Towson University
Department of Economics
Working Paper Series

Working Paper No. 2016-03

Differential Mortality and the Progressivity of Social Security

by Shantanu Bagchi

August, 2016

Differential Mortality and the Progressivity of Social Security[*]

Shantanu Bagchi[†]

August 15, 2016

Abstract

I examine if the positive correlation between wealth and survivorship has any implications for the progressivity of Social Security's benefit-earnings rule. Using a general-equilibrium macroeconomic model calibrated to the U.S. economy, I show that the optimal Social Security arrangement is largely insensitive to wealth-dependent mortality risk. This is because while a more progressive benefit-earnings rule provides increased insurance for households with relatively unfavorable earnings histories, labor supply, savings, and therefore survivorship, their relatively high mortality risk heavily discounts the utility from old-age consumption. I find that these two effects roughly offset each other, yielding nearly identical benefit-earnings rules both with and without differential mortality.

JEL Classifications: E21, E62, H55

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

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

1 Introduction

While the primary justification behind the creation of Social Security in 1935 was to make "more adequate provision for aged persons"[1], today it accounts for 16-17% of total annual federal government expenditures, second only to health expenditures (excluding defense). Economists have traditionally viewed Social Security 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 due to unfavorable labor-market outcomes. Social Security annuities are paid until death, so they insure individuals against the risk of out-living their own savings. As a retirement pension, Social Security benefits are linked to work-life earnings, and the strength of this link determines how unfavorable labor-market events, such as the inability to secure a high-paying job or unemployment, affect retirement consumption.

While linking public pension benefits to work-life income is common within the industrialized world, Social Security is slightly 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. 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 in the U.S., but only about 40-50% for individuals whose wages are higher than the average wage. Therefore, an individual at the bottom of the income distribution receives a higher return on every dollar of Social Security contributions paid, relative to an individual at the top of the income distribution. While this arrangement is designed to provide insurance against unfavorable labor income shocks, its effectiveness depends on a multitude of economic and demographic factors.

From the perspective of a household, an important determinant of the welfare gains 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. Moreover, empirical evidence suggests that there is a significant positive correlation between wealth and life expectancy in the U.S. (Kitagawa and Hauser, 1987). While this phenomenon, referred to as differential mortality, is worth studying in its own right, it has important implications for Social Security. This is because the positive correlation between wealth and survivorship has the potential of undoing the progressivity built into Social Security's current benefit-earnings formula (Coronado et al., 2002, 2011).

In this paper, I quantitatively examine if differential mortality has any implications for the progressivity of Social Security, in an economy with rational life-cycle permanent-income agents. To do this, I construct an overlapping-generations macroeconomic model with incomplete markets, an unfunded public pension system that closely mimics Social Security, and households that experience uninsurable labor income and wealth-dependent mortality risks. Social Security provides partial insurance against these risks, because households do not have access to private insurance markets. 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 some key features of the U.S. economy, such as overall capital accumulation, pattern of labor supply over the life cycle, the earnings distribution relative to Social Security's taxable maximum, 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 fully proportional (i.e. with zero implicit insurance) to completely flat (i.e. with full insurance). Finally, I examine if these welfare implications are sensitive to the positive correlation between wealth and survivorship. I also identify the macroeconomic effects of these experiments, such as their implications for the labor market, capital accumulation,

---

[1]The legislative history of the Social Security Act of 1935, available on http://www.ssa.gov/history.

national income, and the government's budget.

My computations suggest that the welfare implications of Social Security's benefit-earnings rule are largely insensitive to differential mortality. In the presence of wealth-dependent mortality risk, the progressivity of the benefit-earnings rule has two competing effects on welfare. On the one hand, a more progressive benefit-earnings rule provides better work-retirement consumption smoothing for households with relatively unfavorable earnings histories, labor supply, savings, and therefore survivorship. However, on the other hand, their relatively higher mortality risk causes these households to heavily discount the utility from old-age consumption. I find that these effects roughly offset each other: the optimal Social Security arrangement is nearly identical both with and without wealth-dependent mortality risk. In both cases, the optimal benefit-earnings rule warrants benefits that are nearly flat and unrelated to past work-life income. While this arrangement has positive insurance effects for households with unfavorable earnings histories, it also imposes higher implicit tax rates on households with relatively favorable earnings histories, distorting their labor supply. I find that the welfare gains from the insurance effects outweigh the welfare losses from the labor supply distortions, both with and without wealth-dependent mortality risk.

My computations also predict that with the flat-benefit arrangement, expected Social Security benefits increase by as much as 42-46% for households with unfavorable earnings histories, and decline by as much as 21-23% for households with relatively favorable earnings histories, but the overall size of Social Security remains roughly unchanged. Finally, 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 baseline and the flat-benefit arrangement, capital stock, labor, national income, and government expenditures, all increase by 1-2%, both with and without wealth-dependent mortality risk.

This paper contributes to three separate strands of literature. First, it contributes to a large literature that attempts to characterize the optimal redistribution scheme in a heterogeneous-agent economy, accounting for distortions to consumption, saving, and labor supply.[2] Two papers in this literature that highlight the importance of earnings history-dependent tax systems are Grochulski 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 optimum 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. Two studies from this literature that are closest to the current paper are Huggett and Ventura (1999) 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.[3] In general, they do not find substantial welfare improvements in switching from the current U.S. program to the two-tier structure. Nishiyama and Smetters (2008), on the other hand, find that while the progressive linking of earnings with retirement benefits in the U.S. has beneficial insurance effects, it also in-

---

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

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

troduces 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 flat. However, these studies also ignore differential mortality, and their computational experiments do not allow for a clean interpretation of the welfare effects of Social Security's benefit-earnings link.

Finally, this paper also contributes to an empirical literature that attempts to measure 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 formula. 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.

The rest of the paper is organized as follows: Section 2 introduces the model, and Sections 3 and 4 describe the baseline calibration and its results. I describe the computational experiments in Section 5, and examine their results in Sections 6 and 7. Finally, I conclude in Section 8.

2 The model

The unit of the current 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 depends on the household's relative position in the wealth distribution. The household does not have access to markets where it can purchase insurance against the labor income and mortality risks.

At each date, surviving households earn labor income if they work, and they also receive Social Security benefits after the full retirement age. Firms operate competitively and produce output using capital, labor and a constant returns to scale technology. The government provides public goods and Social Security; the public goods purchases are funded using general tax revenues, and Social Security is funded through a payroll tax on labor income. Social Security plays two roles in this model economy: it provides intergenerational transfers from the young to the old, and it also provides partial insurance against labor income and differential mortality risks.

2.1 Preferences

Households derive utility both from consumption and leisure. A household's labor supply decision at each instant 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. The period utility

4

function is given by

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

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 utility terms), 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

$$U = E\left[\sum_{s=0}^{T} \beta^s \left\{\Pi_{j=0}^{s-1} Q\left(j, wp(j)\right)\right\} u\left(c(s), 1-h(s), P(s)\right)\right], \tag{2}$$

where β is the discount factor, and $Q\left(j, wp(j)\right)$ is the probability at age j of surviving to the next period, which depends on age, as well as the household's relative position on the cross-sectional wealth distribution at age j, measured by its wealth percentile $wp(j)$.

2.2 Income

Conditional on labor force participation, a household earns before-tax wage income $y(s, \varphi) = h(s)w(s)e(s, \varphi)$ at age s, where $w(s)$ is the wage rate, and $e(s, \varphi)$ is a labor productivity endowment that depends on age and a stochastic productivity shock φ. This wage income is subject to two separate taxes: a progressive labor income tax $T_y(\cdot)$, and a payroll tax $T_{ss}(\cdot)$ for Social Security that is proportional up to the maximum taxable earnings of $\bar{y}$. After-tax wage income at age s is therefore given by

$$y^{at}(s, \varphi) = y(s, \varphi) - T_y(y(s, \varphi)) - T_{ss}(y(s, \varphi); \bar{y}) \tag{3}$$

Finally, a household's asset holdings at age s earn a risk-free interest rate r, which is subject to a proportional capital income tax at rate τ_k. The after-tax interest rate faced by the household is therefore given by $(1 - \tau_k)r$.

