maintain a given level of government expenditures under more progressive benefit-earnings rules, their experiment distorts labor supply by significantly more than what is caused by the changing benefit-earnings rule.

To summarize, I find that the welfare ranking of the alternative benefit-earnings rules is, in fact, sensitive to differential mortality. A more progressive benefit-earnings rule offers a lower rate of return from Social Security for households with relatively unfavorable earnings histories, and

	π_c	σ	β	η	κ_1	κ_2	κ_3	τ_y	τ_k	$\underline{c}$	δ
Initial baseline	18.0	4	0.997	0.41	0.3186	1.788	5.48	0.126	0.1	\$28,600	0.053
Accidental bequests	18.0	4	1.057	0.4	0.3186	1.788	5.48	0.131	0.3	\$30,250	0.058

Table 9: Unobservable parameter values under the baseline calibration and with accidental bequests.

their higher mortality risk also causes them to heavily discount the utility from higher old-age consumption. I find that these effects are quantitatively important: in the presence of differential mortality, aggregate welfare is maximized with a benefit-earnings rule *less* progressive than the current U.S. rule. Because these effects are absent when differential mortality is eliminated from the baseline model, welfare is maximized with a benefit-earnings rule *more* progressive than the current U.S. rule. I also find that these policy experiments do not have significant effects on other macroeconomic variables: capital, labor, and national income do not change by more than a single percentage point.

9 The importance of accidental bequests

As mentioned earlier, one consequence of missing annuity markets in a life-cycle framework with mortality risk is that households leave behind accidental bequests. Because households expect to survive to the next period with some probability, they accumulate assets to finance future consumption. However, some households do not survive to the next period, and their assets give rise to these accidental bequests. In the baseline model, I assume that these accidental bequests are taxed away by the government, or that they are subjected to an estate tax of 100%. How these accidental bequests are treated, however, can have important consequences for the insurance effects of Social Security. An often-used assumption in the literature is that these accidental bequests are redistributed back to the surviving population in a lump-sum fashion each period (De Nardi et al., 1999; Fehr et al., 2013; Kitao, 2014). As Caliendo et al. (2014) have shown, if one accounts for how higher mandatory saving through Social Security crowds out life-cycle saving, or these accidental bequests, then Social Security has zero effect on life-cycle wealth. In other words, the accidental bequest can itself create a redistribution mechanism, the insurance effects of which can completely undo the insurance effects of Social Security. Therefore, in this section, I examine the sensitivity of the quantitative predictions of the baseline model with respect to how accidental bequests from the deceased households are treated within the model.

To do this, I compute a modified version of the baseline model in which the accidental bequests, rather than being subject to the 100% estate tax, are redistributed back to the surviving population in a lump-sum fashion each period. With this modification, I calibrate the model to the same targets as before: a capital-output ratio of 3.5, an average of 40 hours per week per worker, empirical life-cycle profiles for labor force participation and weekly hours, a VSL consistent with empirical estimates, an investment-to-output ratio of 21%, and a ratio of government expenditures to GDP of 20%. I keep all the observable parameters unchanged at their initial baseline values. I report the values for the unobservable parameters for which the model closely matches the targets in Table 9, and I report the corresponding model performance in Table 10.

Three facts are clear from Table 9. First, the accidental bequests crowd out life-cycle saving, as a result of which one needs a discount factor larger than unity (which effectively implies a negative discount rate) to get realistic capital-output ratios in equilibrium. Second, the elimination of the estate tax requires a significantly larger capital tax rate to yield a realistic level of government expenditures share in GDP. Finally, even with accidental bequests, Social Security benefits need

	Target	Initial baseline	Accidental bequests
Capital-output ratio	3.5	3.34	3.10
Investment-output ratio	0.21	0.21	0.21
Avg. hours of market work per week per worker (25-55)	40	38	45
Share of govt. expenditures in GDP	0.2	0.21	0.19
Ratio of SSI expenditures to GDP	0.003	0.0036	0.0036
Value of Statistical Life	\$4 m.–\$10 m.	\$9 m.	\$8 m.
Interest rate	–	0.0517	0.056
Benefit adjustment factor	–	0.667	0.627

Table 10: Key macroeconomic variables under the baseline calibration and with accidental bequests.

