Is the Social Security Crisis Really as Bad as We Think?

Bagchi, Shantanu

Towson University

August 2013

Is the Social Security Crisis Really as Bad as We Think?*

Shantanu Bagchi†

May 1, 2014

Abstract

Because they ignore the household-level and macroeconomic adjustments associated with longevity improvements, the actuarial projections of the Social Security Administration overestimate the Social Security crisis. Using a general-equilibrium model with heterogeneous agents and incomplete markets, I show that accounting for these adjustments, a significantly smaller decline in benefits is needed to balance the Social Security budget. Households respond to the longevity improvements by delaying retirement and Social Security benefit collection, working more hours, and by also saving more. In general equilibrium, these effects lead to a natural expansion of Social Security's tax base and generate significant delayed retirement credits, which the actuarial estimates completely overlook.

JEL Classifications: E21, H55, J22

Keywords: Social Security; longevity improvement; general equilibrium; delayed retirement; delayed retirement credit

*I would like to thank Frank Caliendo, Jim Feigenbaum, Scott Findley, Nick Guo, Kartik Athreya, Carlos Garriga, and Sagiri Kitao for helpful comments and suggestions. The participants of the 2009 QSPS Summer Workshop, the 2010 Midwest Macroeconomics Meetings, and the audience at the job market presentations at Towson and Georgia Southern University also provided useful feedback.

†Department of Finance and Economics, P.O. Box 8152, Georgia Southern University, Statesboro, GA 30460-8152. Phone: +1-912-478-0404. Email: sbagchi@georgiasouthern.edu

1 Introduction

Over the last few decades, mitigating the effect of longevity improvements on social insurance programs has been a major policy concern in the developed world. According to the 1998 World Health Report, 26 countries will have an average life expectancy at birth of more than 80 years by 2025.[1] As is well known, these demographic developments will significantly worsen the fiscal status of unfunded public pension systems in these countries. In the U.S., life expectancy at birth is projected to increase to slightly over 81 years for males, and to roughly 85 years for females by 2075 (Shrestha, 2006). According to the 2011 Social Security Trustees Report, the current payroll tax rate for the Old-Age and Survivors Insurance (OASI) program is sufficient to pay only 77% of scheduled benefits in 2036, and only 74% of scheduled benefits in 2085.

One difficulty with the projections of the Social Security Administration (SSA) is that they are purely actuarial in nature, and are therefore subject to the Lucas critique.[2] From an actuarial standpoint, it is easy to see that increased longevity will have a significant effect on future Social Security spendings: households that expect to live longer will also collect benefits for longer. With unchanged revenues, the only way to keep Social Security solvent in such a case would be to reduce the future benefit per retiree. The fact that only 77% and 74% of the scheduled benefits in 2036 and 2085 will be payable with the current contribution rate, implies that the benefit per retiree will have to be cut by 23-26% by the end of the century.

However, the story is not so simple from an economic standpoint, because the improvement in longevity will have additional effects on future Social Security revenues and spendings. First, a higher life expectancy will directly increase labor supply, because there will be more workers alive at any age, and also because it may give households an incentive to delay retirement and work more hours. Second, if households risk out-living their assets, then a higher life expectancy will also induce higher saving, and therefore stimulate the aggregate capital stock. In general equilibrium, these effects will lead to a natural expansion of Social Security's tax base. Moreover, an important determinant of the level of Social Security benefits paid to a household is the age at which it starts collecting the benefits. The effect of a higher life expectancy on future Social Security spendings will also depend on how the collection decision responds to the longevity improvement. The actuaries of the SSA use a variety of assumptions about how the household-level and macroeconomic variables relevant for Social Security will change over time, but they overlook these endogenous adjustments to the longevity improvement. In their actuarial estimates, the assumed changes in the relevant household-level and macroeconomic variables are completely exogenous.

In this paper, I compute the extent by which the SSA's actuarial projections

[1]The World Health Reports are published by the World Health Organization.

[2]See Ljungqvist (2008) for a discussion on the Lucas critique.

overestimate the negative impact of longevity improvements on the fiscal status of Social Security. To do this, I construct an overlapping-generations macroeconomic model with heterogeneous households, mortality risk, incomplete markets, and a realistic social security program. In the model, Social Security provides partial insurance against mortality risk, and also against an unfavorable labor productivity shock. Households in the model face a progressive labor income tax schedule similar to the U.S., and they are also allowed to choose when to start collecting Social Security benefits. 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 (both with respect to labor force participation and hours per week), the income distribution, and the share of government expenditures in GDP.

Using this calibrated model, I examine how a singular improvement in life expectancy affects the budget-balancing level of Social Security benefits. First, I compute a partial-equilibrium experiment that replicates the actuarial methodology used by the SSA: I hold all the household-level and macroeconomic variables fixed at the baseline level, and incorporate a longevity improvement consistent with the SSA' projection. I find that Social Security benefits decline by about 23% under this experiment, which is very close to the decline estimated by the SSA. Then, I compute a general-equilibrium experiment in which all the household-level and macroeconomic variables adjust to the longevity improvement. I find that the percentage decline in the benefits required to keep Social Security solvent is significantly smaller in this case. For households that do not delay benefit collection, benefits decline by at most 15%, which is only two-thirds of the decline needed to balance the budget when these adjustments are not accounted for (23%). For households that delay collection, I find that the delayed retirement credit (associated with delaying collection beyond the normal retirement age) increases Social Security benefits by as much as 32%. Therefore, using a model that satisfies the Lucas critique, I show that the Social Security crisis in the U.S. may not be as bad as we think.

In the current model, a higher life expectancy leads to a natural expansion of Social Security's tax base through its effect on aggregate labor supply, and also on the wage rate. I find that in general equilibrium, labor supply increases by 9.2% from the baseline, both because workers survive for longer, and also because they delay retirement by about two years on the average, and work about 90 minutes more per week averaged over the life cycle. Capital stock also increases (by about 21%) as households save more, both because they supply more labor and therefore earn higher income, and also because they have to smooth consumption over a longer expected lifespan. Together, these changes lead to a 3.7% increase in the equilibrium wage rate. Given that Social Security's tax base is simply the product of the wage rate and labor supply, these constitute a roughly 13% expansion in the future Social Security revenues. Once this expansion is accounted for, budget balancing requires

Social Security benefits to decline by a significantly smaller percentage. Moreover, households delay Social Security benefit collection by as much as two-and-a-half years on the average, which generates significant delayed retirement credits. I find that for some households, these credits are large enough to actually increase Social Security benefits from the baseline level.

The critical point of this paper is that the household-level responses to the longevity improvement have an important effect on the level of Social Security benefits. However, in a general equilibrium environment such as this, these household-level responses are themselves, in turn, determined by the level of Social Security benefits. Specifically, the longevity improvement affects household behavior through two separate channels: a primary channel, which captures the *ceteris paribus* effect of the longevity improvement itself, and also a secondary channel, which captures the negative effect of the longevity improvement on the Social Security benefits. To separate these two effects, I also compute a partial-equilibrium experiment where I hold all the macroeconomic variables (including Social Security benefits) fixed at the baseline, and incorporate the longevity improvement. I find that under this experiment, households respond by retiring earlier (by about a year-and-a-half on the average), but also by slightly increasing their weekly hours during their peak productivity years. Because this leads to a slight reduction in overall labor supply, life-cycle motives encourage households to save considerably more. Therefore, the findings from this experiment suggest that the household-level responses observed in general equilibrium are largely governed by the negative effect of the longevity improvement on Social Security benefits, rather than the longevity improvement itself.

Finally, I compute another general-equilibrium experiment in which I also account for the future improvements in health that are likely to accompany the improvements in longevity. To approximate this, I assume that labor productivity declines less rapidly with age under the longevity improvement. I find that accounting for the health improvements has almost no additional effects on the macroeconomic aggregates. Aggregate labor supply increases by 9.4%, and the wage rate increases by 3.4%, so the combined effect on Social Security's tax base is almost identical in this case. However, the health improvements cause households to further delay collection of Social Security benefits, relative to when they are not accounted for. The average delay in collection in this case is about three years, compare to two-and-a-half years when the health improvements are ignored. The delayed retirement credit leads to Social Security benefits declining by at most 11% under this experiment, which is only half of the 23% decline estimated using the SSA's actuarial methodology.

In terms of the nature of the exercise undertaken, this paper is very similar in spirit to Chen and İmrohoroğlu (2012), who quantitatively examine the implications of different expenditure projections, such as those provided by the Congressional Budget Office (CBO) and The National Commission on Fiscal Responsibility and

Reform, on the future debt-to-GDP ratio in the U.S. They demonstrate that the CBO's projections are likely to underestimate the future debt-to-GDP ratio, as they ignore the endogenous labor and capital responses to the changes in the marginal tax rates used in the projections. In this paper, I show that the SSA's projections are likely to overestimate the Social Security crisis in the U.S., because they ignore the very same endogenous responses to the improvements in longevity.

Economists have long emphasized the importance of studying Social Security reform using models that account for the endogenous general-equilibrium effects of aggregate shocks in an economy. For example, De Nardi et al. (1999) demonstrate that the SSA's projections about the future tax rates required to keep Social Security solvent in the U.S. may be overly optimistic, as they overlook the distortions imposed by those higher tax rates on household behavior. They show that higher taxes are likely to discourage labor supply and saving, which is likely to have a quantitatively important effect on the future income rate of the program. Jeske (2003) demonstrates that the privatization of Social Security can be beneficial for all future generations even in the presence of aggregate shocks, if the general-equilibrium effects of the privatization are accounted for. He shows that privatization of social security is likely to increase private saving and therefore the aggregate capital stock, which would lead to an improvement in welfare large enough to compensate against even large aggregate shocks. Therefore, the general finding is that the implications of Social Security reform can be markedly different depending on whether or not the associated general-equilibrium effects have been accounted for.[3] The conclusions of the current paper have a very similar flavor: the measured extent of the Social Security crisis can be markedly different if these general-equilibrium effects are carefully considered.[4]

From the perspective of an individual, labor supply over the life-cycle can be characterized by two margins: the extensive margin or the participation decision, which determines the fraction of lifetime spent in employment, and the intensive margin, which determines the hours of work supplied while participating. Both of these margins have been found to be empirically relevant in an aggregate sense, and also individually. For example, Wallenius (2009) finds that differences in social security account for 35-40% of the differences in aggregate hours between the U.S. and Belgium, France and Germany. Additionally, Prescott et al. (2009) note that about half of the differences in aggregate hours worked between the U.S. and continental Europe can be accounted for by the differences in the fraction of lifetime worked. However, as Rogerson and Wallenius (2009) demonstrate, it is necessary

[3]Other studies that have used general-equilibrium models to examine alternative proposals to reform Social Security in the U.S. include Huang et al. (1997), Huggett and Ventura (1999), Conesa and Garriga (2003, 2008), and Kitao (2012).

[4]General equilibrium effects have also been found to be quantitatively important in resolving other macroeconomic questions, such as the non-monotonicity of life-cycle consumption (Bullard and Feigenbaum, 2007; Feigenbaum, 2008a) and the welfare-improving role of unfunded Social Security in a fully rational economy (İmrohoroğlu et al., 1995).

to account for both the intensive and the extensive margins to correctly estimate the elasticity of labor supply to tax and transfer programs. They show that large aggregate elasticities are consistent with small micro or life-cycle elasticities if both the extensive and intensive margins are accounted for.[5] Given these facts, in the current paper I consider a labor supply decision in which both the extensive and the intensive margins are operative. Households face a fixed cost of labor force participation, and during every period of employment they choose the hours of work.[6] This participation cost generates a reservation wage in the current model, because of which the labor supply function is discontinuous in after-tax wages.

