Can Removing the Tax Cap Save Social Security?*

Shantanu Bagchi†

December 29, 2016

Abstract

The maximum amount of earnings in a calendar year that can be taxed by Social Security is currently set at $118,500. In this paper, I examine if removing this cap can solve Social Security's future budgetary problems. Using a calibrated general-equilibrium life-cycle consumption model, I show that when this cap is removed, benefits need to decline by less than 4% to keep Social Security solvent, compared to by almost 12% when the cap is held fixed at its current level. Households for whom the cap expires respond by working and saving less, which reduces labor supply, capital stock, and output, and also reverses some of the initial expansion in Social Security's revenues. Elimination of the cap alters the pattern of redistribution implicit in Social Security, and also imposes larger distortions on labor supply and saving, which reduces overall welfare.

JEL Classifications: E21, E62, H55

Keywords: Social Security; tax cap; mortality risk; labor income risk; incomplete markets; general equilibrium

*I would like to thank Frank Caliendo, Matt Chambers, Jim Feigenbaum, Aspen Gorry, Juergen Jung, and participants at the 2014 SEA Meetings and the 2015 NTA Meetings for their useful comments and suggestions.

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

1 Introduction

Mitigating the effect of longevity improvements on social insurance programs is currently a major policy concern in the developed world. Life expectancy in the U.S. is projected to increase to almost 85 years by the end of the century (Shrestha, 2006), which will increase the current dependency ratio from 22% to almost 45% by 2100.[1] Because of these demographic developments, the current payroll tax rate for the Old-Age and Survivors Insurance and the Federal Disability Insurance (OASDI) programs will be sufficient to pay only 79% and 73% of the scheduled benefits in 2034 and 2089 respectively.[2]

This projected insolvency of America's social security program, commonly known as the Social Security crisis, has engendered a large amount of theoretical and quantitative research over the last few decades. Focusing specifically on the Old-Age and Survivors Insurance (OASI) program, this literature has rigorously examined a number of reform policy options, such as changing Social Security's contribution rate, the retirement age, the link between Social Security benefits and contributions, and also a phased transition to a fully funded system.[3]

This literature, however, has been silent on the ability of Social Security's tax cap in improving the program's projected fiscal status. The maximum amount of earnings in a calendar year that is taxed by Social Security is currently set at $118,500. This limit, adjusted regularly in proportion to wage growth, has been a salient feature of the OASI program since its inception. While Social Security's projected insolvency and growing income inequality in the U.S. (Heathcote et al., 2010) has renewed calls for abolishing this cap and subjecting all earnings to the payroll tax, the economic implications of such a policy move are largely unknown.[4]

While removing Social Security's tax cap and subjecting all earnings to the payroll tax is likely to boost the program's revenues, there are at least two reasons why the overall effect of this policy change on Social Security's fiscal status is not clear. First, Social Security benefits are calculated based on a measure of average earnings through the work life, and earnings only up to the cap are counted towards the benefits. If this historical link between the cap on taxes and creditable earnings is retained, then removing the cap will also increase Social Security benefit payments, with most of these payments going to high-income households that are no longer subject to the cap. Second, while removing the cap will directly expand Social Security's tax base, households for whom the cap expires will also respond to the higher taxes by working and saving less. These responses, if strong enough, can have potentially significant macroeconomic effects that will reduce aggregate labor supply, capital stock, national income, and Social Security's tax base. Therefore, the ability of the tax cap in improving Social Security's fiscal status will depend on the quantitative importance of these effects.

In this paper, I examine if removing the cap on the amount of taxable earnings can solve Social Security's future budgetary problems. Specifically, I compare the change in benefits needed to keep

[1] Defined as the ratio of population above age 65 to that between ages 20-64. Source: the U.S. Census Bureau, http://www.census.gov/population/www/projections/

[2] Source: The 2015 Social Security Trustees Report.