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	0.334	0.2/0.2/0.2	6089.259
0.45/0.32/0.15	0.726	0.33/0.23/0.11	6087.048
Baseline (0.9/0.32/0.15)	0.627	0.56/0.2/0.09	6086.649
1.8/0.32/0.15	0.485	0.87/0.16/0.07	6085.374
Uniform benefits	–	–	6085.083

Table 11: Replacement rates and aggregate welfare under alternative benefit-earnings rules with differential mortality risk and accidental bequests.

to be scaled down, indicated by a benefit adjustment factor of less than unity, so that Social Security's budget constraint is satisfied in equilibrium with the current U.S. benefit-earnings rule and the model's calibrated labor income process.

Next, I evaluate the overall welfare consequences of modifying Social Security's benefit-earnings rule with this modified model, both with and without differential mortality. First, I report the results with differential mortality in Table 11. The table shows that with accidental bequests, aggregate welfare under differential mortality is maximized when benefits are a fully proportional (or linear) function of past earnings. In fact, welfare increases (decreases) monotonically as benefit-earnings rules less (more) progressive than the current U.S. rule are adopted. This should not be surprising, as the accidental bequest provides an additional redistribution mechanism, as a result of which Social Security's implicit redistribution is not as useful.

However, it is important to note that the overall quantitative results under differential mortality are very similar to those from the initial baseline model with the 100% estate tax. In either case, a benefit-earnings rule *less* progressive than the current U.S. rule maximizes welfare under realistic labor income and differential mortality risks. Therefore, I find that regardless of how accidental bequests are treated within the model, the insurance effects of Social Security's benefit-earnings rule are not strong enough to warrant a benefit-earnings rule that is more progressive than the current U.S. rule.

To examine the importance of differential mortality with respect to the welfare effects of Social Security's benefit-earnings rule, I now take the modified model with accidental bequests and then eliminate the positive correlation between earnings and longevity from it. As before, I make two adjustments in this process. First, I take the 2001 U.S. Life Tables in Arias (2004), and then rescale the death rates so that life expectancy is constant across the models with and without differential mortality. I use these re-scaled death rates to generate the survivor functions for all households under this experiment. Second, I adjust the old-age consumption floor so that the scale of redistribution *outside* Social Security is constant, in the sense that the share of SSI expenditures

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	1.315	0.2/0.2/0.2	6242.281
0.45/0.32/0.15	0.726	0.33/0.23/0.11	6242.392
Baseline (0.9/0.32/0.15)	0.622	0.56/0.2/0.09	6243.202
1.8/0.32/0.15	0.475	0.86/0.15/0.07	6244.290
Uniform benefits	–	–	6242.819

Table 12: Replacement rates and aggregate welfare under alternative benefit-earnings rules without differential mortality risk and with accidental bequests.

in GDP is identical under the two models. With these adjustments, the aggregate welfare effects of Social Security's benefit-earnings rule are reported in Table 12.

It is clear from the table that even with accidental bequests, modifying Social Security's benefit-earnings rule has very different welfare implications when differential mortality is eliminated. As before, overall welfare is maximized under the replacement rates 1.8/0.32/0.15, which is considerably more progressive than the current U.S. rule, and welfare decreases monotonically with benefit-earnings rules less progressive than the current U.S. rule. Therefore, eliminating differential mortality causes a significant reordering of the welfare ranking of the benefit-earnings rules even when accidental bequests are redistributed back to the surviving population.

10 Alternative estimates of differential mortality

In the baseline calibration, I use estimates from Cristia (2009) to calibrate the negative correlation between earnings and mortality risk. Cristia (2009) uses data from the 1984, 1993, 1996, and 2001 panels of the Survey of Income and Program Participation (SIPP), also matched with several files administered by the Social Security Administration (SSA), to compute mortality ratios by income groups. I use these mortality ratios by earnings to adjust the average age-specific death rates from the 2001 U.S. Life Tables (Arias, 2004). However, the literature on accurately measuring the extent of differential mortality in the U.S. has recently expanded, as a result of which there are a number of more recent papers that empirically estimate this relationship. In this section, I examine if the initial computational results are sensitive the estimates of differential mortality used to calibrate the model.