The rest of the paper is organized as follows. In Section 2, I describe in detail how the SSA measures the fiscal status of Social Security. I introduce the model in Section 3, and describe the baseline calibration and the results in Sections 4 and 5. In Section 6, I describe the demographic experiment, and also compute the actuarial estimates of the Social Security crisis from the model. In Sections 7, 8, and 9, I discuss the results of the general-equilibrium experiment, the partial-equilibrium experiment of a *ceteris paribus* longevity improvement, and the general-equilibrium experiment with health improvements respectively. Finally, I conclude in Section 10.

2 The SSA's methodology

The Social Security Administration annually reports the financial health of the Old-Age and Survivors Insurance (OASI) and the Disability Insurance (DI) Trust Funds in the SSA Trustees Report. Actuarial status of the OASDI program is calculated both in the short-(10 years) and long-range (75 years), using specific definitions for the program income and cost rates. Also, both the short- and long-range estimates are presented under three alternative sets of assumptions: low-cost, intermediate, and high-cost. The intermediate assumptions represent the Board of Trustees' best estimate of the future course of the population and the economy, whereas the low-cost and the high-cost assumptions represent more optimistic and more pessimistic estimates, respectively. According the 2011 report, non-interest income in the OASDI program is projected to be sufficient to pay only about 77% of scheduled benefits in 2036, and 74% of scheduled benefits in 2085 based on the intermediate assumptions.

The future financial status of the OASDI program depends on several key variables, such as mortality, average earnings, labor force participation rates, and in-

[5] Focusing only on the extensive margin of household labor supply, Ortiz (2009) finds that roughly 90% of the differences in the employment-to-population ratio at ages 60-64 across the OECD can be explained by the differences in the institutional features of social security.

[6] Cogan (1981) showed that a convenient way to explain the distribution of annual hours of work around part-time and full-time work is to include a fixed participation cost in an otherwise standard model.

flation. In the Trustees' Report, the SSA makes specific assumptions about how the values of each of these variables change over time. For example, average life expectancy at birth in the U.S. is assumed to reach 81.3, 85, and 88.9 years by 2085 under the low-cost, intermediate, and high-cost assumptions. Also, between 2020 and 2085, average nominal U.S. earnings are assumed to grow at the rate of 3.6%, 4%, and 4.4% per annum respectively. Among other factors, these increases in earnings reflect trend increases of 0.1%, 0.0%, and -0.1% per annum in the average hours worked in the U.S.

Even though the Trustees Report accounts for how the macroeconomic variables relevant for Social Security change over time, their assumed changes are completely exogenous. For example, the labor force participation rate projections of the SSA reflect the trend effect of increases in life expectancy, higher assumed disability prevalence rates, and an increasing proportion of males who never marry. Improved life expectancy will trivially increase the labor force participation rate, simply because there will be more workers alive at every age. However, if households respond to the improved life expectancy by working more, then the labor force participation rate will increase even further. Also, if the households risk out-living their assets, then a higher life expectancy will also induce higher saving. To measure the combined effect of these changes on the future financial status of Social Security, one needs an equilibrium model in which all of these adjustments are endogenous.

3 The model

Time is discrete and at each instant a new cohort is born and the oldest cohort dies. Cohort size grows at the rate of n per annum, maximum lifespan is T, and each household in a cohort faces an unconditional probability $Q(s)$ of surviving to age s.

Households smooth consumption and labor supply over the life cycle by accumulating a risk-free asset: physical capital. Private annuities markets are closed by assumption, because of which households are unable to fully insure themselves against mortality risk.[7] This constraint causes deceased households at every age to 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]

[7]Assuming closed private annuities markets is standard in this line of literature, and is also empirically consistent because in reality very few people annuitize. This phenomenon is referred to as the "non-annuitization" puzzle, because a standard life-cycle model predicts that households ought to invest exclusively in annuities if they are fairly priced. Explanations behind this puzzle include existence of pre-annuitized wealth in retirees' portfolios, actuarially unfair prices, bequest motives, and uncertain health expenses. See, for example, studies such as Pashchenko (2013), Dushi and Webb (2004), Mitchell et al. (1999), Lockwood (2012), and Turra and Mitchell (2004).

[8]How these accidental bequests are handled within the model can have 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

At each date, surviving households earn labor income if they work, and they also choose when to start collecting Social Security benefits. Firms operate competitively and produce output using capital, labor and a constant returns to scale technology. The government purchases public goods and provides Social Security. The public goods purchases are funded using the proceeds from the estate tax and also taxes on capital and labor income, and Social Security is funded through a payroll tax on labor income. Social Security plays two roles in this model economy: it provides partial insurance against mortality risk, and also against unfavorable shocks to labor income.

3.1 Preferences

Households derive utility both from consumption and leisure. A household's 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. The period utility 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, θ_P is the fixed 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. Expected lifetime utility from the perspective of a household born at date t is

$$U = \sum_{s=0}^{T} \beta^s Q(s)\, u\left(c(t+s,t), 1-h(t+s,t), P(t+s,t)\right) \tag{2}$$

where β is the discount factor. Also, since I normalize the period time endowment to unity, $0 \leq h(t+s,t) \leq 1$.

3.2 Income

Conditional on labor force participation, a household born at date t earns before-tax wage income $y(t+s,t) = h(t+s,t)w(t+s)e(\varphi,s)$ at age s, where $w(t+s)$ is the wage rate, and $e(\varphi,s)$ is a labor productivity endowment that depends on age

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 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.

and a productivity shock φ. This wage income is subject to two separate taxes: a progressive labor income tax, and a proportional payroll tax for Social Security. Following Storesletten et al. (2012) and Karabarbounis (2012), I assume that the labor income tax function is given by

$$T(y) = y - (1-\tau_y)y^{1-\tau_1}, \tag{3}$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (3) reduces to a proportional tax function with a marginal tax rate of τ_y. The payroll tax rate for Social Security is τ_{ss}. With these specifications, after-tax wage income at age s is given by

$$\begin{aligned} y^{at}(t+s,t) &= y(t+s,t) - T(y(t+s,t)) - \tau_{ss}y(t+s,t) \\ &= (1-\tau_y)y(t+s,t)^{1-\tau_1} - \tau_{ss}y(t+s,t). \end{aligned} \tag{4}$$

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 $(1-\tau_k)r$.

3.3 Social Security

In the model, Social Security provides retirement benefits to households, and the benefit paid to a particular household depends on several factors. First, the base benefit amount (also called the Primary Insurance Amount or the PIA) is progressively linked to a measure of past work-life income, called the Average Indexed Monthly Earnings (AIME). Second, based on the chosen collection date (T_c), the PIA receives an actuarial adjustment. Households that start collecting earlier than the normal retirement age (T_n) receive a permanent reduction in their PIA, called the early retirement penalty, and the households that delay collection beyond the normal retirement age receive a permanent increase in their PIA, called the delayed retirement credit. Households can start collecting as early as the early retirement age (T_e), but they must start collecting by the maximum retirement age (T_m). Third, if a household starts collecting Social Security benefits before the normal retirement age, then (s)he is subject to an earnings test, which adjusts the PIA downward if the household's earnings are higher than a threshold level ($\bar{e}$). Finally, the government also adjusts the level of benefits so that the budget for Social Security is balanced.

3.4 A household's optimization problem

A household born at date t faces the following optimization problem

$$\max_{c,h,P,T_c} U = \sum_{s=0}^{T} \beta^s Q(s) \left[\frac{\left\{c(t+s,t)^{\eta}(1-h(t+s,t))^{1-\eta}\right\}^{1-\sigma}}{1-\sigma} - \theta_P \cdot P(t+s,t) \right] \tag{5}$$

subject to

$$c(t+s,t) + k(t+s+1,t) = (1+(1-\tau_k)r)\,k(t+s,t) + y^{at}(t+s,t) \tag{6}$$

$$\begin{aligned} y^{at}(t+s,t) = (1-\tau_y)\,(h(t+s,t)w(t+s)e(\varphi,s))^{1-\tau_1} \\ -\tau_{ss}h(t+s,t)w(t+s)e(\varphi,s) + \Theta(s-T_c)b(t+s) \end{aligned} \tag{7}$$

$$0 \leq h(t+s,t) \leq 1 \tag{8}$$

$$k(t,t) = k(t+T+1,t) = 0 \tag{9}$$

where

$$\Theta(x) = \begin{cases} 0 & x \leq 0 \\ 1 & x > 0 \end{cases}$$

is a step function.

3.5 Technology and factor prices

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

$$Y(t) \ = \ K(t)^{\alpha} L(t)^{1-\alpha}, \tag{10}$$

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

$$r = MP_K - \delta = \alpha \left[\frac{K(t)}{L(t)}\right]^{\alpha-1} - \delta \tag{11}$$

$$w(t) = MP_L = (1-\alpha)\left[\frac{K(t)}{L(t)}\right]^{\alpha} \tag{12}$$

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

3.6 Aggregation

Let us define the set of productivity shocks at age s as $\Phi(s)$, where $\varphi \in \Phi(s)$. Then, the aggregate capital stock and labor supply are given by

$$K(t) = \sum_{s=0}^{T} N(t-s)Q(s) \int_{\Phi(s)} k(t, t-s-1; \varphi)\ \mathrm{d}\varphi \tag{13}$$

$$L(t) = \sum_{s=0}^{T} N(t-s)Q(s) \int_{\Phi(s)} h(t, t-s; \varphi)e(\varphi, s)\ \mathrm{d}\varphi. \tag{14}$$

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

$$B(t) = (1+r)\left[\sum_{s=0}^{T}\{N(t-s)Q(s) - N(t-s-1)Q(s+1)\}\int_{\Phi(s)} k(t,t-s-1;\varphi)\ \mathrm{d}\varphi\right]$$
$$-\sum_{s=0}^{T}\left(N(t-s+1) - N(t-s)\right)Q(s)\int_{\Phi(s)} k(t+1,t-s;\varphi)\ \mathrm{d}\varphi. \qquad (15)$$

Note that in a model with mortality risk and population growth, the number of households between two successive ages changes because of two reasons: only a fraction of each cohort survives to the following age, and over time cohorts get successively larger. The first term on the right-hand side of (15) gives the total assets left behind because of these two reasons. Therefore, to isolate the assets that are left behind purely because of mortality risk (i.e. by the households that die between ages s and $s+1$), the second term on the right-hand side of (15) reduces total assets by the part that is attributable only to population growth. In the absence of population growth, i.e. when $N(t-s) = N(t-s-1) = N$, (15) collapses to

$$B(t) = N \times (1+r)\left[\sum_{s=0}^{T}\{Q(s) - Q(s+1)\}\int_{\Phi(s)} k(t,t-s-1;\varphi)\ \mathrm{d}\varphi\right] \qquad (16)$$

which can be rewritten as

$$B(t) = N \times (1+r)\left[\sum_{s=0}^{T} h(s)Q(s)\int_{\Phi(s)} k(t,t-s-1;\varphi)\ \mathrm{d}\varphi\right] \qquad (17)$$

where $h(s) = -\left(Q(s+1) - Q(s)\right)/Q(s)$ is the hazard rate of dying between age s and $s+1$. It is easy to see that the right-hand side of (17) now only contains the assets left behind by the households that do not survive to the following ages.

The budget-balancing condition for Social Security is given by

$$\tau_{ss}w(t)L(t) = \sum_{s=0}^{T} N(t-s)Q(s)\int_{\Phi(s)} \Theta(s-T_c;\varphi)b(t-s;\varphi)\ \mathrm{d}\varphi. \qquad (18)$$

Finally, the government also adjusts the labor income tax parameter τ_y and the capital income tax rate τ_k such that total tax revenues from labor income, capital income, and the accidental bequests, are sufficient to finance its expenditures

$$B(t) + \tau_k rK(t) + \sum_{s=0}^{T} N(t-s)Q(s)\int_{\Phi(s)} T\left(y(t,t-s;\varphi)\right)\ \mathrm{d}\varphi = G(t), \qquad (19)$$

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

3.7 Competitive equilibrium

A competitive equilibrium in the current model is characterized by a collection of

1. cross-sectional consumption allocations $\{c(t,t-s;\varphi)\}_{s=0}^{T}$, labor force participation decisions $\{P(t,t-s;\varphi)\}_{s=0}^{T}$, and labor hours allocations $\{h(t,t-s;\varphi)\}_{s=0}^{T}$,
2. Social Security benefit collection decisions $T_c(\varphi)$,
3. an aggregate capital stock $K(t)$ and labor $L(t)$,
4. a rate of return r and a wage rate $w(t)$, and
5. Social Security benefits $b(t-s;\varphi)$

that

1. solves the households' optimization problems,
2. maximizes the firms' profits,
3. equilibrates the factor markets, and
4. balances the government's budgets.