[3] See, among others, studies such as Auerbach and Kotlikoff (1987), Huang et al. (1997), De Nardi et al. (1999), Altig et al. (2001), Nishiyama and Smetters (2005), Conesa and Garriga (2008), and Kitao (2014).

[4] Abolishing the taxable maximum of Social Security was an important part of the 2016 U.S. presidential election campaigns. Senator Bernie Sanders (I-VT), who ran a strong campaign for the Democratic nomination, has introduced specific legislation on this topic in the U.S. Senate. His proposal, titled *The Social Security Expansion Act* (Senate Bill S.731), includes a provision to apply Social Security's payroll tax rate to earnings above $250,000 in 2016 and later, keeping the limit fixed and not indexed to wage growth. The proposal claims that based on projected wage growth in the U.S., the current-law taxable maximum will exceed $250,000 by 2034, after which the payroll tax will apply to all earnings. See Sanders (2015) for further details on Senate Bill S.731.

Social Security solvent under projected future U.S. demographics when this cap is eliminated, to that when the cap is held fixed at its current level. To do this, I construct an overlapping-generations macroeconomic model with incomplete markets, an unfunded public pension system that mimics Social Security, and households that experience realistic mortality and labor income risks. Social Security provides partial insurance against these risks, because households do not have access to private insurance markets. Households in the model also face a progressive labor income tax schedule, factor markets are competitive, firms maximize profit, and the government provides public goods and Social Security. I calibrate this model to match 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 tax cap, and the share of government expenditures in GDP. Finally, I incorporate an empirically reasonable improvement in longevity into the calibrated model, and then compute the change in benefits needed to keep Social Security's budget balanced. In the computations, I allow for all the household-level and macroeconomic adjustments to the longevity improvement, as well as to the changes in Social Security.

Intuitively, the idea of removing the cap on taxes to generate more Social Security revenues seems appealing: the additional distortions caused by this policy are likely to be small, relative to the "across-the-board" policy changes usually considered in the literature, such as increases in the payroll tax rate or cuts in the benefits. However, while removing this cap may come with little additional distortions, this policy change will fundamentally alter the pattern of redistribution implicit in Social Security. In an environment where Social Security partially replaces missing insurance markets, any change in this implicit redistribution will also affect how welfare gains or losses are distributed across households. Therefore, in this paper, I also evaluate the overall welfare effects of removing the cap on Social Security taxes, as well as the distributional consequences of this policy change.

In general, I find that removing the cap on the amount of earnings subject to the payroll tax can partially solve Social Security's future budgetary problems. With the payroll taxes and benefits based on current law, benefits need to decline by almost 12%, on the average, to keep Social Security solvent under the longevity improvement. However, when the cap is removed from the amount of earnings subject to the payroll tax and also creditable towards benefits, I find that the average budget-balancing decline in benefits is only 3.7%. In equilibrium, subjecting all earnings to the payroll tax increases Social Security's revenues by 14%, but counting all earnings towards future benefits causes much of these extra revenues to be spent in paying benefits to wealthy retirees who are no longer subject to the cap. Moreover, households for whom the cap expires respond by working and saving less. This reduces aggregate labor supply by 4.4%, capital stock by 9.3%, and output by 6.2% from the current status-quo under the longevity improvement. Overall, these changes reverse almost one-third of the initial expansion in Social Security's revenues from the elimination of the cap.

I find that the fiscal advantages to Social Security are larger when the cap is removed only from the amount of earnings subject to the payroll tax, but retained on how much of those earnings are creditable towards benefits. In this case, the average budget-balancing change in benefits is actually positive: a 0.4% increase from the baseline. As in the previous case, removal of the cap increases Social Security's revenues by 14% under the longevity improvement, but retaining it on the creditable earnings causes these revenues to be spent overwhelmingly in paying benefits to low- and medium-income retirees. As before, households who are no longer subject to the cap respond by working and saving less, which reduces labor supply, capital stock, and output by 4.4%, 9.2%, and 6.1% respectively, relative to when the cap is held fixed at its current level.