Among others, Bosworth et al. (2016) is a recent study on old-age inequality and the longevity gap between the rich and the poor. In this paper, the authors use data on earnings, death rates, and a wide array of socio-economic variables such as educational attainment, from the SIPP and also the Health and Retirement Study (HRS), matched with SSA data on career earnings, benefit payments, and individual mortality. Using this data, they estimate a logit regression of mortality risk, for ages 50 and above, of the form

$$\frac{h_{it}}{1-h_{it}} = \exp\left(\beta_{ij} \times X_{ijt}\right), \tag{36}$$

where $h_{it} = Pr(Y_{it} = 1/Y_{it} = 0)$ is the hazard that person i will die in year t, and X_{ijt} is a vector of potential determinants of mortality risk. In incorporating their empirical estimates into the baseline model, I focus on age, birth cohort, household earnings (mid-career earnings relative to neighboring cohorts), and educational attainment (years of schooling, relative to neighboring cohorts). Finally, I re-scale the intercept in their mortality risk estimates so that death rates are smooth across age 50.[18] I report the values for the unobservable parameters for which the model

[18]Prior to age 50, I use the death rates from the 2001 U.S. Life Tables (Arias, 2004).

	π_c	σ	β	η	κ_1	κ_2	κ_3	τ_y	τ_k	$\underline{c}$	δ
Initial baseline	18.0	4	0.997	0.41	0.3186	1.788	5.48	0.126	0.1	\$28,600	0.053
Bosworth et al. (2016)	119.0	4	0.917	0.45	0.3186	1.788	5.48	0.126	0.17	\$18,150	0.042

Table 13: Unobservable parameter values under the baseline calibration and with accidental bequests.

	Target	Initial baseline	Bosworth et al. (2016)
Capital-output ratio	3.5	3.34	3.7
Investment-output ratio	0.21	0.21	0.19
Avg. hours of market work per week per worker (25-55)	40	38	37
Share of govt. expenditures in GDP	0.2	0.21	0.21
Ratio of SSI expenditures to GDP	0.003	0.0036	0.0037
Value of Statistical Life	\$4 m.–\$10 m.	\$9 m.	\$7.5 m.
Interest rate	–	0.0517	0.0535
Benefit adjustment factor	–	0.667	0.369

Table 14: Key macroeconomic variables under the baseline calibration and with mortality estimates from Bosworth et al. (2016).

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	0.203	0.12/0.12/0.12	1300.403
0.45/0.32/0.15	0.433	0.19/0.14/0.06	1301.395
Baseline (0.9/0.32/0.15)	0.369	0.33/0.12/0.06	1301.272
1.8/0.32/0.15	0.286	0.51/0.09/0.04	1301.145
Uniform benefits	–	–	1301.020

Table 15: Replacement rates and aggregate welfare under alternative benefit-earnings rules with differential mortality estimates from Bosworth et al. (2016).

closely matches the targets in Table 13, and I report the corresponding model performance in Table 14. The table shows that the equilibrium characteristics of the baseline model are largely insensitive to the empirical estimates of differential mortality use in the calibration. As before, the values of the parameters for which the model reasonably matches the targets are well within the range encountered in the literature.

Next, I compute equilibria of the model calibrated with mortality estimates from Bosworth et al. (2016) under alternative benefit-earnings rules, both with and without differential mortality. I first report in Table 15 the overall welfare results with differential mortality. As before, I report the benefit adjustment factor in the second column, the effective replacement rates in the third column, and ex-ante expected utility in the last column. The table shows that the overall welfare consequences of the alternative benefit-earnings rules are largely insensitive to the estimates of differential mortality used to calibrate the model. Even with the mortality estimates from Bosworth et al. (2016), overall welfare is maximized under the replacement rates 0.45/0.32/0.15, which is less progressive than the current U.S. rule. Therefore, I find that even with more recent estimates of differential mortality, the insurance effects of Social Security's benefit-earnings rule are weak, warranting a less progressive relationship between past earnings and Social Security benefits.