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.

$$\begin{aligned} C(t)+K(t+1)-(1-\delta)K(t)+G(t) &= C(t)+(n+\delta)K(t)+G(t) \\ &= w(t)L(t)+(r+\delta)K(t) \\ &= Y(t) \end{aligned} \tag{20}$$

In computing this equilibrium, I set calendar time to $t=0$ and I also normalize the initial newborn cohort size to $N(0)=1$.

4 Calibration

4.1 Demographics

I first set the demographic parameters of the model. I assume that households enter the model at actual age 25, which corresponds to the model age of zero. To get the baseline survival probabilities, I use Feigenbaum's (2008a) sextic fit to the mortality data in Arias (2004), which is given by

$$\begin{aligned} \ln Q(s) \;=\; & -0.01943039+\left(-3.055\times10^{-4}\right)s+\left(5.998\times10^{-6}\right)s^2 \\ & +\left(-3.279\times10^{-6}\right)s^3+\left(-3.055\times10^{-8}\right)s^4+\left(3.188\times10^{-9}\right)s^5 \\ & +\left(-5.199\times10^{-11}\right)s^6, \end{aligned} \tag{21}$$

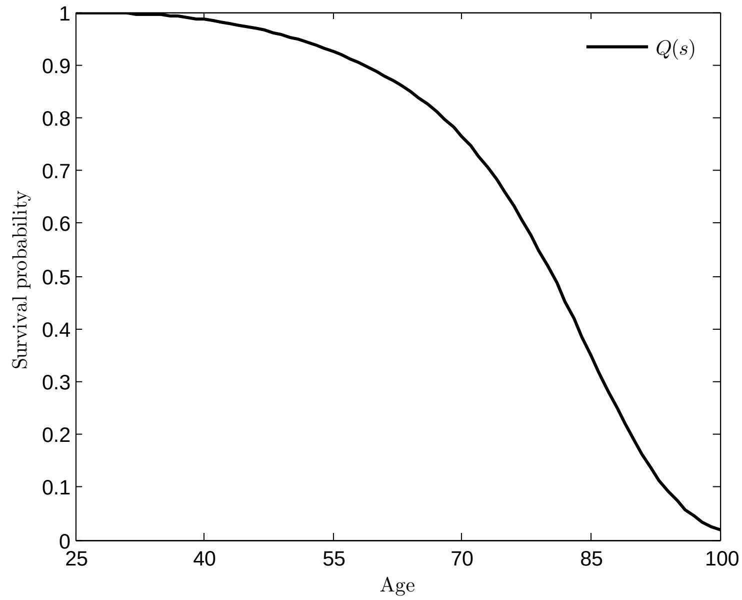

Figure 1: Survival probabilities from Feigenbaum's (2008a) sextic fit to the mortality data in Arias (2004).

where s is model age. The 2001 U.S. Life Tables in Arias (2004) are reported up to actual age 100, so I set the maximum model age to $T = 75$. I plot the resulting survivor function in Figure 1. Under these survival probabilities, the model life expectancy at birth turns out to be 78.6 years, which is slightly higher than the current projection of 77.9 years by the Centers for Disease Control and Prevention (CDC).[9] This divergence is because households in the model survive to age 25 with certainty (as the model age of zero corresponds to the actual age of 25), whereas in the real world they survive to age 25 with roughly 98% probability.

I set the population growth rate to $n = 1\%$, which is consistent with the U.S. demographic history and also with the literature. This population growth rate, along with the above survival probabilities, yields an aged-dependency ratio of 29% in the model, which is slightly higher than the 23% for the U.S. This should not be surprising, as the working-age population in the U.S. is measured as the population between ages 20 to 64, but in the model it is measured as the population between ages 25 to 64.

4.2 Social Security

I calibrate the Social Security program in the model to closely match the U.S. program. First, to compute the base Social Security benefit (also known as the Primary

[9]See http://www.cdc.gov/nchs/fastats/lifexpec.htm.

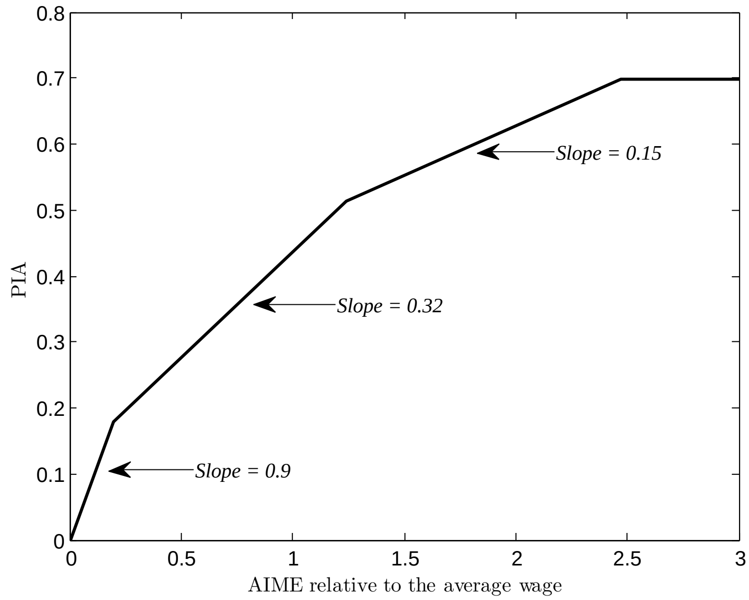

Figure 2: Benefit formula in the U.S.

Insurance Amount or 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 work-life income. The Social Security Administration measures what is known as the Average Indexed Monthly Earnings (AIME) for every covered individual, and then 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, there is an adjustment in the fraction of the AIME replaced by Social Security. For example, in the year 2000, the OASI benefit was 90% of the AIME for the first \$531, 32% of the next \$2671, and 15% of the remaining up to the maximum creditable 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. Note that the progressivity in the benefit rule is captured by the fact that the replacement rate is decreasing in the AIME (see Figure 2).

Second, I set the early, normal, and the maximum retirement ages in the model to $T_e = 37$, $T_n = 41$, and $T_m = 45$, which correspond to the actual ages of 62, 66, and 70 in the U.S. Based on when a household starts collecting Social Security benefits, there is a permanent adjustment in the base PIA. There is an early retirement penalty of 25%, 20%, 13.3%, or 6.7% if collection begins at age 62, 63, 64, or 65 respectively, and there is a delayed retirement credit of 8% for every single year of delaying benefit collection beyond the normal retirement age up to the maximum retirement age.

Third, households are also subject to an earnings test if they start collecting benefits before the normal retirement age. In 2008, the earnings threshold was \$13,560, and benefits were reduced by \$1 for every \$2 of earnings in excess of this threshold. Relative to the average household income in U.S. in 2008, this earnings threshold turns out to be about 27%, following which I set $\bar{e} = 0.27$. Then, I impose a tax of 50% to any benefits above this threshold at every age that the household collects benefits before the normal retirement age.

Finally, I set the payroll tax rate for Social Security to the current full OASI rate of $\tau_{ss} = 10.6\%$.

4.3 Labor productivity endowment

To calibrate the labor income process, I assume that the productivity endowment at age s can be multiplicatively decomposed into a deterministic age-dependent component $\epsilon(s)$, and a stochastic productivity shock φ as follows:

$$e(\varphi, s) = \varphi\epsilon(s). \tag{22}$$

I parameterize the age-dependent component $\epsilon(s)$ using hourly earnings data from Hansen (1993). However, as it is well-known, productivity measured from wage data suffers from sample selection bias, especially at the later ages when a large number of households begin to retire. For this reason, I fit a quartic polynomial to the log of the efficiency data in Hansen (1993) only for ages 25-65, which gives

$$\begin{aligned} \ln\epsilon(s) \quad = \quad & -3.273 \times 10^{-5} + \left(3.7484 \times 10^{-2}\right) s + \left(-1.7541 \times 10^{-3}\right) s^2 \\ & + \left(3.4625 \times 10^{-5}\right) s^3 + \left(-2.7949 \times 10^{-7}\right) s^4 \end{aligned} \tag{23}$$

where s is model age and $s \leq 40$. Beyond actual age 65 (i.e. for $s > 40$), I use the following quadratic function

$$\ln\epsilon(s) = -f_0 - f_1 s - 0.01s^2 \tag{24}$$

and parameterize f_0 and f_1 such that $\epsilon(s)$ is continuous and once differentiable at age $s = 40$.[10] The resulting age-dependent component of labor productivity is plotted in Figure 3.

To calibrate the distribution of the stochastic productivity shock, further assumptions are needed regarding how and when this shock is realized. Broadly speaking, there are two alternative ways of modeling labor productivity shocks in the literature. In the "permanent" specification, the shocks are realized before an agent enters the model (ex-ante), and they are preserved perfectly over the life cycle (Zhao, 2014). Alternatively, in the "idiosyncratic" specification, the shocks are

[10] The values that satisfy these conditions are $f_0 = 15.4789$ and $f_1 = -0.7918$.

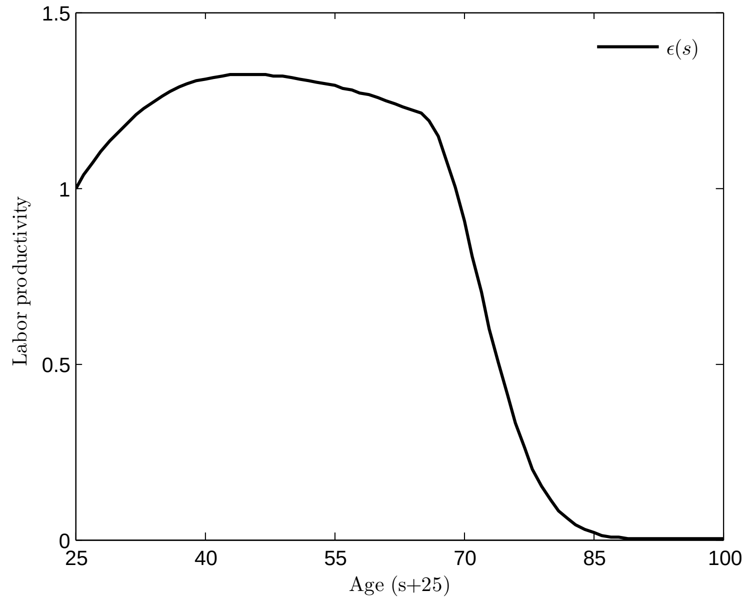

Figure 3: The age-dependent component of labor productivity estimated from Hansen (1993).

realized after the agent enters the model (ex-post), and they may or may not be preserved over time.[11] The key difference between these two specifications is that in the ex-ante and permanent case, households accumulate wealth due to purely life-cycle motives, whereas in the ex-post and idiosyncratic case, there are also precautionary motives at work.

The choice between these two alternative specifications for the productivity shock essentially depends on the frequency of consumption smoothing relevant for the question at hand. While the ex-post specification is more suitable for studying intertemporal allocation at medium (year to year or across the business cycle) or low (across the working life) frequencies, the ex-ante specification is sufficient in the current context. This is because in an environment without borrowing constraints, Social Security matters for consumption smoothing only at a very low frequency: across the working and retirement phases of the life cycle. While precautionary saving motives have been important in explaining the empirical phenomenon of consumption expenditures tracking income (Nagatani, 1972; Skinner, 1988; Gourinchas and Parker, 2002; Feigenbaum, 2008b), and also generating realistic wealth distributions from life-cycle consumption models (Hubbard and Judd, 1987; Huggett, 1996), it has been shown that even in a model with borrowing constraints, the manner in which

[11]See, for example, studies such as İmrohoroğlu et al. (1995), Huggett (1996), and İmrohoroğlu and Kitao (2009).