It is useful to note here that because they are unable to insure themselves against mortality risk, 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%.[4]

2.3 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

---

[4]How these accidental bequests are handled within the model has important consequences for its quantitative predictions. A common assumption in the literature is that these accidental bequests are evenly distributed back to the surviving population. However, it has been recently shown that with this assumption, Social Security fails to provide any insurance against mortality risk. Caliendo et al. (2014) demonstrate that if one accounts for how Social Security affects the accidental bequest that households leave (and also receive) in equilibrium, then higher mandatory saving through Social Security crowds out these accidental bequests, and therefore has zero effect on life-cycle wealth. Moreover, with this assumption, the accidental bequests create an additional layer of redistribution in the model that does not exist in reality. Because a higher life expectancy increases saving, it also increases accidental bequests and therefore has a pure income effect on *all* households in equilibrium.

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, also called the Primary Insurance Amount (PIA), is then calculated as a piecewise linear function of the AIME. Finally, the government scales benefits up or down proportionally so that Social Security's budget is balanced.[5]

2.4 A household's optimization problem

The state vector of each household 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 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 represented as

$$V(s, x) = \max_{c, k', h, P} \left\{ u(c, 1 - h, P) + \beta \, Q(s, x) \, E \left[V(s + 1, x') \right] \right\} \tag{4}$$

subject to

$$c + k' = (1 + (1 - \tau_k)r) \, k + y^{at}(s, \varphi) + \Theta(s - T_r) \, b(AIME) \tag{5}$$

$$y^{at}(s, \varphi) = h(s)w(s)e(s, \varphi) - T_y \left(h(s)w(s)e(s, \varphi) \right) - T_{ss} \left(h(s)w(s)e(s, \varphi); \bar{y} \right) \tag{6}$$

$$k' \geq 0 \tag{7}$$

$$AIME' = AIME \ \forall \ s \geq T_r \tag{8}$$

$$0 \leq h \leq 1, \tag{9}$$

$$\tag{10}$$

where

$$\Theta(s - T_r) = \left\{ \begin{array}{ll} 0 & s < T_r \\ 1 & s \geq T_r \end{array} \right.$$

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.

2.5 Technology and factor prices

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

$$Y(t) = K(t)^\alpha L(t)^{1-\alpha}, \tag{11}$$

---

[5]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, 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). This is due to 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.

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

$$r = MP_K - \delta = \alpha \left[\frac{K(t)}{L(t)} \right]^{\alpha-1} - \delta \qquad (12)$$

$$w(t) = MP_L = (1 - \alpha) \left[\frac{K(t)}{L(t)} \right]^{\alpha} \qquad (13)$$

where δ is the depreciation rate of physical capital and $w(t)$ is the wage rate at time t.

2.6 Aggregation

The population structure in the model 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 at any instant t are given by

$$K(t) = \sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, k(s+1; x)\mu_s(x) \qquad (14)$$

$$L(t) = \sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, h(s; x)e(s, x)\mu_s(x), \qquad (15)$$

where $N(t-s)$ is the size of the age-s cohort.

The total value of the accidental bequests by households who die on date t is given by

$$\text{Beq}(t) = (1 + (1 - \tau_k)r) \left[\sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} (k(s+1; x) - k(s; x)) \, \mu_s(x) \right]$$

$$- n \sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, k(s+1; x)\mu_s(x), \qquad (16)$$

and the budget-balancing condition for Social Security is given by

$$\sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, T_{ss}\left(h(s; x)w(s)e(s, x); \bar{y}\right) \mu_s(x)$$

$$= \sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, \Theta(s - T_r)b(x)\mu_s(x). \qquad (17)$$

Finally, the government also adjusts the labor income tax function $T_y(\cdot)$ and the capital income tax rate τ_k such that the total tax revenues from labor income, capital income, and the accidental bequests are sufficient to finance its expenditures

$$\text{Beq}(t) + \tau_k rK(t) + \sum_{s=0}^{T} N(t-s) \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \, T_y\left(y(s; x)\right) \mu_s(x) = G(t), \qquad (18)$$

where $G(t)$ is the exogenously given level of government expenditures.

2.7 Competitive equilibrium

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

1. cross-sectional consumption allocations $\{c(s;x)\}_{s=0}^{T}$, participation decisions $\{P(s;x)\}_{s=0}^{T}$, and labor hours allocations $\{h(s;x)\}_{s=0}^{T}$,

2. an aggregate capital stock $K(t)$ and labor $L(t)$,

3. a rate of return r and a wage rate $w(t)$,

4. Social Security benefits $b(x)$, 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 expenditure at time t equals consumption plus net investment plus government purchases, which is equal to the total income earned from capital and labor at time t.

$$
\begin{aligned}
C(t) + K(t+1) - (1-\delta)K(t) + G(t) &= C(t) + (n+\delta)K(t) + G(t) \\
&= w(t)L(t) + (r+\delta)K(t) \\
&= Y(t) \tag{19}
\end{aligned}
$$

I consider only steady-state equilibria of this model, so I set calendar time to $t = 0$ and also normalize the initial newborn cohort size to $N(0) = 1$.

3 Calibration

3.1 Demographics

To set the demographic parameters of the model, I first assume that households enter the model at the actual age of 25 (model age of zero), and 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 survival probabilities, I use estimates of differential mortality from Attanasio and Hoynes (2000), who use the 1984 and 1987 panels of the Survey of Income and Program Participation (SIPP) to estimate an empirical model that relates age and the position in the wealth distribution to mortality outcomes. Attanasio and Hoynes (2000) estimate four-month survival rates beyond age 50 as a function of age and wealth percentile, using two definitions of wealth: total net worth, which includes financial equity, home equity, business equity, and IRA/Keogh accounts minus unsecured debt, and also financial wealth,

8
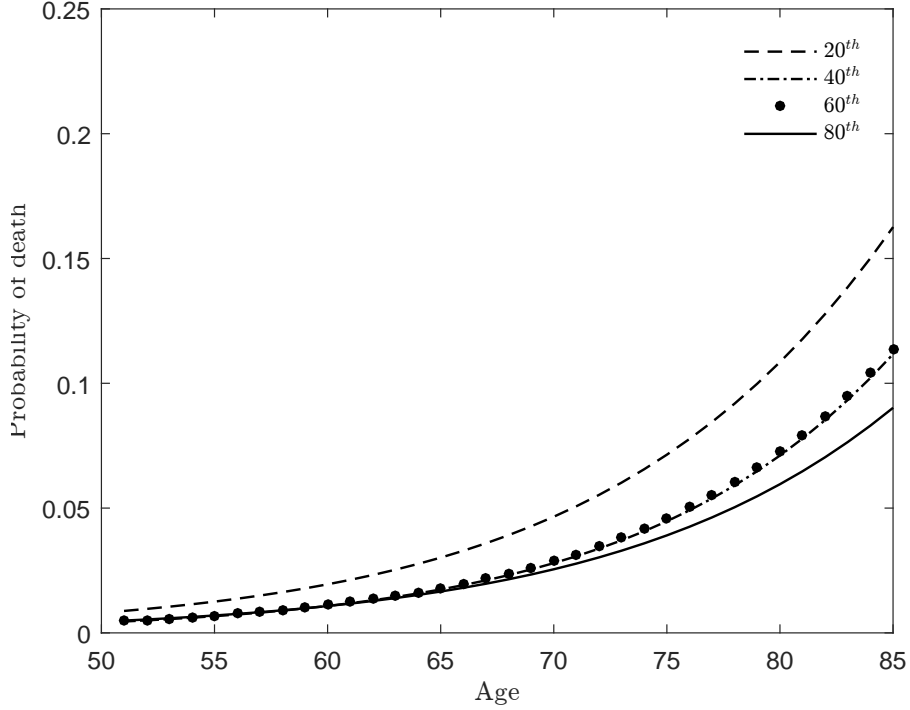
Figure 1: Death rates by age for select wealth percentiles from Attanasio and Hoynes (2000).

which is closer to the wealth definition in the CEX. Their estimated four-month survival function is given by

$$
\begin{aligned}
PS^4\left(j, wp(j)\right) &= \frac{1}{1 + \exp\left(f\left(j, wp(j)\right)\right)}, \\
f\left(j, wp(j)\right) &= -7.517 + (-19.773) \times wp(j) + (31.197) \times wp(j)^2 + (-14.575) \times wp(j)^3 \\
&\quad + (0.095) \times j \times wp(j) + (-0.098) \times j \times wp(j)^2 + (0.074) \times j, \quad (20)
\end{aligned}
$$

where $j > 50$ is actual age, and $wp(j)$ is the household's wealth percentile at age j. I take this estimated survival function directly to the model.[6,7] I plot the estimated annual death rates as a function of age for select wealth percentiles, and also as a function of wealth percentiles for select ages, in Figures 1 and 2 respectively. It is clear from the figures that mortality risk is increasing in age and decreasing in wealth percentile.