Finally, to examine how differential mortality matters for the welfare ranking of these alternative benefit-earnings rules, I replace the mortality estimates from Bosworth et al. (2016) with those from the 2001 U.S. Life Tables for all households. As before, I re-scale the death rates and the old-age consumption floor so that life expectancy and the SSI expenditures-to-GDP ratio are constant

Replacement rates	Adjustment factor	Effective replacement rates	W
Proportional/linear	0.205	0.12/0.12/0.12	1312.662
0.45/0.32/0.15	0.433	0.20/0.14/0.07	1313.176
Baseline (0.9/0.32/0.15)	0.374	0.34/0.12/0.06	1313.223
1.8/0.32/0.15	0.292	0.53/0.09/0.04	1313.228
Uniform benefits	–	–	1312.692

Table 16: Replacement rates and aggregate welfare under alternative benefit-earnings rules without differential mortality estimates from Bosworth et al. (2016).

across the models with and without differential mortality. I report the overall welfare consequences of modifying Social Security's benefit-earnings rule with this model in Table 16. The table shows a pattern that is now familiar: modifying the benefit-earnings rule has only a small effect on welfare in equilibrium, and when differential mortality is eliminated, welfare is maximized under a benefit-earnings rule *more* progressive than the current U.S. rule. Taken together, these two experiments suggest that the welfare ranking of Social Security's benefit-earnings rules, both with and without differential mortality, is not sensitive to the estimates of differential mortality used to calibrate the baseline model.

11 The welfare criterion

So far, I have evaluated all the welfare consequences of Social Security's benefit-earnings rule, both with and without differential mortality using the ex-ante expected lifetime utility criterion defined in (34). While this "utilitarian" measure is the most commonly used welfare criterion in the quantitative public finance literature, there are alternative ways to evaluate social welfare in the context of the current question. First, the primary justification behind the creation of Social Security in 1935 was to make "more adequate provisions for aged persons", which suggests that Social Security's primary aim was to alleviate old-age poverty.[19] Second, due to the presence of differential mortality in the current model, the welfare of the elderly poor receives an even smaller weight in the expected lifetime utility criterion. Given these two facts, one could argue that the welfare criterion used to evaluate Social Security in the context of the current paper should attach more weight to the utility of low-income households. Therefore, in this section, I evaluate if the initial computational results are sensitive to the welfare criterion used to evaluate Social Security. In particular, I use a Rawlsian "maximin" criterion (Rawls, 1974; Weymark, 2016), measured by the effect of the alternative benefit-earnings rules on the lowest-utility households, i.e. those with the worst realization of the permanent productivity fixed effect $U(p = 0.54)$.

I report the welfare effects of Social Security's benefit-earnings rules under the "maximin" criterion in Tables 17 and 18. The welfare effects with differential mortality are reported in Table 17, and those without differential mortality are reported in Table 18. It is clear from Table 17 that the welfare ranking of alternative benefit-earnings rules under the "maximin" welfare criterion is identical to that under the initial ex-ante expected lifetime utility criterion with differential mortality. Therefore, even when more weight is attached to the utility of low-income households, a benefit-earnings rule less progressive than the current U.S. rule maximizes welfare. Consistent with the mechanisms from the simple two-period model, I find that adopting a more progressive benefit-earnings rule has an overall negative effect on households with $p = 0.54$ in the presence of differential mortality.

[19] See the legislative history of the Social Security Act of 1935 at http://www.ssa.gov/history.

Replacement rates	Adjustment factor	Effective replacement rates	$U(p = 0.54)$
Proportional/linear	0.339	0.2/0.2/0.2	367.331
0.45/0.32/0.15	0.767	0.35/0.25/0.12	369.346
Baseline (0.9/0.32/0.15)	0.667	0.6/0.21/0.1	368.878
1.8/0.32/0.15	0.525	0.94/0.17/0.08	368.715
Uniform benefits	–	–	368.643

Table 17: Replacement rates and aggregate welfare under the "maximin" welfare criterion with differential mortality risk.

Replacement rates	Adjustment factor	Effective replacement rates	$U(p = 0.54)$
Proportional/linear	0.352	0.21/0.21/0.21	402.336
0.45/0.32/0.15	0.755	0.34/0.24/0.11	402.224
Baseline (0.9/0.32/0.15)	0.654	0.59/0.21/0.1	402.213
1.8/0.32/0.15	0.507	0.91/0.16/0.08	402.248
Uniform benefits	–	–	401.857

Table 18: Replacement rates and aggregate welfare under the "maximin" welfare criterion without differential mortality risk.