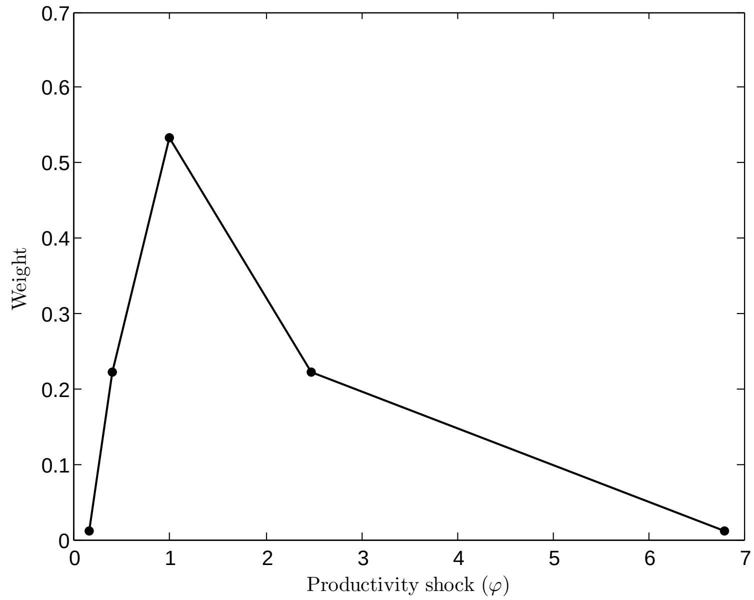

Figure 4: The distribution of the stochastic productivity shock.

the labor productivity shocks are modeled has little effect on the macroeconomic and distributional effects of Social Security (Huggett and Ventura, 1995, 1999).

I assume that the stochastic productivity shock is realized before the households enter the model and is permanent in nature, i.e. $\Phi(s) = \Phi(0)\ \forall s$ and

$$Prob\left(\varphi' = \varphi_j \mid \varphi = \varphi_i\right) = \begin{cases} 1 & \text{if} \quad i = j \\ 0 & \text{if} \quad i \neq j \end{cases} \tag{25}$$

To calibrate its distribution, I follow Zhao (2014) and assume that $\ln \varphi \sim N\left(0, \sigma_\varphi^2\right)$, and then set $\sigma_\varphi^2 = 0.45$, which is the variance of log male annual wages in Heathcote et al. (2010). Also, I use Gaussian quadrature to transform the continuous distribution into a 5-point discrete distribution for computational convenience (see Figure 4).

4.4 Tax function

With the assumed income tax function in (3), 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 pre-tax income, and finally 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:

$$\hat{\tau_1} = 0.06411. \tag{26}$$

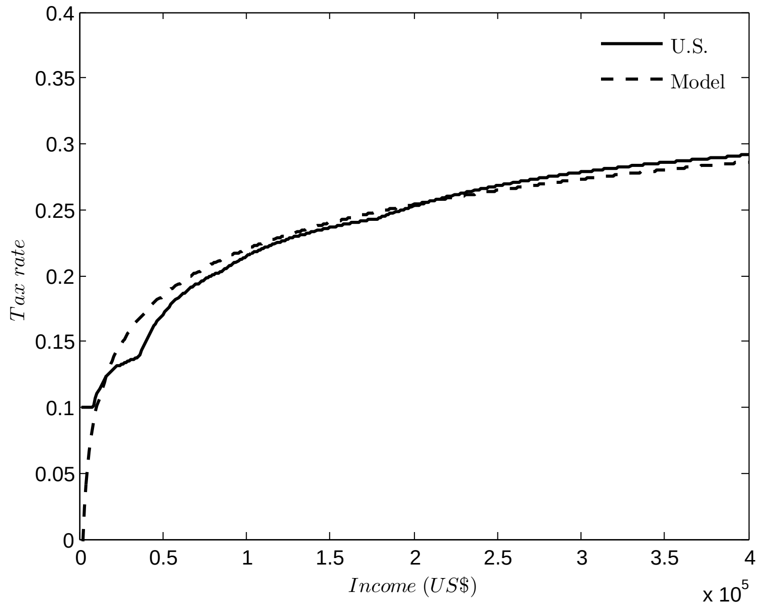

Figure 5: The labor income tax schedule.

I plot the estimated tax function along with the U.S. schedule in Figure 5.

4.5 Technology

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

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 to match certain macroeconomic targets. First, so that overall wealth accumulation in the model matches the U.S. economy, I target 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 34 hours per week spent on market work between ages 25 to 55. I adopt both of these empirical facts as targets. Finally, to pin down the tax parameters τ_y and τ_k, I assume that $\tau_y = \tau_k$ and target a share of government expenditures in GDP of 15%.

σ	β	η	$\tau_y (= \tau_k)$	θ_P
4	0.9866	0.295	0.176	0.15

Table 1: Unobservable parameter values under the baseline calibration.

	Target	Model
Capital-output ratio	3.0	2.98
Avg. hours of market work per week between ages 25-55	34	32.7
Share of govt. expenditures in GDP	0.15	0.16

Table 2: Model performance under the baseline calibration.

5 Baseline results

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 $\sigma^c = 1 + \eta(\sigma - 1) = 1.89$, which lies within the range frequently encountered in the literature. The model-generated values for the targets under the baseline calibration are reported in Table 2, and the cross-sectional labor force participation and labor hours data (conditional on participation) are reported in Figures 6 and 7.[12] Note that in calculating labor hours per week, I net out eight hours per day as sleep time.

It is clear from Figures 6 and 7 that the model does a reasonable job of matching observed labor supply behavior in the U.S., both along the extensive and the intensive margins. As for the extensive margin, the model replicates the rapid decline in labor force participation between ages 55 and 70 quite well. However, it slightly overestimates the participation rates prior to age 62, and underestimates the participation rates after age 62. This is because of two reasons. First, prior to age 62, market work is the only source of income for households in the model, whereas in reality households also receive various forms of assistance from the government that are not tied to work, such as food stamps and supplemental security income. Due to this reason, labor force participation rates between ages 30-50 are close to 98% in the model, whereas in the U.S. they range between 87-90%.

Second, the model underestimates participation rates after age 62 because it abstracts from illiquid private pension wealth. As French (2005) points out, two most important determinants of the labor force exit decision are the accrual rates of illiquid private pension wealth, and the tax disincentives associated with collecting Social Security benefits before the normal retirement age. In most private pension plans, there is a formula component that increases pension wealth with age, up to the normal retirement age (Gustman et al., 2000). This incentive, which delays

[12] The cross-sectional mean of a variable $x(t, t-s; \varphi)$ is calculated using the formula $\bar{x}(t, t-s) = \int_{\Phi(s)} x(t, t-s; \varphi)\, \mathrm{d}\varphi$.

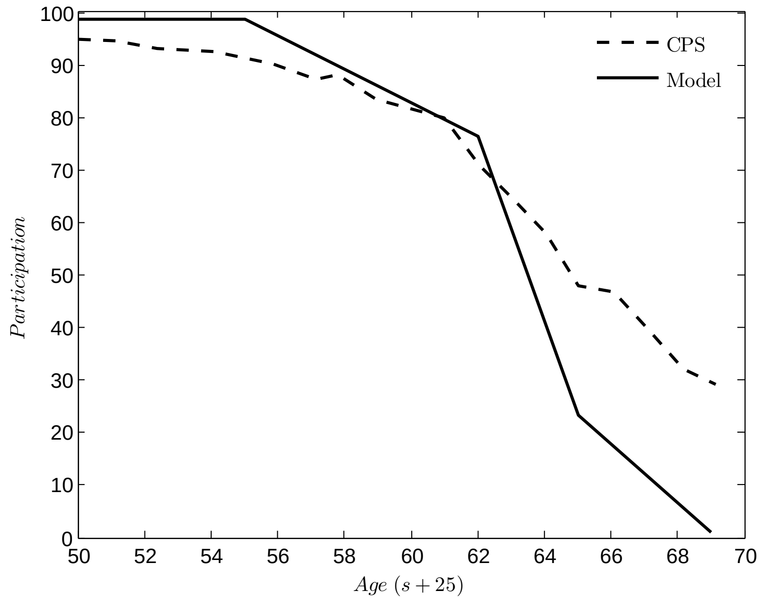

Figure 6: Cross-sectional labor force participation rates under the baseline calibration.

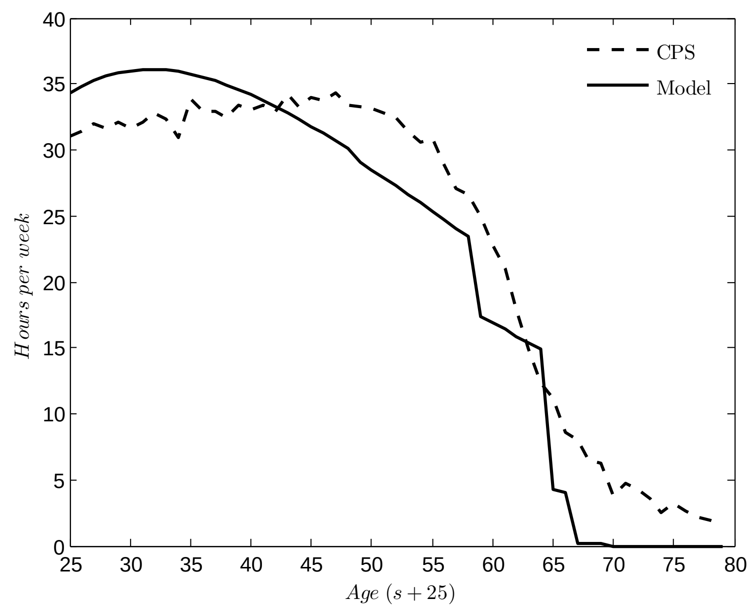

Figure 7: Cross-sectional mean of labor hours per week (conditional on participation) under the baseline calibration.

the labor force exit decision beyond the early retirement age of 62, is absent in the current model.[13]

As for the intensive margin of labor supply or the hours per week (conditional on participation), Figure 7 shows that the baseline hours profile initially increases, but then declines with age, with more rapid declines between ages 60 and 70. While this captures the overall pattern of hours over the life cycle observed in the data quite well, the model hours profile peaks relatively early − at age 35, compared to about age 45 in the data. This is because the shape of the age-dependent component of labor productivity $\epsilon(s)$ is the single-most important determinant of the shape of the hours profile generated from the model, and the productivity profile estimated from Hansen (1993) peaks relatively early. While it is standard practice in the literature to treat the age-dependent component $\epsilon(s)$ as observable, there are studies that treat this as an unobservable structural parameter, and calibrate it to match the shape of the hours profile in the data (Bullard and Feigenbaum, 2007; Bagchi and Feigenbaum, 2014). As one would expect, the calibrated productivity profiles in these studies peak significantly later in life (about age 52 or even later), relative to the productivity profile estimated from Hansen (1993), which peaks at age 42.[14]

The key factor behind the rapid decline in labor force participation beyond age 62 is the earnings test associated with working while collecting Social Security benefits. In Figure 8, I plot the fraction of households collecting Social Security benefits in the baseline equilibrium as a function of the collection age, along with the corresponding data from the SSA. Two key facts with respect to Social Security benefit collection in the U.S. are (*i*) about 45-47% of households start collecting at the early retirement age of 62, and (*ii*) almost 96% start collecting by the normal retirement age of 66. While the model matches the rapid increase in the number of households collecting Social Security benefits between ages 62-66, it somewhat underestimates the fraction of households collecting benefits at the early retirement age. This number is around 24% in the model, which is about half of that observed in the data. This is due to the fact that the model abstracts from an important determinant of the collection decision: the positive correlation between earnings and life expectancy observed in the U.S. (Waldron, 2005). The expected lifespan at age 62 is identical for poor and wealthy households in the model, but considerably less for the poorer households in reality. This causes households at the bottom of the income distribution in the model to delay collection until about age 70, which is

[13] Another factor that causes labor force participation rates in the baseline model to near zero by age 70 is the rapid decline in the age-dependent component of labor productivity $\epsilon(s)$ between ages 65 and 70. In Section 9, I compute an alternative version of the model in which this decline is less rapid. As one would expect, labor force participation rates in this version of the model decline to only about 24% by age 70.