3.2   Social Security

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

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

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 annually

---

[6]The annual survival rates are obtained by multiplying three consecutive four-month survival rates.

[7]Prior to age 50, I calibrate the model using the average age-specific death rates from the 2001 U.S. Life Tables in Arias (2004).
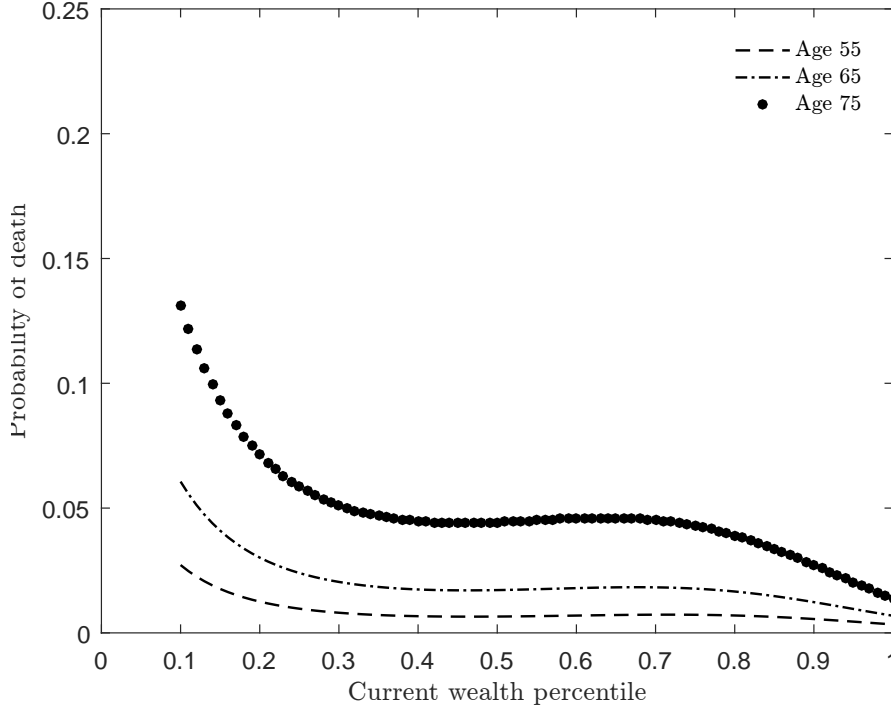
Figure 2: Death rates by wealth percentile for select ages from Attanasio and Hoynes (2000).

relative to the average wage in the U.S. For example, the taxable maximum was set at \$76,200 in the year 2000, but was adjusted to \$106,800 in 2010 and \$113,700 in 2013. During the same period, the national average wage index increased from \$32,155 to \$41,674, and finally to \$44,888.[8] 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 Social Security benefit amount (also known as the PIA), I incorporate the U.S. benefit-earnings rule into the model. The benefit-earnings rule in the U.S. is a concave (piecewise linear) function of past work-life income, the AIME. The Social Security Administration (SSA) calculates the AIME, and then it calculates the PIA as a fraction of the AIME.

Depending on how large or small the AIME for an individual is relative to the average wage in the economy, the SSA adjusts the fraction of the AIME that PIA replaces. For example, in the year 2000, the OASI benefit was 90% of the AIME for the first \$531, 32% of the next \$2,671, and 15% of the remaining up to the maximum taxable earnings. As shown by Huggett and Ventura (1999), these dollar amounts come out to be roughly 20%, 124%, and 247% of the average wage in the economy. These percentage amounts are referred to as the "bend points" of the benefit rule, and I take them directly to the model, while adjusting them proportionally so that Social Security's budget is balanced in equilibrium. It is worth noting that the progressivity in the benefit rule is captured by the fact that the "replacement rate" is decreasing in the AIME (see Figure 3).

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.

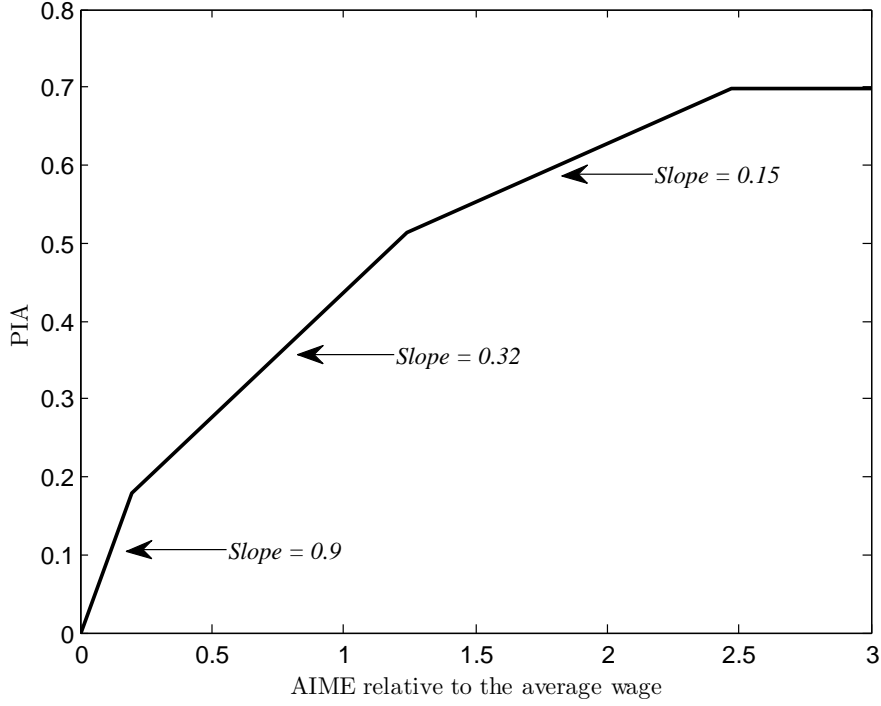[8]See http://www.ssa.gov/oact/cola/awiseries.html for more details.

Figure 3: Benefit formula in the U.S.

3.3 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

$$\log e(s, \varphi) = \epsilon(s) + \varphi, \tag{21}$$

where $\epsilon(s)$ is a deterministic age-dependent component, and φ is the stochastic component, given by

$$\varphi_t = p + z_t + \nu_t \tag{22}$$

$$z_t = \rho z_{t-1} + \upsilon_t, \tag{23}$$

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

I parameterize $\epsilon(s)$ 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. The resulting $\epsilon(s)$, normalized with respect to productivity at age 25, is plotted in Figure 4.