My computations predict that eliminating the cap has a negative effect on overall welfare. Social Security becomes slightly less progressive than the baseline when the cap is removed from the amount of earnings subject to the payroll tax and also creditable towards benefits. This policy changes has negative insurance effects for households with relatively unfavorable earnings histories; their welfare losses are equivalent to 1–1.3% of period consumption. Moreover, this experiment leads to an expansion in the relative size of Social Security, which further distorts labor supply and saving, especially for households no longer subject to the cap. Because of these two effects, this policy change causes a decline in overall welfare, relative to the current status-quo.

I find that the overall welfare loss is smaller in magnitude when the cap is removed only from the amount of earnings subject to Social Security's payroll tax. Retaining the cap on the amount of creditable earnings makes Social Security slightly more progressive than the baseline. This generates stronger insurance effects for households with relatively unfavorable earnings histories. Under this policy change, these households experience smaller welfare losses: equivalent to only 0.6–0.7% of period consumption. But this experiment also causes an expansion in the relative size of Social Security, compared to the current status-quo under the longevity improvement. Accounting for these effects, this policy change also yields an overall welfare loss relative to the current status-quo, but a slight welfare gain relative to when the cap is removed from both taxes and benefits.

Finally, I compare the consequences of removing the tax cap to two other reform policies commonly considered in the literature: separately increasing Social Security's payroll tax rate, and also its full retirement age. First, I find that with Social Security's current taxable maximum and the full retirement age, a payroll tax rate of 13% is needed to keep expected benefits unchanged. This policy has slightly better macroeconomic outcomes than removing the tax cap. In this case, aggregate labor supply, capital stock, and output decline by only 2%, 7.2%, and 2.3% respectively from their corresponding status-quo levels. However, the "across-the-board" tax increase causes a large distortionary welfare loss, equivalent to 2.3–3.5% of period consumption for all households. Second, I find that with Social Security's current payroll tax rate and the taxable maximum, a full retirement age of 69 is needed to keep expected benefits roughly unchanged. The macroeconomic effects of this policy change are even more favorable: declines of 0.6%, 2.1%, and 0.2% in labor supply, capital stock, and output, respectively, from their status-quo levels. Because this policy change keeps the relative size of Social Security roughly unchanged, overall welfare in this case is lower than the current status-quo under the longevity improvement, but slightly higher than when the tax cap is removed.

The literature on Social Security reform in the U.S. has generally concluded that because of increasing life expectancies and falling population growth rates, it will be costly maintaining Social Security benefits at their current level. With benefits being paid out as defined by current law, significant payroll tax increases may be required to balance Social Security's budget in the long run (De Nardi et al., 1999). Kitao (2014) finds that keeping the program self-financed with the current contribution rate will require benefit reductions in the form of reducing the replacement rates by one-third, delaying the normal retirement age by seven years, or letting the benefits decline one-to-one with income. Each of these options, however, will have significant consequences on household consumption, savings, labor force participation, and also labor hours over the life cycle. Conesa and Garriga (2009) show that some of these distortions can be minimized with an age-dependent labor income tax structure, a removal of the compulsory retirement age, and by increasing the level of government debt during the demographic transition. The findings in this paper build on these results and demonstrate that the annual cap on the amount of earnings subject to the Social Security tax can also play an important role in this discussion.

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 define the longevity improvement and also evaluate its consequences on Social Security's current status-quo in Section 5. In Sections 6 and 7, I describe the two main computational experiments of this paper and their quantitative results. I examine two additional computational experiments in Section 8: separately increasing Social Security's payroll tax rate and its full retirement age, and I conclude in Section 9.

2 The model

The unit of the current model is a 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 and mortality risk, but it does not have access to markets to purchase insurance against these risks.