[14] Because of the progressive labor income tax function, average tax rates over the life cycle start out low, peak roughly when hours peak, and then decline. In the baseline, the peak labor income tax rate is as low as 2.1% for households at the bottom of the income distribution, and is as high as 24% for households at the top of the income distribution.

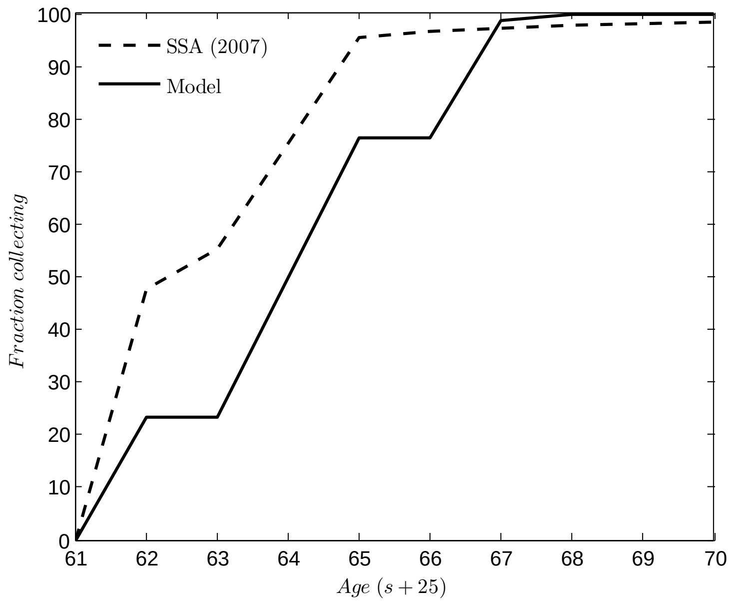

Figure 8: Cumulative proportion of households collecting Social Security benefits by age under the baseline calibration.

significantly later than when they start collecting in the real world.[15]

Finally, the cross-sectional age-consumption profile in the baseline calibration is reported in Figure 9. It is clear from the figure that consumption in the baseline equilibrium does not match data (Gourinchas and Parker, 2002) very well. This should not be surprising, though, as it is well-known that no fully rational model can replicate the empirical life-cycle consumption profile in Gourinchas and Parker (2002) in a general equilibrium with Social Security.[16] This is because Social Security reduces private saving and therefore the aggregate capital stock, which leads to significantly higher interest rates in general equilibrium (which is about 5.7% in the baseline model and 3.4% in Gourinchas and Parker (2002)). A higher interest rate causes consumption to increase much more rapidly in early life, because of which the peaks in life cycle consumption in the model are too large and occur much later than what is found in data. For example, cross-sectional mean consumption in the baseline calibration peaks at roughly age 58, with a ratio of peak to initial consumption of about 1.44. The empirical consumption profile in Gourinchas and Parker (2002) peaks much earlier roughly at age 50, and is much flatter with a peak-to-initial consumption ratio of about 1.1.

[15]See İmrohoroğlu and Kitao (2010) for a general-equilibrium model with health-dependent mortality and Social Security benefit claiming.

[16]Bullard and Feigenbaum (2007) find a calibration that reasonably matches the empirical consumption profile, but using a model without Social Security.

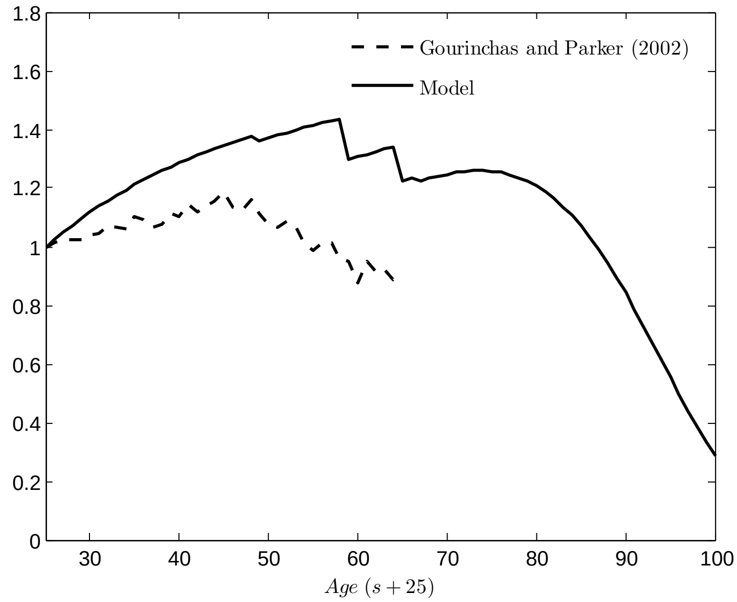

Figure 9: Cross-sectional mean of consumption (normalized by consumption at birth) under the baseline calibration.

6 The longevity improvement

A straightforward way to incorporate a one-time improvement in longevity in the current model is to reduce the baseline age-specific death rates based on the following formula:

$$h_n(s) = h(s) - \gamma s^{\nu}, \tag{27}$$

where $h(s)$ is the baseline death rate at age s, and γ and ν are parameters that can be chosen to match a specific life expectancy target. Note that these age-specific improvements in mortality replicate the fact that old-age survivorship in the U.S. has increased at a faster rate in the later half of the twentieth century, making the population survival curve more rectangular (Arias, 2004).

I choose values for γ and ν such that model life expectancy under the augmented survival probabilities matches the 2011 Trustees Report's average period life expectancy projection for the year 2085. The SSA reports the future financial status of Social Security under three alternative assumptions: low-cost, intermediate, and high-cost, of which, the intermediate assumptions reflect the SSA's Trustees' "best estimates of future experiences". To match the projected life expectancy of 85 years under the intermediate assumptions, I set $\gamma = 1 \times 10^{-5}$ and $\nu = 1.8509$. The survivor function corresponding to the new demographic parameters is compared to that in the baseline calibration in Figure 10.

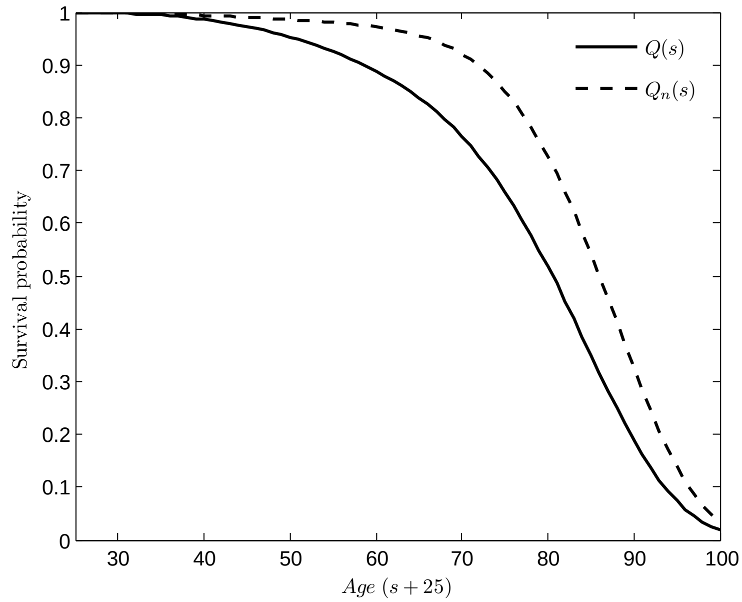

Figure 10: Baseline and the projected survival probabilities.

Based on the actuarial estimates in the 2011 Social Security Trustees' Report, non-interest income in the OASDI program is projected to be sufficient to pay only 75% of scheduled benefits in 2085 under the intermediate assumptions. This implies that if Social Security benefits and contributions continue to be based on current law, then by 2085 benefits will have to be reduced by 25% to keep the program solvent. It is useful to verify to what extent the current model matches this estimate of the Social Security crisis, if the SSA's actuarial methodology is applied on the current model. To do this, I calculate the percentage change in the Social Security benefits required to balance the program's budget with the projected survival probabilities, while holding all the household-level and macroeconomic variables fixed at their baseline equilibrium values (Case 1).[17]

It turns out that under Case 1, Social Security benefits in the model need to decline by about 23% to keep the program solvent with the projected survival probabilities. There is a small increase of 2% in Social Security's tax base under Case 1, but this is purely an effect of reduced mortality on the labor force.[18] On the other

[17]It is useful to note here that the baseline steady-state characteristics of the current model are also very similar to the assumptions used in the SSA's actuarial estimates. In the SSA's estimates, average hours worked are assumed to grow at the rate of 0.0% per annum, and the labor force participation rate is assumed to be roughly constant. Both of these conditions are satisfied under Case 1.

[18]Note from equation (18) that Social Security's tax base is nothing but the wage rate times aggregate labor.

	Baseline	Case 2
Labor	13.3	14.5
Capital	71.5	86.6
Real interest rate	5.74%	4.97%
Wage	1.17	1.21
Capital-output ratio	2.98	3.19

Table 3: Macroeconomic variables in the initial baseline, and also under Case 2.

hand, if Social Security benefits and contributions continue at their baseline equilibrium rates, the longevity improvement causes total Social Security expenditures under Case 1 to increase by 32%. The only way to manage this fiscal imbalance is to permanently reduce the benefits by 23%. Given that this is very close to the 25% decline estimated by the SSA, we can conclude that the current model does a good job of matching the SSA's projections of the Social Security crisis in 2085 using their actuarial methodology.[19]

7 The general-equilibrium estimates

To examine how accounting for the endogenous household-level and macroeconomic adjustments to the longevity improvement affect the measured extent of the Social Security crisis, I incorporate the projected survival probabilities into the baseline model, but this time compute the budget-balancing change in Social Security benefits (from the baseline) in general equilibrium (Case 2). In Table 3, I report the equilibrium values of relevant macroeconomic variables, both in the initial baseline and with the new survival probabilities. It is clear from the table that the improvement in longevity leads to a significant increase in labor supply, aggregate capital stock, and the wage rate, and a decline in the real interest rate. Note that this expansion in labor supply consists of adjustments both along the extensive, as well as the intensive margins. Aggregate output also increases, but the increases in the capital stock are proportionally larger, because of which the capital-output ratio increases from the baseline. Together, the increase in the wage rate and labor supply under Case 2 lead to a 13% increase in the size of Social Security's tax base, which is significantly larger than the 2% increase under Case 1.

With the SSA's actuarial methodology, the only factor relevant for the level of Social Security benefits is how the longevity improvement increases the number of surviving retirees in the economy. However, when one accounts for the household-level and macroeconomic adjustments to the longevity improvement, the effect on Social Security benefits is more complicated. In this case, Social Security benefits

[19]Note that because the actuarial methodology ignores how households respond to the improvement in longevity, the decline in Social Security benefits is identical for *all* household types under Case 1.

φ	0.15	0.40	1	2.48	6.79
Case 1	-23%	-23%	-23%	-23%	-23%
Case 2	-14.9%	-9.8%	7.5%	-13.0%	31.7%

Table 4: Percentage changes in Social Security benefits under Cases 1 and 2 for each household type.

depend not only on the number of surviving retirees, but also on how the tax base and the benefit collection decision are affected by the longevity improvement.