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 and transitory shocks to $\sigma_p^2 = 0.124$ and $\sigma_\nu^2 = 0.04$ respectively, and variance of the white-noise disturbance to $\sigma_\upsilon^2 = 0.015$. I use Gaussian quadrature to approximate the distribution of the permanent shock using a three-point discrete distribution, and I approximate the persistent shock using a five-state first-order discrete Markov process following Tauchen and Hussey (1991).
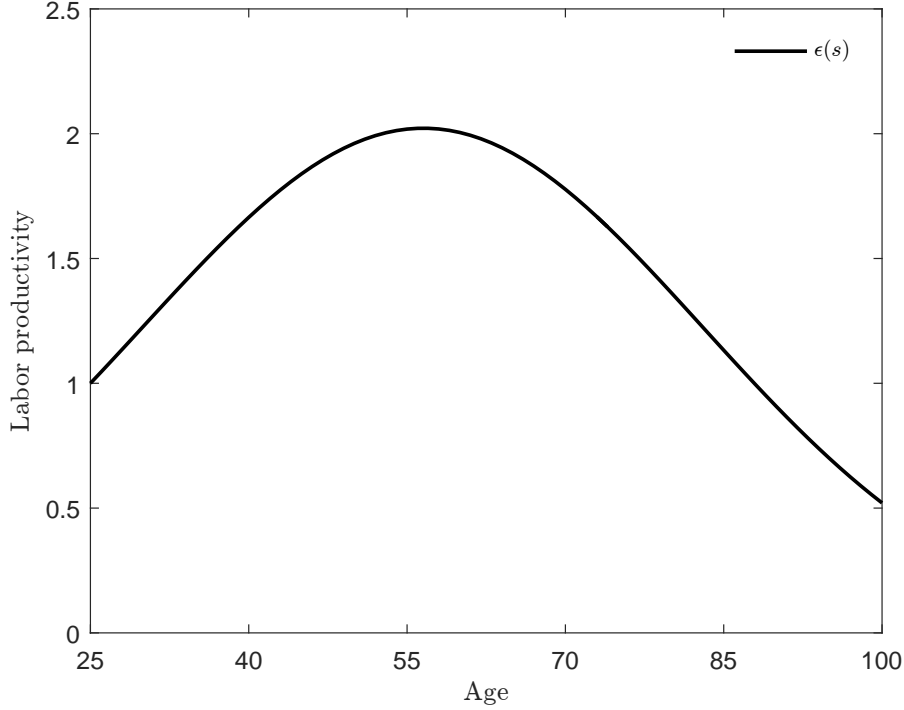
Figure 4: The age-dependent component of labor productivity from Kitao (2014).

3.4 Income tax

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

$$T_y(y) = y - (1 - \tau_y)y^{1-\tau_1}, \tag{24}$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (24) reduces to a proportional tax function with a marginal rate of τ_y. With this income tax function, after-tax labor income is log-linear in before-tax labor income. To estimate the parameters of this tax function, I take the 2012 tax rate schedule for a single filer in the U.S., compute the after-tax income for each level of before-tax income, and then regress the log of after-tax income on the log of before-tax income. This yields the following estimate for the parameter τ_1, which controls the progressivity of the tax code:

$$\hat{\tau}_1 = 0.06411. \tag{25}$$

I plot the average tax rates that emerge from the estimated tax function along with those from the U.S. tax schedule in Figure 5. Note that because these are the average rates, they are slightly lower than the marginal tax rates in the U.S. tax schedule. The top marginal tax rate in the U.S. tax schedule is 39.6%, but the top average rate is only around 28% because only a small fraction of income is subject to the top marginal rate.

3.5 Technology

The historically observed value of capital's share in total income in U.S. ranges between 30-40%, so I set $\alpha = 0.35$. Also, following Stokey and Rebelo (1995), I set the depreciation rate to $\delta = 0.06$.
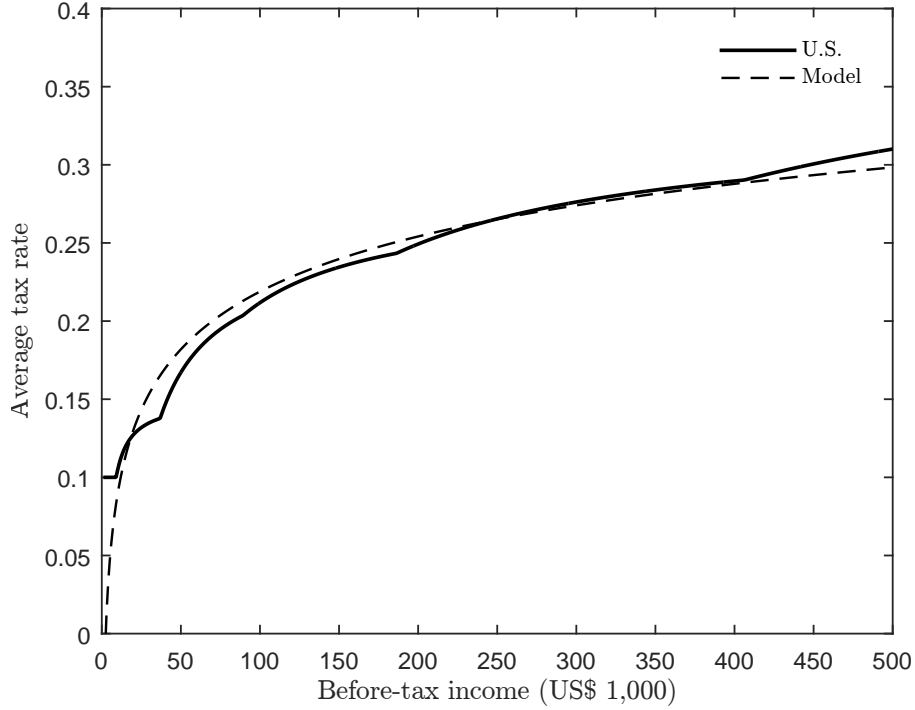
Figure 5: The average tax rates in the model, compared to those from the U.S. tax schedule.

3.6 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 σ, β, and η, the age-dependent labor force participation cost $\theta_P(s)$, and also the income tax parameters τ_y and τ_k, 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 get an equilibrium capital-output ratio of 3.0. 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.

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.[9]

Finally, I set $\tau_y = \tau_k$ and calibrate them such that the model matches 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.

---

[9]I set the maximum disposable time to 16 hours per day.

σ	β	η	κ_1	κ_2	κ_3	$\tau_y(=\tau_k)$
4	0.9706	0.425	6.12×10^{-8}	3.43×10^{-7}	2.98	0.163

Table 1: Unobservable parameter values under the baseline calibration.
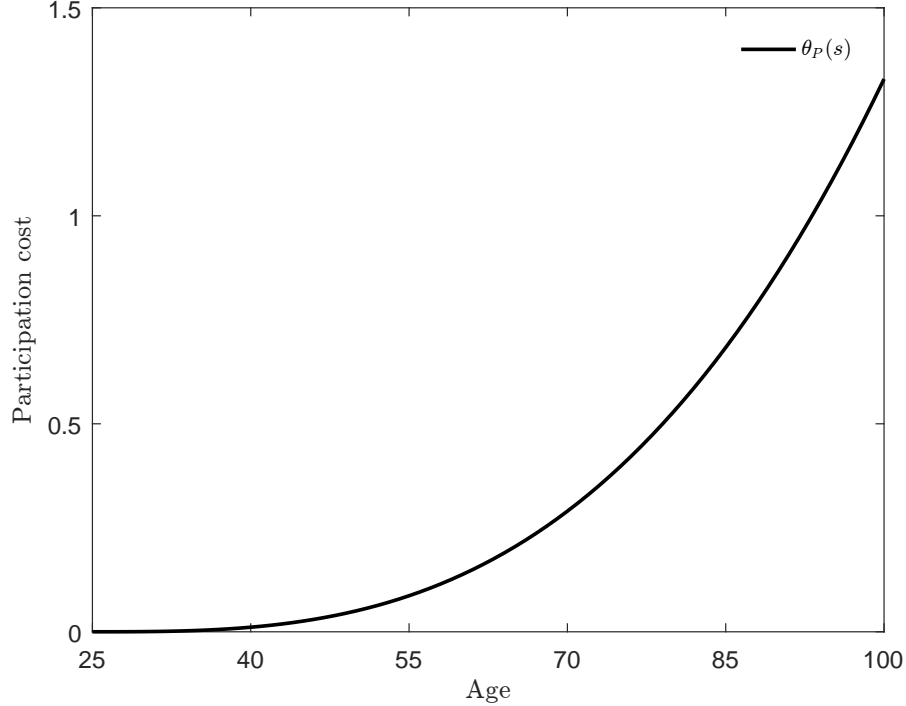
Figure 6: Age-dependent labor force participation cost $\theta_P(s)$.

4   Baseline economy

The unobservable parameter values under which the baseline equilibrium reasonably matches the above targets are reported in Table 1. Note that with leisure in period utility, the relevant inverse elasticity for consumption is $\sigma^c = 1 + \eta(\sigma - 1) = 2.3$, 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 (see Figure 6).