At each date, a surviving household earns labor income if it works, and it also receives 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 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 function is given by

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

where η is the share of consumption, σ is the inverse of the 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, I normalize the period time endowment to unity, i.e. $0\leq h\leq 1$.

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

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

where β is the discount factor, and $Q(s)$ is the unconditional probability of surviving up to age s.

2.2 Income

Conditional on labor force participation, a household earns before-tax wage income $y(s,\varphi)=h(s)we(s,\varphi)$ at age s, where w 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 for Social Security $T_{ss}(\cdot)$ 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%.[5]

2.3 Social Security

The government pays Social Security benefits to households after the full retirement age (T_c), and the amount of benefits paid to a particular household depends on its earnings history. For each household, the government calculates an average of past earnings (up to the maximum taxable earnings), referred to as the Average Indexed Monthly Earnings (AIME). The Social Security benefit, 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.[6]

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

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

[5] 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, 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.

[6] 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.

subject to

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

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

$$0 \leq h(s) \leq 1 \tag{7}$$

$$k' \geq 0 \tag{8}$$

$$AIME' = \begin{cases} \left[AIME \times (s-1) + \min\left\{h(s)we(s, \varphi), \bar{y}\right\}\right]/s & s < T_c \\ AIME & s \geq T_c \end{cases}, \tag{9}$$

where

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

is a step function. Households are born with and die with zero assets, i.e. $k(0) = k(T+1) = 0$, and prior to age T_c, 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

$$Y \;=\; K^{\alpha} L^{1-\alpha}, \tag{10}$$

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

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

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

where δ is the depreciation rate of physical capital.

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 are given by

$$K = \sum_{s=0}^{T} N(s)Q(s) \int k(s+1; x)\, d\mu_s(x) \tag{13}$$

$$L = \sum_{s=0}^{T} N(s)Q(s) \int h(s; x)e(s, x)\, d\mu_s(x), \tag{14}$$

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

The total value of the accidental bequests from households who die on a given date is

$$BEQ = (1+r)\left[\sum_{s=0}^{T}\{N(s)Q(s) - N(s+1)Q(s+1)\}\int k(s+1;x)d\mu_s(x)\right]$$
$$-n\sum_{s=0}^{T} N(s)Q(s)\int k(s+1;x)d\mu_s(x). \tag{15}$$

The first term in equation (15) gives the total capital held by all cohorts transitioning from a given age to the next, and the second term adjusts this amount to whatever is attributable to mortality risk. This is because cohort size changes in the current model for two reasons: mortality risk, as well as population growth.

The budget-balancing condition for Social Security is given by

$$\sum_{s=0}^{T} N(s)Q(s)\int T_{ss}\left(h(s;x)we(s,x);\bar{y}\right)d\mu_s(x) = \sum_{s=0}^{T} N(s)Q(s)\int \Theta(s-T_c)b(x)d\mu_s(x). \tag{16}$$

Finally, the government adjusts its level of non-Social Security expenditures (G) to match the total tax revenues from labor income, capital income, and the accidental bequests

$$BEQ + \tau_k r\sum_{s=0}^{T} N(s)Q(s)\int k(s;x)d\mu_s(x) + \sum_{s=0}^{T} N(s)Q(s)\int T_y\left(y(s;x)\right)d\mu_s(x) = G. \tag{17}$$

2.7 Competitive equilibrium

A competitive equilibrium in this model 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 and labor L,
3. a rate of return r and a wage rate w,
4. Social Security benefits $b(x)$ and government expenditures G, and
5. a measure of households $\mu_s(x)\ \forall\ s$,

that

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

In equilibrium, total expenditure equals consumption plus net investment plus government purchases, which is equal to the total income earned from capital and labor.

$$\begin{aligned} C + K' - (1-\delta)K + G &= C + (n+\delta)K + G \\ &= wL + (r+\delta)K \\ &= Y. \end{aligned} \tag{18}$$

Finally, I consider only a steady-state equilibrium, i.e. when all aggregate quantities grow at the rate of population growth and all per-capita quantities are constant. I also normalize the initial newborn cohort size to $N(0) = 1$.