Without differential mortality, however, the welfare results in Table 18 are not as clear cut. Utility of households with $p = 0.54$ is maximized when benefits are a fully proportional (linear) function of past earnings, but it is non-monotonic with respect to the progressivity of the benefit-earnings rule. These results closely mirror the consumption equivalences reported in Table 7. This is because in this counterfactual model, households with $p = 0.54$ have a slightly better earnings history, on the average, compared to those with $p = 1.0$ and 1.84. Consequently, benefit-earnings rules less progressive than the current U.S. rule lead to welfare improvements for households with $p = 0.54$.

Overall, I find that when a Rawlsian "maximin" criterion is used to evaluate Social Security, the effect of differential mortality on the welfare ranking of alternative benefit-earnings rules is slightly different, compared to when the traditional ex-ante expected lifetime utility criterion is used. With the "maximin" criterion, welfare is maximized with a fully proportional (linear) benefit-earnings rule, whereas with the ex-ante expected lifetime utility criterion, welfare is maximized with the replacement rates 0.45/0.32/0.15. Both of these benefit-earnings rules are less progressive than the current U.S. rule.

12 Conclusions

While linking public pension benefits to work-life income is common within the industrialized world, Social Security is unusual in the sense that there is an explicit progressive link between average earnings over the work life, and the benefits paid out to an individual. The rationale behind this link is that it provides partial insurance against uninsurable shocks to labor income, such as unemployment or the inability to secure a high-paying job. However, recent empirical evidence shows that there is a significant positive correlation between earnings and life expectancy, which has the potential of undoing the progressivity built into Social Security's benefit-earnings rule. In this paper, I examine if differential mortality has any implications for the welfare effects of Social Security's benefit-earnings rule.

To do this, I first use a simple two-period overlapping generations model with incomplete markets to illustrate that for households with lower earnings, (a) Social Security's internal rate

of return from a progressive benefit-earnings rule is lower in the presence of differential mortality, compared to when mortality risk is symmetric and uncorrelated to earnings, and (b) the expected utility from old-age consumption is lower due to their higher mortality risk. Next, I evaluate these effects using a rich multi-period calibrated general-equilibrium model with rational life-cycle permanent-income households, incomplete markets, and uninsurable labor income and differential mortality risks. I use this model to compute the welfare ranking of alternative benefit-earnings rules, ranging from a fully proportional (or linear) function of past earnings (i.e. with zero implicit redistribution), to a fixed benefit that is uniform across all retirees and completely unrelated to past earnings (i.e. with full redistribution), and then examine if the welfare rankings are sensitive to the positive correlation between earnings and survivorship.

My computational results suggest that the welfare ranking of these benefit-earnings rules is, in fact, sensitive to differential mortality. I find that welfare is maximized with a benefit-earnings rule *less* progressive than the current U.S. rule with differential mortality, but with a *more* progressive rule when differential mortality is eliminated. Moreover, I find that this result is robust with respect to whether accidental bequests originating from the assets of the deceased are taxed away or redistributed to the surviving population, as well as with respect to the estimates of differential mortality used to calibrated the quantitative model. Finally, differential mortality does not appear to affect the welfare ranking of the benefit-earnings rules only when a Rawlsian "maximin" criterion is used to evaluate the welfare effects, rather than the traditional ex-ante expected utility criterion commonly encountered in the literature.

Both Grochulski and Kocherlakota (2010) and Michau (2014) show that an earnings history-dependent tax-and-transfer scheme can be used to implement the socially optimal allocation in a heterogeneous-agent economy with private information. The findings from this paper suggest that differential mortality may be an important determinant of the welfare implications of such a history-dependent tax system. The presence of differential mortality reduces Social Security's internal rate of return for households with unfavorable earnings histories, and it also limits their welfare gains from better work-retirement consumption smoothing. My findings suggest that these two effects are quantitatively important: benefits that are more strongly related to past work-life income maximize aggregate welfare in the presence of differential mortality.