The model-generated values for key macroeconomic variables under the baseline calibration are reported in Table 2 along with their targets, and the cross-sectional labor force participation and labor hours data (conditional on participation) are reported in Figures 7 and 8. Note that the benefit-earnings rule adjustment factor in the baseline calibration is larger than unity, which implies that the bend points of the benefit-earnings rule are adjusted upward to balance Social Security's budget in the baseline equilibrium.

It is clear from Figures 7 and 8 that the baseline calibration does a reasonable job of matching observed labor supply behavior in the U.S. It replicates the rapid decline in participation after age 50 quite well, and it also reasonably matches the general declining trend of weekly hours over the life cycle. However, the current model fails to replicate the mild-hump shape in the hours profile, and it also generally underestimates both participation and weekly hours at later ages. There are two potential ways to improve the model's fit along these dimensions. First, I treat

	Target	Model
Capital-output ratio	3.0	2.97
Avg. hours of market work per week per worker (25-55)	40	39.2
Share of govt. expenditures in GDP	0.20	0.203
Interest rate	-	0.058
Ratio of Social Security expenditures to GDP	-	0.056
Benefit-earnings rule adjustment factor	-	1.16

Table 2: Key macroeconomic variables under the baseline calibration.
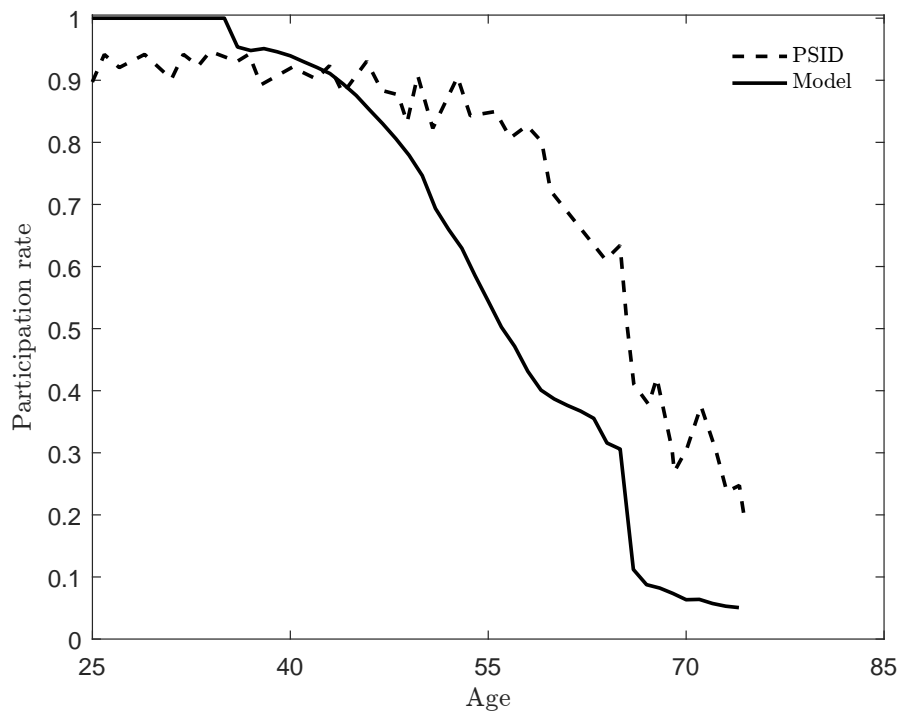
Figure 7: Cross-sectional labor force participation rates under the baseline calibration.
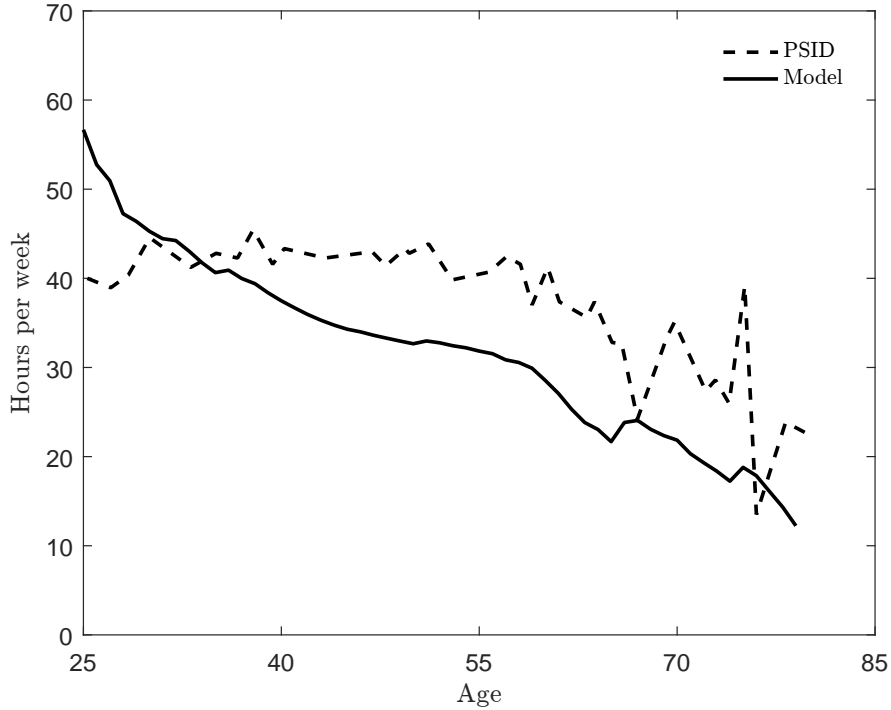
Figure 8: Cross-sectional mean of labor hours per week (conditional on participation) under the baseline calibration.

the age-dependent component of labor productivity $\epsilon(s)$ as observable, whereas in reality it is an unobservable structural parameter. Treating $\epsilon(s)$ as an unobservable parameter could potentially eliminate any selection bias arising from measuring it as residual wages (Bullard and Feigenbaum, 2007; Bagchi and Feigenbaum, 2014). Second, households in the current model smooth consumption across the work life and retirement (the life-cycle motive), and also across the stochastic realizations of the idiosyncratic productivity shock (the precautionary motive). However, both life-cycle and precautionary motives are less important at later ages, especially because the idiosyncratic productivity shock is highly persistent. Introducing a third factor that determines behavior, such as a bequest motive, could potentially induce older households to increase labor supply in the model. In fact, the absence of a bequest motive, and also any health risk, explains why the current model underestimates asset holdings at later ages, as seen in Figure 9 (De Nardi et al., 2010).

It is worthwhile at this point to examine the distribution of earnings in the baseline calibration, relative to the maximum taxable earnings for Social Security. First, in the baseline calibration, Social Security's tax base is about 82% of total earnings, which is very close to the current U.S. ratio of 83% reported by the Social Security Administration.[10] Second, from the perspective of a household, whether or not the cap on Social Security taxes binds depends on three key factors: the stochastic labor productivity shock, its implications for the household's life-cycle pattern of labor supply, and finally the interaction of labor supply with the life-cycle endowment profile. Unconditionally, the cap is more likely to bind for households with a favorable productivity shock, and conditional on a particular realization of the shock, the cap is more likely to bind when before-tax labor income is near or at its peak in the life cycle. In Figure 10, I report the fraction of workers