In Table 4, I compare the percentage change in Social Security benefits (from the initial baseline) for each household type under Cases 1 and 2. As the table shows, there is considerable variation in how the longevity improvement affects the level of Social Security benefits under Case 2. The largest decline in Social Security benefits is experienced by the households at the bottom of the income distribution ($\varphi = 0.15$): their benefits decline by about 15% from the baseline level, whereas households at the top of the income distribution ($\varphi = 6.79$) actually experience an increase (roughly 32%). However, even the largest decline in benefits under Case 2 (15%) is only two-thirds of the decline under Case 1 (23%). Averaging across all households at an age when everyone collects benefits, the decline under Case 2 is only 0.9%, which is even less than a percentage point.

Because the results in Table 4 are jointly determined by the effect of the longevity improvement on the size of the tax base, and also on the date of benefit collection, it is useful to try and understand the importance of each mechanism individually. First, note that the 13% increase in the tax base under Case 2 can be decomposed into a 9.2% increase in labor supply from the baseline, which includes adjustments both along the extensive and the intensive margin, and a 3.7% increase in the wage rate. In Figures 11 and 12, I compare the cross-sectional labor force participation rates, and the weekly hours (conditional on participation) under Case 2 with those in the baseline.

First, let us consider the extensive margin of labor supply, or the labor force participation decision (Figure 11). It is clear from the figure that a large number of households in the model delay retirement as a result of the longevity improvement. Households do not begin to exit the labor force until age 61 under Case 2, which is significantly later than age 55 in the baseline. Also, because the participation rate drops to almost zero by age 70, the decline in labor force participation is more rapid in this case. On the average, retirement is delayed by about two years, with households in the top half of the income distribution experiencing the larger delays.

Second, consider the intensive margin of labor supply, or the hours per week of market work, conditional on labor force participation (Figure 12). The figure demonstrates a clear increase in the weekly hours between ages 33-70: on the average, labor supply increases by about 90 minutes per week over the life cycle. As in the baseline, we see a rapid decline in hours per week at later ages, but in this case

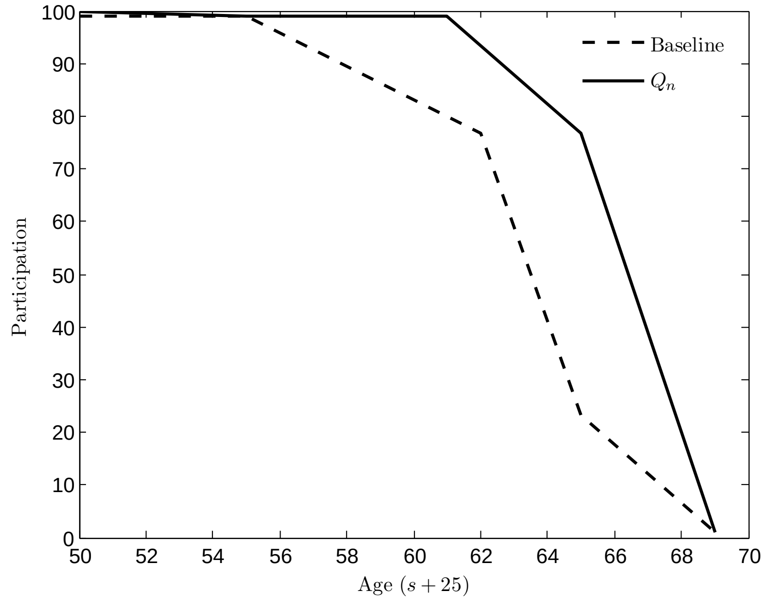

Figure 11: Cross-sectional labor force participation rates, baseline and under Case 2.

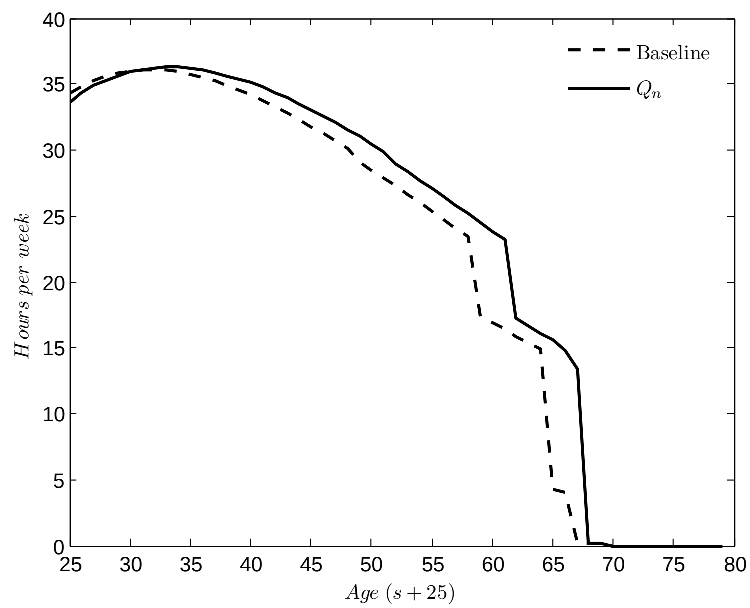

Figure 12: Cross-sectional means of hours per week (conditional on participation), baseline and under Case 2.

the decline does not begin until age 62. Even though the age-dependent component of the labor productivity $\epsilon(s)$ starts to decline after age 42, delayed retirement and the increased hours per week in this range deliver an overall increase of 9.2% in labor supply under Case 2.

If the capital stock were fixed at the baseline level, this increase in labor supply would reduce the equilibrium wage rate by about 3%. Taken together, these changes would lead to an expansion of only 6% in Social Security's tax base. However, capital stock also increases from the baseline, as household saving also responds to the improved longevity. A *ceteris paribus* improvement in longevity encourages households to save more due to life-cycle motives. This effect is weakened when labor supply responds positively, but the positive income effect from the increased labor supply encourages further saving. Finally, if a household experiences a decline in their Social Security benefits, then life-cycle motives encourage them to save even more. Averaging across all ages, these effects collectively increase the mean cross-sectional asset holdings by about 16% from the baseline under Case 2. These changes in asset holdings translate into a 21% increase in the aggregate capital stock, and a 3.7% increase in the equilibrium wage rate. Taking the labor supply and wage rate adjustments together, Social Security's tax base increases by roughly 13% under Case 2.

While the household-level labor supply and saving decisions affect the tax base, or the income side of Social Security, the decision to start collecting Social Security benefits affects the expenditure side of the program. On the one hand, delaying collection reduces the number of retirees who have to be paid Social Security benefits, but on the other hand, the actuarial adjustment increases the benefit paid to each retiree who delays collection. To understand the effect of the longevity improvement on the collection decision, I compare in Figure 13 the fraction of households collecting Social Security benefits at every age in the baseline with those under Case 2.

It is clear from the figure that a large number of households delay collecting Social Security benefits with the longevity improvement. The fraction of households who start collecting at the early retirement age of 62 under Case 2 is roughly identical to the baseline, but it remains unchanged until after the normal retirement age. At age 67, only 22% of the households collect Social Security benefits under Case 2, which is significantly smaller than the 98% in the baseline equilibrium, but after age 67 this fraction increases rapidly and reaches 100% by age 70.[20]

Decomposing the collection decision by household type, I find that while there is no effect on the households with $\varphi = 0.15$ and 2.48, the longevity improvement delays collection by one, four, and seven years for households with $\varphi = 0.40$, 1, and 6.79 respectively. Given this fact, it is not surprising that these households experience the smallest declines in their Social Security benefits under Case 2 (in fact, benefits actually increase from the baseline for $\varphi = 1$ and 6.79). For these

[20]The response to delay collecting Social Security benefits due to the longevity improvement that I find here is consistent with İmrohoroğlu and Kitao (2010).

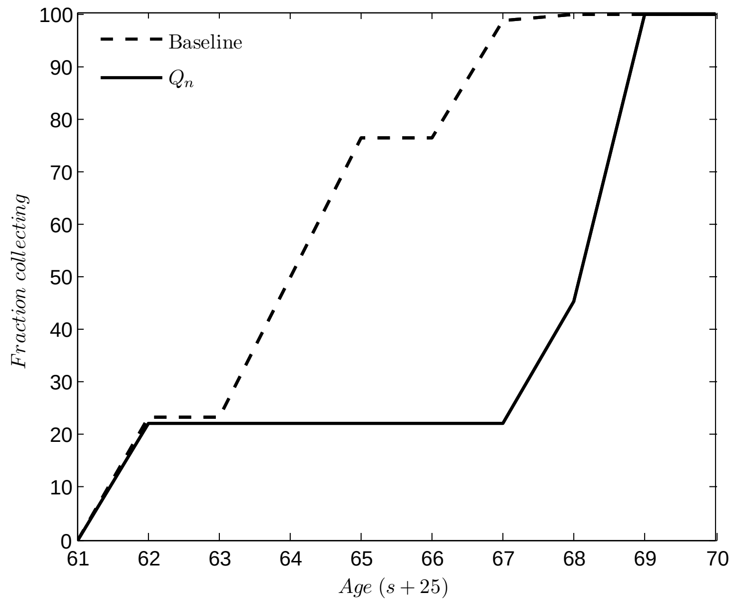

Figure 13: Fraction of retirees collecting benefits, baseline and under Case 2.

households, the delayed retirement credit associated with delaying collection somewhat compensates for the decline in Social Security benefits, and in some cases even reverses it.

To summarize, the above experiment demonstrates that in general equilibrium, the improvement in longevity causes households to delay retirement and work more hours, save more over the life cycle, and also delay collecting Social Security benefits. Accounting for all of these responses, the largest decline in Social Security benefits is only two-thirds of the amount estimated by the SSA.

8 Improved longevity, or shrinking Social Security?

While the above results come from a general-equilibrium model calibrated to the U.S. economy, it would appear to the casual observer that the model's predictions with respect to household behavior are inconsistent with data. In the model, the longevity improvement causes delayed retirement, increases weekly hours, and also increases saving. However, historical U.S. time-series data shows that while there has been a steady improvement in life expectancy in the later half of the twentieth century (Arias, 2004), old-age labor force participation rates have actually declined during this period (Fullerton Jr., 1999). Between 1950-1970, there was a small increase in the overall participation rates between ages 55-64 (from 57% to 62%), but this was mostly due to the fact that the participation rates for women in this age bracket increased rapidly: from 27% in 1950 to 43% in 1970. During this same period, the overall participation rates for ages 65 and above decreased from 27% to 17%, and continued to decline to about 12% in 1998.[21]

It turns out that this inconsistency is due to the fact that the experiment defined under Case 2 is quite different from the U.S. economic experience in the later half of the twentieth century. In the U.S., this period was characterized by improving life expectancies, and also by an expansion of Social Security, such as a gradual increase in the payroll tax from a combined rate of about 5% in 1960, to about 9% in 1980 and 10.6% in 2000, and the introduction of the automatic cost-of-living adjustments (COLA) starting in 1975. In contrast, in the experiment defined under Case 2, there are no changes to the institutional features of Social Security, and only longevity improves. Holding everything else fixed, this longevity improvement actually makes Social Security less generous relative to the baseline. The general-equilibrium adjustments identified under Case 2, therefore, capture the effect of the longevity improvements on household behavior through two different channels: a primary effect of the longevity improvement itself, and a secondary effect due to shrinking Social Security benefits.

In this section, I compute a new partial equilibrium experiment to separate these two effects. In this experiment, I hold the wage rate, the interest rate, and also the Social Security benefits fixed at the initial baseline level, but allow the households to respond along the labor supply and saving margins to the improvement in longevity (Case 3).[22] I compare the labor force participation rates, as well as the weekly hours of work (conditional on participation) under this experiment and the baseline equilibrium, in Figures 14 and 15 respectively.