Appendix A

The type$-l$ household's utility maximization problem is

$$\underset{c_{l,1},c_{l,2}}{\text{Max}} \quad U = u(c_{l,1}) + \beta\ Q_l\ u(c_{l,2}) \tag{37}$$

subject to

$$c_{l,1} + \frac{c_{l,2}}{1+r} = we_l, \tag{38}$$

where $we_l = y_{l,1} - t_{l,1} + \frac{y_{l,2}-t_{l,2}}{1+r}$. Substituting the budget constraint in the utility function, the objective function is

$$U = u(c_{l,1}) + \beta\ Q_l\ u\left((we_l - c_{l,1})(1+r)\right). \tag{39}$$

The first-order condition for this problem is

$$\begin{aligned} u'(c_{l,1}) - \beta\ Q_l(1+r)\ u'\left((we_l - c_{l,1})(1+r)\right) = 0 \\ \implies F\left(c_{l,1}; \beta, Q_l, r, we_l\right) = 0. \end{aligned} \tag{40}$$

Because

$$F_{c_{l,1}} = u''(c_{l,1}) + \beta\ Q_l(1+r)^2\ u''\left((we_l - c_{l,1})(1+r)\right) < 0, \tag{41}$$

the Implicit Function Theorem can be used to express the optimal values of the first- and second-period consumptions as

$$c^*_{l,1} = c_{l,1}\left(\beta, Q_l, r, we_l\right)$$
$$c^*_{l,2} = \left(we_l - c_{l,1}\left(\beta, Q_l, r, we_l\right)\right)(1+r),$$

and the optimal value function as

$$V(\beta, Q_l, r, we_l) = u\left(c_{l,1}\left(\beta, Q_l, r, we_l\right)\right) + \beta Q_l u\left(\left(we_l - c_{l,1}\left(\beta, Q_l, r, we_l\right)\right)(1+r)\right). \tag{42}$$

Applying the Envelope Theorem on this value function, the welfare gain from the improved consumption smoothing is given by

$$V_{we_l}(\beta, Q_l, r, we_l) = \beta Q_l u'\left(\left(we_l - c_{l,1}\left(\beta, Q_l, r, we_l\right)\right)(1+r)\right) > 0, \tag{43}$$

which is nothing but the expected marginal utility of old-age consumption of the type$-l$ household, weighted by its probability of surviving into retirement. Therefore, the higher mortality risk experienced by the type$-l$ households heavily discounts the marginal utility of the improved work-retirement consumption smoothing.

References

Arias, E., 2004. United States Life Tables, 2001. National Vital Statistics Reports 52 (14).

Bagchi, S., 2015. Labor Supply and the Optimality of Social Security. Journal of Economic Dynamics and Control 58, 167–185.

Bagchi, S., Feigenbaum, J., 2014. Is smoking a fiscal good? Review of Economic Dynamics 17 (1), 170–190.

Boskin, M. J., Kotlikoff, L. J., Puffert, D. J., Shoven, J. B., 1987. Social Security: A Financial Appraisal Across and Within Generations. Working Paper 1891, National Bureau of Economic Research.

Bosworth, B., Burtless, G., Zhang, K., 2016. Later Retirement, Inequality in Old Age, and the Growing Gap in Longevity Between Rich and Poor. Economic Studies at Brookings.

Bullard, J., Feigenbaum, J., 2007. A Leisurely Reading of the Life-Cycle Consumption Data. Journal of Monetary Economics 54 (8), 2305–2320.

Caliendo, F. N., Guo, N. L., Hosseini, R., 2014. Social security is NOT a substitute for annuity markets. Review of Economic Dynamics 17 (4), 739–755.

Conesa, J. C., Garriga, C., 2009. Optimal response to a transitory demographic shock in social security financing. Federal Reserve Bank of St. Louis Review 91 (1), 33–48.

Conesa, J. C., Krueger, D., 1999. Social Security Reform with Heterogeneous Agents. Review of Economic Dynamics 2, 757–795.

Coronado, J. L., Fullerton, D., Glass, T., 2002. Long-Run Effects of Social Security Reform Proposals on Lifetime Progressivity. University of Chicago Press, pp. 149–206.

Coronado, J. L., Fullerton, D., Glass, T., 2011. The Progressivity of Social Security. The B.E. Journal of Economic Analysis & Policy 11 (1), 1–45.

Cremer, H., Lozachmeur, J.-M., Pestieau, P., 2004. Social security, retirement age and optimal income taxation. Journal of Public Economics 88 (11), 2259–2281.

Cristia, J. P., 2009. Rising mortality and life expectancy differentials by lifetime earnings in the united states. Journal of Health Economics 28 (5), 984–995.

De Nardi, M., İmrohoroğlu, S., Sargent, T. J., 1999. Projected U.S. Demographics and Social Security. Review of Economic Dynamics 2 (3), 575–615.