---

[10]Social Security's Annual Statistical Supplement, Table 4B.1.
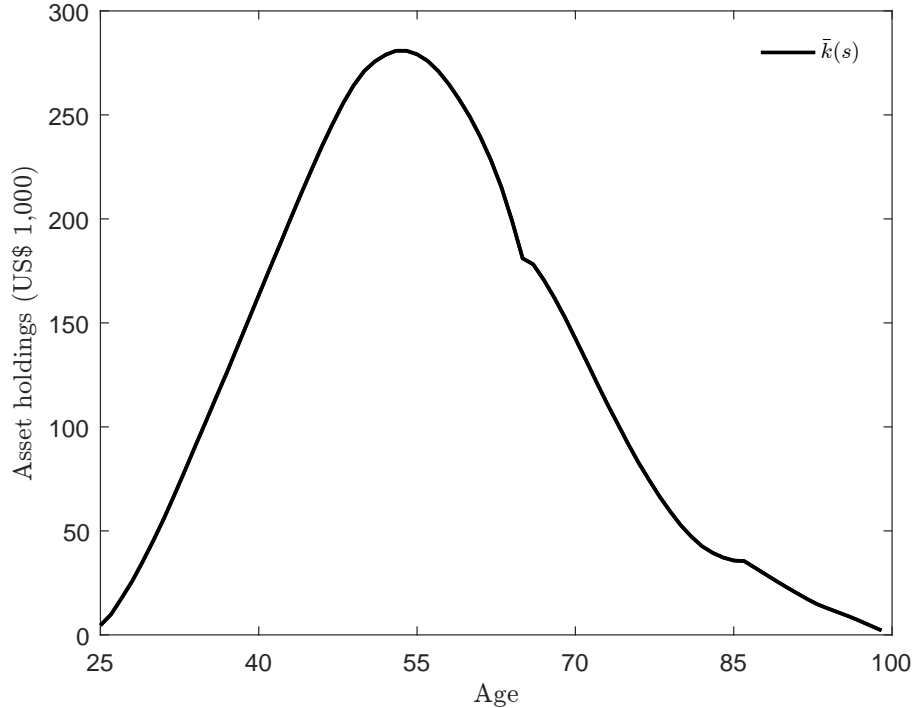
Figure 9: Cross-sectional mean of asset holdings under the baseline calibration.
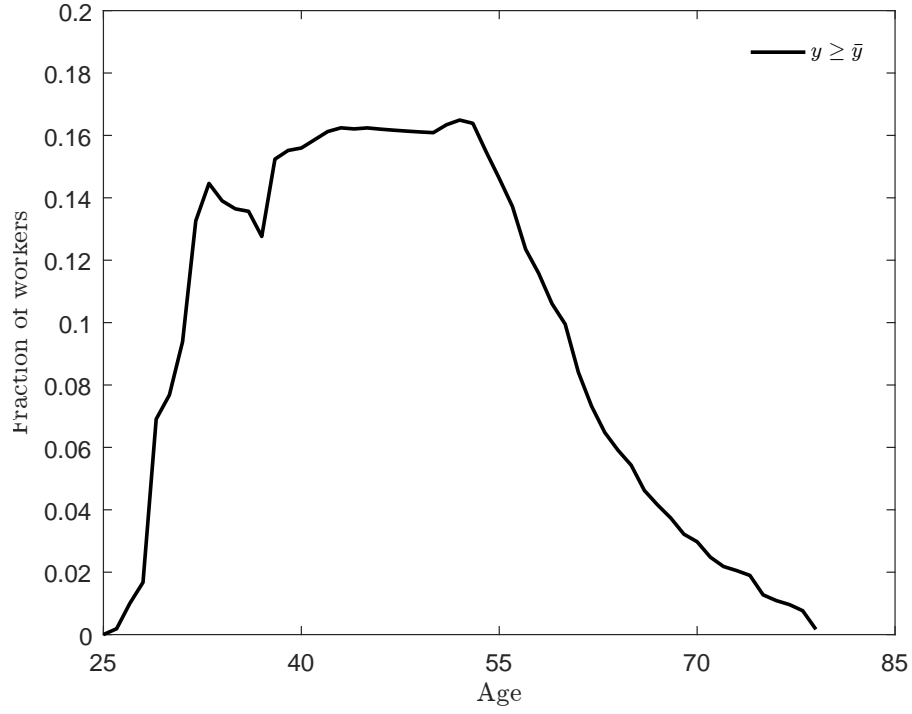
with labor income above the cap as a function of age in the baseline calibration, which shows that this ratio peaks out at 16% at age 47, roughly where labor income reaches its maximum in the life cycle.

Finally, in Figure 11, I report the fraction of households surviving to each age for each realization of the permanent productivity shock (p), based on their expected death rates, which in turn depend on their asset holdings. It is clear from the figure that households with relatively favorable productivity shocks (and therefore on higher wealth percentiles) are more likely to survive to later ages. For example, about 45% of the households with $p = 1.84$ survive up to age 80, compared to only 29% of households with $p = 0.54$. Also, the extent of differential mortality increases with age: the average survival rate of households with $p = 1.84$ is almost identical to that of households with $p = 0.54$ at age 60, but about 50% higher at age 80, and almost three times as large at age 90.

5   The experiments

The goal of this paper is to examine if differential mortality has any implications for the progressivity of Social Security's benefit-earnings rule. Essentially, this experiment involves computing 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 wealth-dependent 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 experiment? 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 optimal structure of Social Security benefits

Figure 10: Fraction of workers with labor income above the taxable maximum ($\bar{y}$) in the baseline calibration.
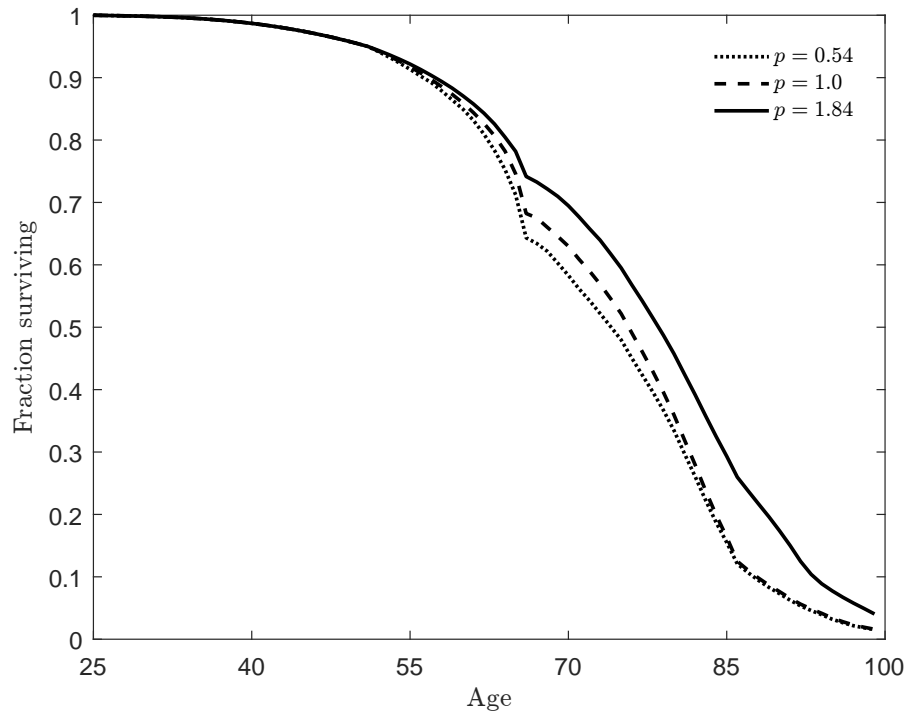
Figure 11: Fraction of households surviving to each age in the baseline calibration.

in a model such as this, I focus only on benefit-earnings rules that are structurally similar to the current U.S. rule. As discussed earlier, the current benefit-earnings rule in Social Security provides 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 3). With this structure, the progressivity of the benefit-earnings rule is largely controlled by the first bend point. Reducing this replacement rate strengthens the link between Social Security benefits and past work-life income, thereby reducing the progressivity of the benefit-earnings rule, and increasing this replacement rate weakens the link between Social Security benefits and past work-life income, thereby increasing the progressivity of the benefit-earnings rule. Because of this reason, I focus only on this first bend point in my computational experiment, and I consider values for this bend point that yield benefit-earning rules ranging from perfectly proportional (i.e. with zero implicit insurance), to perfectly flat (i.e. with full insurance).