As Figure 14 demonstrates, this experiment yields old-age labor force participation rates that are consistent with the historical trend of labor force participation in the U.S. Under Case 3, the participation rate in the model declines from 98% to 77% at age 58, and from 77% to 23% at age 62. However, while labor force participation declines under this experiment, households' weekly hours (conditional on participation) respond positively (see Figure 15). It is clear from the figure that there is an increase in the weekly hours prior to age 55 under Case 3, and the rapid decline in labor force participation causes hours to decline much faster at later ages. Because the increased hours occur at ages during which productivity is relatively high, overall labor supply declines only marginally $-$ by about a third of a percentage point, in spite of the fact that households exit the labor force much earlier under this experiment.[23]

[21] It is useful to note here that while the personal saving rate measured by the National Income and Product Accounts (NIPA) ranged at about 8-10% between 1960 and 1980, it declined to almost 0% by 2000. Lansing (2005) attributes this decline in the NIPA personal saving rates to measurement errors, such as ignoring capital gains on stock portfolios and real estate, and counting employer pension contributions as income, even though they are not readily available to households.

[22] I also hold the benefit collection decisions fixed at their baseline levels in this experiment.

[23] Because labor supply does not change much, household saving responds much more strongly to

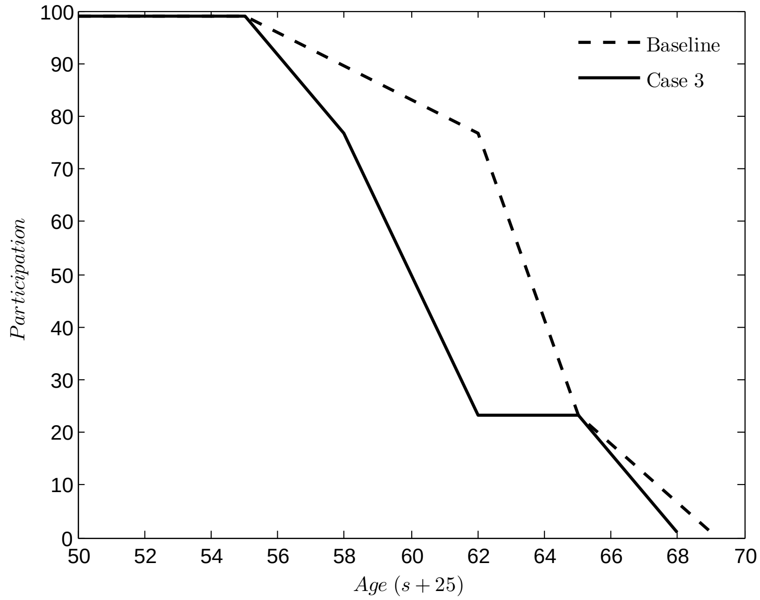

Figure 14: Cross-sectional labor force participation rates, baseline and under Case 3.

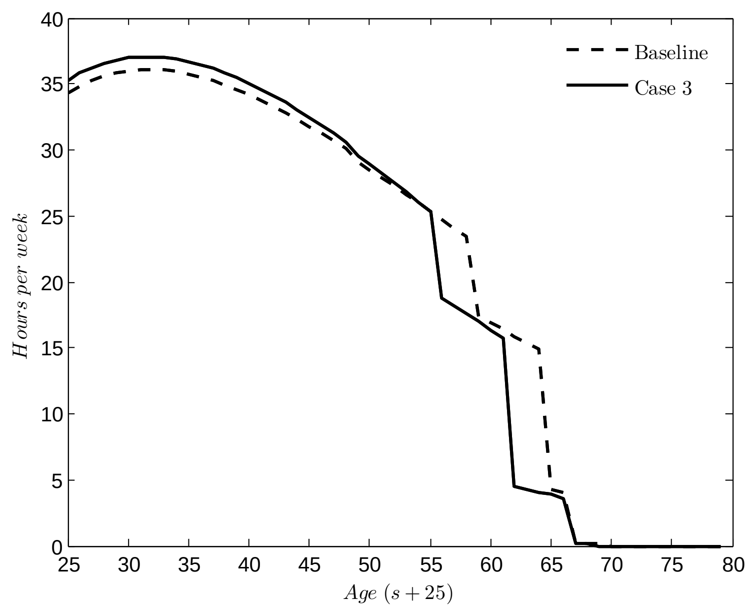

Figure 15: Cross-sectional means of hours per week (conditional on participation), baseline and under Case 3.

Therefore, I find that a *ceteris paribus* improvement in longevity has a small negative effect on overall labor supply. Holding Social Security benefits constant at the baseline level, the longevity improvement by itself causes old-age labor force participation rates to decline, which matches the historical trend in the U.S. quite well, and weekly hours increase slightly from the baseline. Taken together, these findings indicate that it is the shrinking Social Security benefits, rather than the longevity improvement itself, that cause households to delay retirement under Case 2.

9 Improved longevity and health

While the primary focus of the current paper is to account for the effect of improved longevity on households' labor supply and saving decisions, one factor that has been ignored so far is the possible effect of improved longevity on households' health. First, there is considerable evidence that the health of the elderly population in the U.S. has improved considerably over time (Cutler and Landrum, 2012). Moreover, Halliday et al. (2009) find that better health allows households to allocate more time to productive activities, so holding life expectancy constant, households with better health will supply more labor. Therefore, there is a third reason due to which SSA's actuarial projections overestimate the Social Security crisis: they ignore the effect of longevity improvements on the future health status of the U.S. population.

In this section, I compute a new general-equilibrium experiment where I approximate the improvement in future health status by changing the age-dependent component of labor productivity $\epsilon(s)$ at older ages (Case 4). Specifically, I incorporate the new survival probabilities into the baseline model, and I assume that beyond age 65, productivity now declines linearly with age rather than quadratically, given by

$$\ln \epsilon(s) = -f_0 - f_1 s. \tag{28}$$

I parameterize f_0 and f_1 such that $\epsilon(s)$ is continuous and once differentiable at age $s = 40$.[24] I report this new endowment profile in Figure 16, and I compare the equilibrium values of key macroeconomic variables under this experiment with the baseline and Case 2 in Table 5. It is clear from the table that accounting for the health improvements has almost no additional effect on the macroeconomic variables. Across the baseline and Case 4, wages increase by 3.4% and labor supply increases by 9.4%, which leads to a roughly 13% increase in Social Security's tax base. This expansion is almost identical to that documented under Case 2.

the longevity improvement under Case 3. Averaged across all ages, cross-sectional asset holdings increase by 37% in this case, compared to only 16% under Case 2. This increase in asset holdings constitutes an overall increase of 42% in the supply of capital under Case 3, which is twice as large as the 21% increase in the equilibrium capital stock under Case 2.

[24]The values that satisfy these conditions are $f_0 = -0.5211$ and $f_1 = 0.0082$.

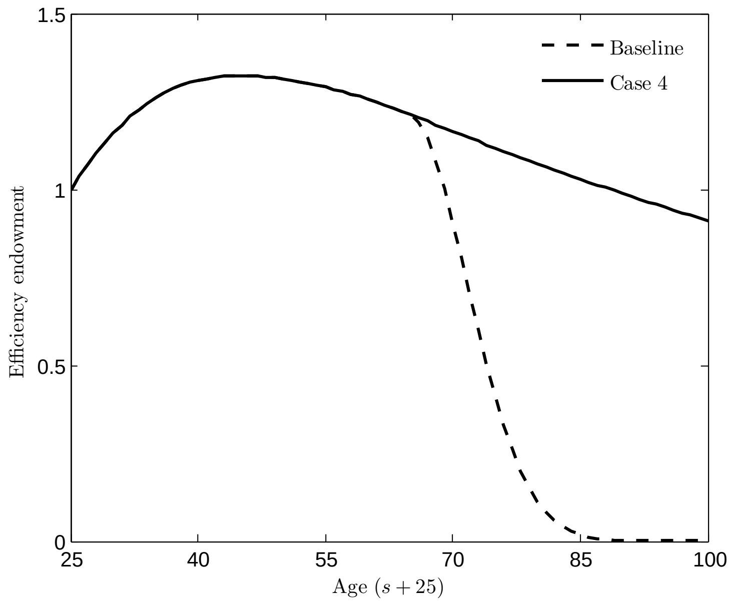

Figure 16: The age-dependent component of labor productivity, baseline and under Case 4.

	Baseline	Case 2	Case 4
Labor	13.3	14.5	14.6
Capital	71.5	86.6	85.5
Real interest rate	5.74%	4.97%	5.02%
Wage	1.17	1.21	1.21
Capital-output ratio	2.98	3.19	3.18

Table 5: Macroeconomic variables in the initial baseline, and under Cases 2 and 4.

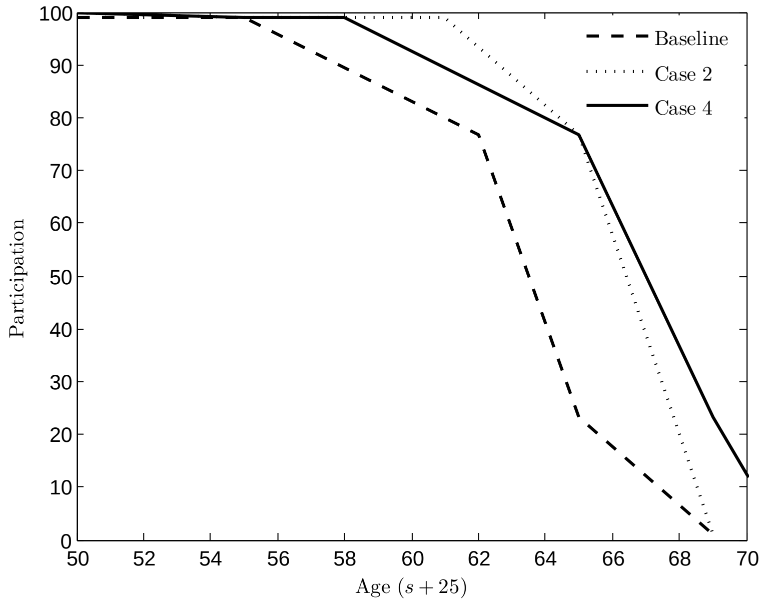

Figure 17: Cross-sectional labor force participation rates, baseline and under Case 4.

Even though accounting for the health improvements generates macroeconomic aggregates that are almost identical to those under Case 2, there are some differences in labor supply behavior across the two experiments, both across the extensive and the intensive margin. As Figure 17 demonstrates, the improvement in longevity leads to an increase in old-age labor force participation even under Case 4, but the improvement in health somewhat counteracts it. The participation rate under this experiment does not begin to decline until age 58, which is later than age 55 under the baseline calibration, but earlier than age 61 under Case 2. Also, the participation rate declines to near zero by age 69 under Case 2, but accounting for the health improvement, it declines to only about 24%. This fact is also reflected in the weekly hours, conditional on participation, under Case 4. Figure 18 shows that the hours per week are virtually indistinguishable between Cases 2 and 4 up to about age 67, but beyond that hours are slightly higher when the health improvement are accounted for. However, these differences do not appear to be macroeconomically significant.

Given that the expansion in Social Security's tax base under Case 4 is roughly identical that under Case 2, one would expect the Social Security benefits under this experiment to respond similarly. To verify this, I report the percentage changes in the Social Security benefits (from the baseline) under Case 4 in Table 6. The table shows that the overall pattern of changes in the Social Security benefits under this experiment is very similar to Case 2: households with $\varphi = 0.15$, 0.4, and

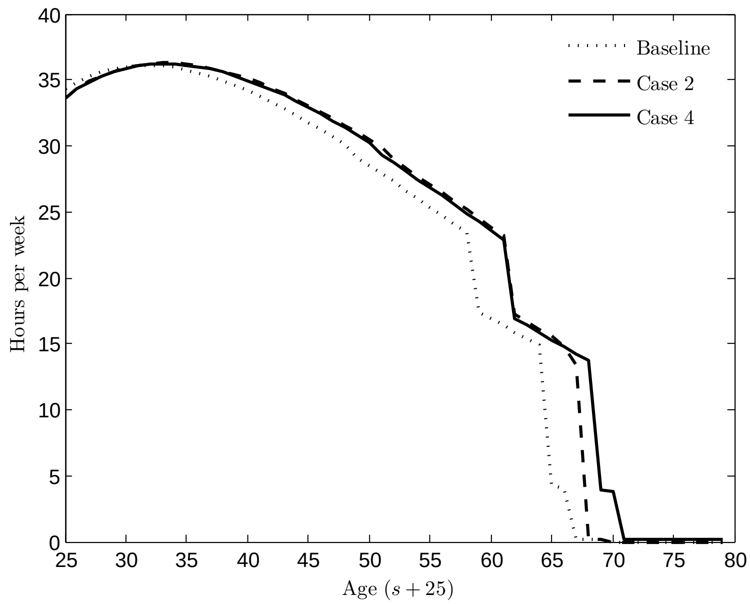

Figure 18: Cross-sectional means of hours per week (conditional on participation), baseline and under Case 4.