Farhi, E., Werning, I., 2013. Insurance and Taxation over the Life Cycle. The Review of Economic Studies 80 (2), 596–635.

Fehr, H., Kallweit, M., Kindermann, F., 2013. Should pensions be progressive? European Economic Review 63, 94–116.

Fuster, L., İmrohoroğlu, A., İmrohoroğlu, S., 2003. A Welfare Analysis of Social Security in a Dynastic Framework. International Economic Review 44 (4), 1247–1274.

Golosov, M., Troshkin, M., Tsyvinski, A., 2011. Optimal Taxation: Merging Micro and Macro Approaches. Journal of Money, Credit and Banking 43, 147–174.

Grochulski, B., Kocherlakota, N., 2010. Nonseparable preferences and optimal social security systems. Journal of Economic Theory 145 (6), 2055–2077.

Guner, N., Kaygusuz, R., Ventura, G., 2012. Taxation and Household Labour Supply. The Review of Economic Studies 79 (3), 1113–1149.

Hall, R. E., Jones, C. I., 2007. The value of life and the rise in health spending*. The Quarterly Journal of Economics 122 (1), 39–72.

Heathcote, J., Perri, F., Violante, G. L., 2010. Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006. Review of Economic Dynamics 13 (1), 15–51.

Hubbard, R. G., Judd, K. L., 1987. Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax. The American Economic Review 77 (4), 630–646.

Huggett, M., Ventura, G., 1999. On the Distributional Effects of Social Security Reform. Review of Economic Dynamics 2 (3), 498–531.

İmrohoroğlu, A., İmrohoroğlu, S., Joines, D. H., 1995. A Life Cycle Analysis of Social Security. Economic Theory 6 (1), 83–114.

İmrohoroğlu, A., İmrohoroğlu, S., Joines, D. H., 2003. Time-Inconsistent Preferences and Social Security. The Quarterly Journal of Economics 118 (2), 745–784.

İmrohoroğlu, S., Kitao, S., 2012. Social Security Reforms: Benefit Claiming, Labor Force Participation, and Long-Run Sustainability. American Economic Journal: Macroeconomics 4 (3), 96–127.

Jeske, K., 2003. Pension systems and aggregate shocks. Federal Reserve Bank of Atlanta Economic Review (Q1), 15–31.

Karabarbounis, M., 2012. Heterogeneity in labor supply elasticity and optimal taxation. In: 2012 Meeting Papers. No. 655. Society for Economic Dynamics.

Kitagawa, E., Hauser, P., 1987. Differential Mortality in the United States: A Study in Socioeconomic Epidemiology. Harvard University Press, Cambridge, MA.

Kitao, S., 2014. Sustainable social security: Four options. Review of Economic Dynamics 17 (4), 756–779.

Meyerson, N., Sabelhaus, J., 2006. Is Social Security Progressive? Economic and Budget Issue Brief, Congressional Budget Office.

Michau, J.-B., 2014. Optimal redistribution: A life-cycle perspective. Journal of Public Economics 111, 1–16.

Nishiyama, S., Smetters, K., 2008. The Optimal Design of Social Security Benefits. Working Papers wp197, University of Michigan, Michigan Retirement Research Center.

Rawls, J., 1974. Some reasons for the maximin criterion. The American Economic Review 64 (2), 141–146.

Saez, E., 2002. Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses. The Quarterly Journal of Economics 117 (3), 1039–1073.

Samuelson, P. A., 1975. Optimum social security in a life-cycle growth model. International Economic Review 16 (3), 539–544.

Sheshinski, E., 2008. Optimum Delayed Retirement Credit. The MIT Press.

Smetters, K., 2003. Is the social security trust fund worth anything? Working Paper 9845, National Bureau of Economic Research.

Storesletten, K., Violante, G., Heathcote, J., 2012. Redistributive Taxation in a Partial Insurance Economy. 2012 Meeting Papers 588, Society for Economic Dynamics.

Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59 (2), 371–396.

Viscusi, W. K., Aldy, J. E., 2003. The value of a statistical life: A critical review of market estimates throughout the world. Journal of Risk and Uncertainty 27 (1), 5–76.

Weymark, J. A., 2016. Social Welfare Functions. Oxford University Press, pp. 126–159.

Zhao, K., 2014. Social security and the rise in health spending. Journal of Monetary Economics 64 (0), 21–37.