It is important to note here even though I focus only on the first bend point in the computational experiment, keeping Social Security's budget balanced with the current payroll tax rate and the taxable maximum requires the second and third bend points to be adjusted as well. For example, increasing the first bend point from 90%, while keeping the second and third bend points 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 second and third bend points must be reduced. This is accomplished automatically in the model through the benefit-earnings rule adjustment factor, which adjusts every alternative benefit-earnings rule to ensure that Social Security's budget is balanced. In fact, as we will see, this adjustment factor even offsets some of the direct change in the first bend point, in addition to the second and third bend points. 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

$$W = \sum_{s=0}^{T} \beta^s \sum_{x} \left\{ \Pi_{j=0}^{s-1} Q\left(j, x(j)\right) \right\} \ u\left(c(s;x), 1 - h(s;x), P(s;x)\right) \mu_s(x) \qquad (26)$$

which is the ex-ante expected lifetime utility. Then, to understand the distributional consequences of these benefit-earnings rules, I define a consumption equivalence ψ for each realization of the permanent productivity shock (p) that solves

$$E\left[\sum_{s=0}^{T} \beta^s \ \left\{ \Pi_{j=0}^{s-1} Q\left(j, wp^C(j)\right) \right\} \ u\left((1 + \psi) c^C(s), 1 - h^C(s), P^C(s)\right) \right] =$$

$$E\left[\sum_{s=0}^{T} \beta^s \ \left\{ \Pi_{j=0}^{s-1} Q\left(j, wp^H(j)\right) \right\} \ u\left(c^H(s), 1 - h^H(s), P^H(s)\right) \right], \qquad (27)$$

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 productivity shock, under each one of my 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.

Bend points	Adjustment factor	Effective bend points	W
Baseline	1.162	1.05/0.37/0.17	-66.98
Proportional	0.697	0.63/0.63/0.63	-67.31
0.45/0.32/0.15	1.675	0.75/0.53/0.25	-67.15
1.8/0.32/0.15	0.720	1.30/0.23/0.11	-66.82
9.0/0.32/0.15	0.182	1.64/0.06/0.03	-66.70
45.0/0.32/0.15	0.038	1.72/0.01/0.006	-66.67
108.0/0.32/0.15	0.016	1.74/0.005/0.002	-66.66
Flat benefits	−	−	-66.59

Table 3: Overall welfare consequences of alternative Social Security benefit-earnings rules with wealth-dependent mortality risk.

6   Computational results

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, with a strong link between benefits and past work-life income. They argue that Social Security's relatively long averaging period of thirty-five years 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 wealth-dependent mortality in Table 3. In the first column, I report the bend points of the benefit-earnings rules 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" bend points in the third column, and in the last column I report overall welfare.

The following two facts are clear from Table 3. First, with wealth-dependent mortality risk, ex-ante expected utility is maximized when benefits are flat, i.e. when they are completely unrelated to work-life income. Reducing the first bend point from its baseline value of 0.9 reduces overall welfare, but increasing it consistently increases overall welfare. Second, modifying the first bend point of the benefit-earnings rule also requires adjusting the second and third bend points, so that Social Security's budget is balanced with the current payroll tax rate and taxable maximum. This adjustment occurs through the benefit-earnings rule adjustment factor, which steadily declines as the first bend point is increased. In addition to offsetting some of the direct change in the first bend point, this leads to a consistent flattening of the benefit-earnings rule, as seen in the declining values of the second and third bend points in the table. To summarize, I find that with wealth-dependent mortality risk, the optimal benefit-earnings rule warrants flat benefits; the insurance effects of a more progressive benefit-earnings rule appear to outweigh the labor supply distortions, which is in stark contrast with the findings of Nishiyama and Smetters (2008).

A potential reason behind the discrepancy in results between Nishiyama and Smetters (2008)

Permanent productivity shock (p)	0.54	1.00	1.84
Proportional	-0.68	-0.27	0.0
0.45/0.32/0.15	-0.38	-0.13	0.0
1.8/0.32/0.15	0.31	0.14	0.0
9.0/0.32/0.15	0.57	0.26	0.0
45.0/0.32/0.15	0.63	0.30	-0.11
108.0/0.32/0.15	0.65	0.30	-0.11
Flat benefits	0.73	0.38	-0.11

Table 4: Consumption equivalences (ψ) under the alternative benefit-earnings rules with wealth-dependent mortality risk.

and the current model is the size of labor supply distortions in their experiments. There are two ways in which Nishiyama and Smetters (2008) potentially overestimate the welfare losses from labor supply distortions. First, 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, such an experiment does not isolate the effect of changing the benefit-earnings rule; 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. Second, Nishiyama and Smetters (2008) also do not report how their baseline model matches the heterogeneity in earnings, relative to Social Security's taxable maximum. If they underestimate the fraction of labor income above the taxable maximum, then they also potentially overestimate the labor supply distortions from Social Security to households with favorable earnings histories.

To assess the distribution of welfare gains and losses under the alternative benefit-earnings rules in the presence of wealth-dependent mortality risk, I report in Table 4 the consumption equivalence (ψ) for each realization of the permanent productivity shock (p) under each computation (in percentage terms). As expected, the table shows that increasing the degree of progressivity generates welfare gains for households more likely to have unfavorable earnings histories, and welfare losses for households more likely to have favorable earnings histories. Adopting the flat-benefit arrangement leads to welfare improvements equivalent to increases of 0.7% and 0.4% in period consumption for households with $p = 0.54$ and 1.0 respectively, and a welfare loss equivalent to a 0.1% reduction in period consumption for households with $p = 1.84$.

To summarize, my computations suggest that in a general-equilibrium environment with uninsurable labor income and wealth-dependent mortality risk, the insurance effects of Social Security's benefit-earnings rule are large enough to warrant flat benefits that are completely unrelated to past work-life income. While this arrangement leads to higher implicit tax rates for households with relatively favorable earnings histories, their distortionary welfare losses are smaller in magnitude.

7   The importance of differential mortality

Let us now address the key question of this paper: how important is the effect of differential mortality on the welfare implications of Social Security's benefit-earnings rule? To measure 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 wealth-dependent mortality risk. Specifically, I

Bend points	Adjustment factor	Effective bend points	W
Baseline	0.922	0.83/0.29/0.14	-66.18
Proportional	0.544	0.49/0.49/0.49	-66.27
0.45/0.32/0.15	1.332	0.60/0.43/0.20	-66.25
1.8/0.32/0.15	0.571	1.03/0.18/0.09	-66.11
9.0/0.32/0.15	0.141	1.27/0.05/0.02	-66.03
45.0/0.32/0.15	0.029	1.33/0.009/0.004	-66.00
108.0/0.32/0.15	0.012	1.34/0.004/0.002	-66.00
Flat benefits	–	–	-66.00

Table 5: Overall welfare consequences of alternative Social Security benefit-earnings rules without wealth-dependent mortality risk.

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. Then, with this modified model, I compute the welfare implications of changing the progressivity in Social Security'e benefit-earnings rule in Table 5. As before, I report the bend points of the benefit-earnings rules being examined in the first column, the corresponding adjustment factor in the second column, and the "effective" bend points in the third column. I report the overall welfare in the last column.

Three facts are clear from the table. First, modifying Social Security's benefit earnings rule has only marginally different welfare implications in the absence of differential mortality. Overall welfare peaks out when the benefit-earnings rule is nearly flat, when the first bend point is increased five-fold from its baseline level, yielding bend points 45.0/0.32/0.15 (effective 1.33/0.009/0.004). Under this rule, expected Social Security benefits for households with $p = 1.84$ are only 2% higher than those with $p = 0.54$, compared to the baseline scenario (with the current U.S. rule), when they are almost twice as large. Second, there is no change in overall welfare when benefit-earnings rules more progressive than 45.0/0.32/0.15 are adopted: the flat-benefit rule yields the same level of welfare as 45.0/0.32/0.15. Finally, elimination of wealth-dependent mortality risk requires a larger decline in the benefit-earnings rule adjustment factor, relative to its corresponding baseline level. This suggests that without differential mortality, the longer survival of retirees with relatively unfavorable earnings histories has a larger effect on Social Security's fiscal status, compared to the cost savings from the reduced survival of retirees with better earnings histories, and therefore higher benefits.

The distributional consequences of the alternative benefit-earnings rules in the absence of differential mortality are reported in Table 6. The table shows a very similar pattern: a more progressive benefit-earnings rule generates welfare gains for households more likely to have relatively unfavorable earnings histories, and welfare losses for households likely to have favorable earnings histories. Adopting the rule with bend points 45.0/0.32/0.15 (and also rules more progressive than 45.0/0.32/0.15) leads to welfare improvements equivalent to increases of 0.6% and 0.1% in period consumption for households with $p = 0.54$ and 1.0 respectively, and a welfare loss equivalent to a 0.3-0.4% reduction in period consumption for households with $p = 1.84$. These effects are only marginally different from those with wealth-dependent mortality risk, which shows that the welfare effects of Social Security's benefit-earnings rule are largely insensitive to the positive correlation between wealth and survivorship.