φ	0.15	0.40	1	2.48	6.79
Case 1	-23%	-23%	-23%	-23%	-23%
Case 2	-14.9%	-9.8%	7.5%	-13.0%	31.7%
Case 4	-11.3%	-3.7%	9.7%	-11.7%	32.3%

Table 6: Percentage changes in Social Security benefits under Cases 1, 2, and 4 for each household type.

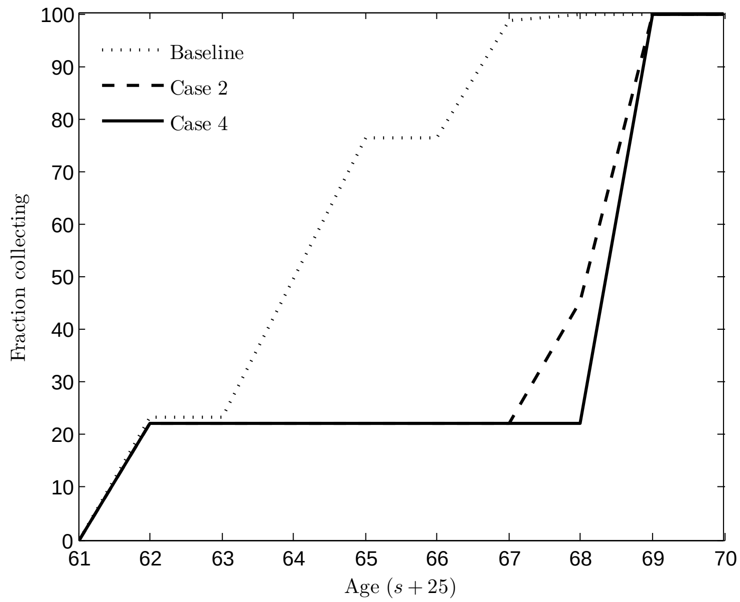

Figure 19: Cumulative proportion of households collecting Social Security benefits by age in the baseline, and under Cases 2 and 4.

2.48 experience a decline in their benefits, whereas households with $\varphi = 1$ and 6.79 experience an increase. However, accounting for the improvements in health leads to benefits declining by a smaller extent relative to Case 2. For example, benefits decline only by about 11% for the households at the bottom of the income distribution under Case 4, compared to almost 15% under Case 2. This decline in Social Security benefits, which is the largest under Case 4, is only half of the decline under Case 1 (23%).

Social Security benefits decline by a smaller percentage under Case 4 because households delay collection even further in this experiment. I report the fraction of households collecting Social Security benefits under Case 3 as a function of age in Figure 19. The figure shows that until age 68, the fraction of households collecting Social Security benefits stays stable at about 22% under Case 4. This is about half of that under Case 2, where the fraction collecting at age 68 is about 46%. However, between ages 68 and 69, the fraction of households collecting increases rapidly to 100. Decomposing by household type, those with $\varphi = 0.15$ delay collection by one year, those with $\varphi = 0.4$ by two years, those with $\varphi = 1$ by four years, and those with $\varphi = 6.79$ by seven years, relative to the baseline. The delayed retirement credits associated with these delays in collection cause the benefits to decline by a smaller percentage under Case 4.

To summarize, I find that while accounting for the improvements in health associated with the longevity improvements has roughly the same effect on Social

Security's tax base, the effect on benefits is slightly different. Benefits generally decline by a smaller percentage when the health improvements are accounted for, because households delay collection even further under this experiment.

10 Conclusions

The SSA's actuarial projections overestimate the Social Security crisis because they ignore the household-level and macroeconomic adjustments associated with improvements in longevity. In this paper, I show that accounting for these adjustments, a significantly smaller decline in benefits is needed to balance the Social Security budget. This is because households respond to a higher life expectancy by delaying retirement and Social Security benefit collection, working more hours, and by also saving more. In general equilibrium, these effects increase the aggregate labor supply, and also the capital stock and the wage rate. Collectively, these changes lead to a natural expansion of Social Security's tax base, and also generate significant delayed retirement credits, which the actuarial estimates completely overlook.

The future insolvency of Social Security in the U.S. has generated a lot of research on how the program can be reformed to be fiscally sustainable, while also providing much of the benefits that it currently provides to the general population. The mechanisms that I consider in this paper are a part of virtually every single macroeconomic model that has been used to evaluate such reform policies. My findings in this paper make a strong case in favor of including these very mechanisms in the statistics routinely used to measure the health of Social Security.

References

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

Bagchi, Shantanu and James Feigenbaum (2014), "Is smoking a fiscal good?" *Review of Economic Dynamics*, 17, 170–190.

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

Caliendo, Frank N., Nick L. Guo, and Roozbeh Hosseini (2014), "Social security is NOT a substitute for annuity markets." *Review of Economic Dynamics*, –.

Chen, Kaiji and Ayse İmrohoroğlu (2012), "Debt and the U.S. Economy." *Working Paper*.

Cogan, John F. (1981), "Fixed Costs and Labor Supply." *Econometrica*, 49, 945–963.

Conesa, Juan C. and Carlos Garriga (2003), "Status quo problem in social security reforms." *Macroeconomic Dynamics*, 7, 691–710.

Conesa, Juan C. and Carlos Garriga (2008), "Optimal Fiscal Policy In The Design Of Social Security Reforms." *International Economic Review*, 49, 291–318.

Cutler, David and Mary Beth Landrum (2012), *Dimensions of Health in the Elderly Population*, 179–201. University of Chicago Press for National Bureau of Economic Research, Chicago.

De Nardi, Mariacristina, Selahattin İmrohoroğlu, and Thomas J. Sargent (1999), "Projected U.S. Demographics and Social Security." *Review of Economic Dynamics*, 2, 575–615.

Dushi, Irena and Anthony Webb (2004), "Household Annuitization Decisions: Simulations and Empirical Analysis." *Journal of Pension Economics and Finance*, 3, 109–143.

Feigenbaum, James (2008a), "Can mortality risk explain the consumption hump?" *Journal of Macroeconomics*, 30, 844 – 872.

Feigenbaum, James (2008b), "Information shocks and precautionary saving." *Journal of Economic Dynamics and Control*, 32, 3917–3938.

French, Eric (2005), "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *The Review of Economic Studies*, 72, 395–427.

Fullerton Jr., Howard N. (1999), "Labor force participation: 75 years of change, 195098 and 19982025." *Monthly Labor Review*, 122, 3–12.

Gourinchas, P.-O. and J. A. Parker (2002), "Consumption over the Life Cycle." *Econometrica*, 70, 47–89.

Gustman, Alan L., Olivia S. Mitchell, Andrew A. Samwick, and Thomas L. Steinmeier (2000), *Evaluating Pension Entitlements*, 309–326. University of Pennsylvania Press for Pension Research Council, Philadelphia.

Halliday, Timothy, Hui He, and Hao Zhang (2009), "Health Investment over the Life-Cycle." *IZA Discussion Papers Series*, Working Paper 4482.

Hansen, G. D. (1993), "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked." *Journal of Applied Econometrics*, 8, 71–80.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010), "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006." *Review of Economic Dynamics*, 13, 15–51.

Huang, He, Selahattin İmrohoroğlu, and Thomas J. Sargent (1997), "Two Computations to Fund Social Security." *Macroeconomic Dynamics*, 1, 7–44.

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

Huggett, M. (1996), "Wealth distribution in life-cycle economies." *Journal of Monetary Economics*, 38, 469–494.

Huggett, Mark and Gustavo Ventura (1995), "Understanding why high income households save more than low income households." Discussion Paper / Institute for Empirical Macroeconomics 106, Federal Reserve Bank of Minneapolis.

Huggett, Mark and Gustavo Ventura (1999), "On the Distributional Effects of Social Security Reform." *Review of Economic Dynamics*, 2, 498–531.

İmrohoroğlu, Ayse, Selahattin İmrohoroğlu, and Douglas H Joines (1995), "A Life Cycle Analysis of Social Security." *Economic Theory*, 6, 83–114.

İmrohoroğlu, Selahattin and Sagiri Kitao (2009), "Labor supply elasticity and social security reform." *Journal of Public Economics*, 93, 867–878.

İmrohoroğlu, Selahattin and Sagiri Kitao (2010), "Social Security, benefit claiming, and labor force participation: a quantitative general equilibrium approach." Technical Report 436, Federal Reserve Bank of New York.

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

Karabarbounis, Marios (2012), "Heterogeneity in labor supply elasticity and optimal taxation." In *2012 Meeting Papers*, 655, Society for Economic Dynamics.

Kitao, Sagiri (2012), "Sustainable Social Security: Four Options." *Working Paper*.

Lansing, Kevin J. (2005), "Spendthrift Nation." *Federal Reserve Bank of San Francisco Economic Letters*.

Ljungqvist, Lars (2008), "Lucas critique." In *The New Palgrave Dictionary of Economics* (Steven N. Durlauf and Lawrence E. Blume, eds.), Palgrave Macmillan, Basingstoke.

Lockwood, Lee M. (2012), "Bequest motives and the annuity puzzle." *Review of Economic Dynamics*, 15, 226–243.

Mitchell, Olivia S., James M. Poterba, Mark J. Warshawsky, and Jeffrey R. Brown (1999), "New Evidence on the Money's Worth of Individual Annuities." *The American Economic Review*, 89, 1299–1318.

Nagatani, Keizo (1972), "Life Cycle Saving: Theory and Fact." *The American Economic Review*, 62, 344–353.

Ortiz, Jorge Alonso (2009), "Social Security and Retirement across OECD Countries." *Working Paper*.

Pashchenko, Svetlana (2013), "Accounting for non-annuitization." *Journal of Public Economics*, 98, 53–67.

Prescott, Edward C., Richard Rogerson, and Johanna Wallenius (2009), "Lifetime aggregate labor supply with endogenous workweek length." *Review of Economic Dynamics*, 12, 23–36.

Rogerson, Richard and Johanna Wallenius (2009), "Micro and macro elasticities in a life cycle model with taxes." *Journal of Economic Theory*, 144, 2277–2292.

Shrestha, Laura B. (2006), "Life Expectancy in the United States." CRS Report for Congress RL32792, Congressional Research Service.

Skinner, Jonathan (1988), "Risky income, life cycle consumption, and precautionary savings." *Journal of Monetary Economics*, 22, 237–255.

Stokey, Nancy L. and Sergio Rebelo (1995), "Growth Effects of Flat-Rate Taxes." *Journal of Political Economy*, 103, 519–550.

Storesletten, Kjetil, Gianluca Violante, and Jonathan Heathcote (2012), "Redistributive Taxation in a Partial Insurance Economy." 2012 Meeting Papers 588, Society for Economic Dynamics.

Turra, Cassio and Olivia S. Mitchell (2004), "The Impact of Health Status and Out-of-Pocket Medical Expenditures on Annuity Valuation." *Michigan Retirement Research Center Research Paper*, WP 2004-086.

Waldron, Hilary (2005), "Heterogeneity in Health and Mortality Risk Among Early Retiree Men." Working Paper 105, Office of Policy, Office of Research, Evaluation, and Statistics, Social Security Administration.

Wallenius, Johanna (2009), "Social Security and Cross-Country Hours: A General Equilibrium Analysis." *Working Paper*.

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