The reason why differential mortality has little effect on the progressivity of Social Security is as follows. In the presence of wealth-dependent mortality risk, Social Security's benefit-earnings rule has two competing effects on welfare. While a more progressive benefit-earnings rule provides better work-retirement consumption smoothing for households with relatively unfavorable earnings

Permanent productivity shock (p)	0.54	1.00	1.84
Proportional	-0.47	0.04	0.28
0.45/0.32/0.15	-0.31	0.01	0.19
1.8/0.32/0.15	0.24	0.03	-0.14
9.0/0.32/0.15	0.50	0.08	-0.30
45.0/0.32/0.15	0.58	0.11	-0.33
108.0/0.32/0.15	0.59	0.11	-0.36
Flat benefits	0.59	0.12	-0.36

Table 6: Consumption equivalences (ψ) under the alternative benefit-earnings rules without wealth-dependent mortality risk.

Permanent productivity shock (p)	0.54	1.00	1.84
Proportional	-0.27	0.0	0.15
0.45/0.32/0.15	-0.19	-0.01	0.09
1.8/0.32/0.15	0.16	0.0	-0.09
9.0/0.32/0.15	0.40	0.03	-0.17
45.0/0.32/0.15	0.46	0.03	-0.20
108.0/0.32/0.15	0.47	0.03	-0.20
Flat benefits	0.46	0.02	-0.21

Table 7: Change in expected Social Security benefits under the alternative benefit-earnings rules with wealth-dependent mortality risk.

histories, lower saving, labor supply, and survivorship, their relatively higher mortality risk also causes them to heavily discount the utility from old-age consumption. These two effects almost offset each other in the current model, generating nearly identical optimal benefit-earnings rules, both with and without differential mortality. A comparison of Tables 4 and 6 shows that for every single benefit-earnings rule considered, the magnitude of welfare gains (losses) are always smaller (larger) in the absence of wealth-dependent mortality risks. This suggests that the consumption-smoothing effects of Social Security are smaller in this case, but not small enough to warrant an optimal benefit-earnings rule that is significantly less progressive than what we obtain in the presence of differential mortality.

So far, we have focused only on the welfare effects of modifying Social Security's benefit-earnings rule. I now turn to the potential macroeconomic implications of such a policy change. Perhaps the most important macroeconomic consequence of the modifying the benefit-earnings rule, both with and without differential mortality, is how it affects the level of Social Security benefits. I report in Table 7 the percentage change in expected Social Security benefits from the baseline with wealth-dependent mortality risk, for each value of the permanent productivity shock. The table shows that as expected, increasing the progressivity in the benefit-earnings rule (i.e. making it flatter and less related to past work-life income) increases the expected benefits for households with relatively unfavorable earnings histories, and reduces it for those with favorable earnings histories. Under the flat-benefit rule, expected benefits for households with $p = 0.54$ increase by almost 46% from the baseline, and decline by 21% for those with $p = 1.84$. These changes are roughly identical even without wealth-dependent mortality risk; between the baseline and bend points 45.0/0.32/0.15, expected benefits increase by 42% for for households with $p = 0.54$, and decline by 22% for households with $p = 1.84$.

I report the effects on other key macroeconomic variables in Table 8, such as aggregate capital, labor, national income, the interest rate, and the share of government expenditures in GDP,

	K	L	GDP	r	G/Y
Proportional	0.984	0.992	0.989	1.016	0.995
0.45/0.32/0.15	0.991	0.996	0.994	1.007	0.997
1.8/0.32/0.15	1.010	1.005	1.007	0.992	1.003
9.0/0.32/0.15	1.019	1.006	1.011	0.988	1.005
45.0/0.32/0.15	1.021	1.008	1.013	0.986	1.006
108.0/0.32/0.15	1.022	1.008	1.013	0.986	1.006
Flat benefits	1.023	1.009	1.014	0.981	1.007

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

relative to the baseline with wealth-dependent mortality risk. The table shows that increasing the progressivity of the benefit-earnings rule leads to a consistent increase in capital, labor, and GDP, a decline in the interest rate, and a roughly constant ratio of government expenditures to GDP. However, it is also clear that these changes are very small in magnitude: across the baseline and the flat-benefit rule, capital, labor, and national income do not increase by more than 1-2%. Overall, modifying the benefit-earnings rule does not appear to have significant general-equilibrium effects, and this result continues to hold even without wealth-dependent mortality risk.

To summarize, I find that the welfare implications of Social Security's benefit-earnings rule are largely insensitive to differential mortality. While a more progressive benefit-earnings rule offers better consumption-smoothing benefits for households with relatively unfavorable earnings histories, their higher mortality risk also causes them to heavily discount the utility from old-age consumption. I find that these two effects roughly offset each other: both with and without wealth-dependent mortality, the optimal Social Security arrangement warrants benefit-earning rules that are nearly flat and unrelated to past work-life income. Under the flat-benefit arrangement, expected Social Security benefits increase by 42-46% for households with relatively unfavorable earnings histories, and decline by 21-23% for those with favorable earnings histories. However, these policy experiments do not appear to have significant macroeconomic effects; capital, labor, and national income do not increase by more than 1-2%.

8 Conclusions

While linking public pension benefits to work-life income is common within the industrialized world, U.S. Social Security is slightly 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 wealth and life expectancy, which has the potential of undoing the progressivity built into Social Security's benefit-earnings rule. In this paper, I quantitatively examine if this differential mortality has any implications for the progressivity of Social Security's benefit-earnings rule. To do this, I construct a calibrated general-equilibrium model economy with rational life-cycle permanent-income households and uninsurable labor income and wealth-dependent mortality risks. I use this model to compute the welfare effects of alternative benefit-earnings rules ranging from fully proportional (i.e. zero implicit insurance) to completely flat (i.e. full insurance), and then examine if they are sensitive to the positive correlation between wealth and survivorship.

My computational results suggest that the welfare implications of Social Security's benefit-

earnings rule are largely insensitive to differential mortality. In the presence of wealth-dependent mortality risk, the progressivity of the benefit-earnings rule has two competing effects on welfare. On the one hand, a more progressive benefit-earnings rule provides better work-retirement consumption smoothing for households with relatively unfavorable earnings histories, labor supply, savings, and therefore survivorship. On the other hand, their relatively higher mortality risk causes these households to heavily discount the utility from old-age consumption. I find that these effects roughly offset each other: the optimal Social Security arrangement is nearly identical both with and without wealth-dependent mortality risk. In both cases, the optimal benefit-earnings rule warrants benefits that are nearly flat and unrelated to past work-life income. While this arrangement has positive insurance effects for households with unfavorable earnings histories, it also imposes higher implicit tax rates on households with relatively favorable earnings histories, distorting their labor supply. I find that the welfare gains from the insurance effects outweigh the welfare losses from the labor supply distortions, both with and without differential mortality.

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 when a "restricted" social optimum with the extant labor income and payroll tax functions for the U.S. is considered, such history-dependence may not necessarily be optimal. In fact, my findings suggest that with the current payroll tax rate and taxable maximum, Social Security benefits should be completely history-*in*dependent, both with and without the positive correlation between wealth and survivorship.

References

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

Attanasio, O. P., Hoynes, H., 2000. Differential Mortality and Wealth Accumulation. The Journal of Human Resources 35 (1), 1–29.

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.

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.

De Nardi, M., French, E., Jones, J. B., 2010. Why Do the Elderly Save? The Role of Medical Expenses. Journal of Political Economy 118 (1), 39–75.

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

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.

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.

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., 2003. Time-Inconsistent Preferences and Social Security. The Quarterly Journal of Economics 118 (2), 745–784.

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.

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

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.

Stokey, N. L., Rebelo, S., 1995. Growth Effects of Flat-Rate Taxes. Journal of Political Economy 103 (3), 519–550.

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.

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