3 Calibration

3.1 Demographics

I first set the demographic parameters of the model. I assume that households enter the model at the actual age of 25, which corresponds to the model age of zero. I obtain the average age-specific death rates in the model from the 2001 U.S. Life Tables in Arias (2004), and because these data are reported up to the actual age of 100, I set the maximum model age to $T = 75$. Finally, I set the population growth rate to $n = 1\%$, which is consistent with the U.S. demographic history and also with the literature.

3.2 Social Security

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

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

and then set the tax rate to $\tau_{ss} = 0.106$, which is the combined tax rate for the Old-Age and Survivors Insurance (OASI) component.[7] The maximum taxable earnings ($\bar{y}$) is adjusted regularly 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. I use this estimate to 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

[7] In reality, this rate is evenly split between the employer and the employee, but the standard hypothesis in the literature is that due to labor-market pressures, the employee bears the full burden of the tax.

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

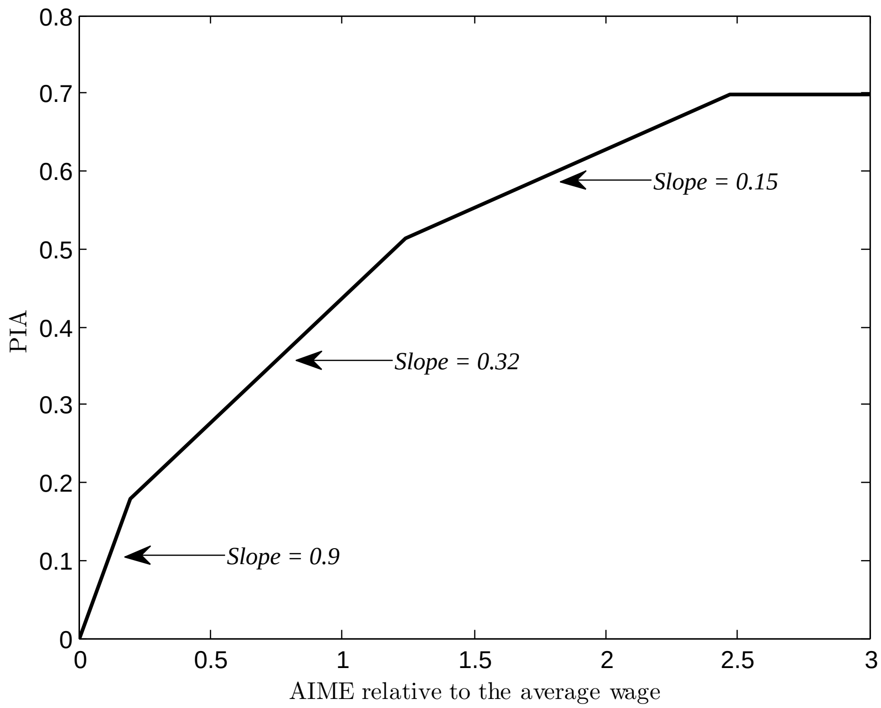

Figure 1: Benefit formula in the U.S.

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-earnings rule is captured by the fact that the "replacement rate" is decreasing in the AIME (see Figure 1).

Finally, I assume that benefit collection in the model begins at age $T_c = 41$, which corresponds to the current full retirement age of 66 in the U.S.

3.3 Labor productivity

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

$$\log e(s, \varphi) = \epsilon(s) + \varphi, \tag{19}$$

where $\epsilon(s)$ is a deterministic age-dependent component, and φ is a stochastic component, which is further decomposed as

$$\varphi_t = p + z_t + \nu_t. \tag{20}$$

Here, $p \sim N(0, \sigma_p^2)$ is a permanent productivity fixed effect, $\nu_t \sim N(0, \sigma_\nu^2)$ is a transitory shock, and a z_t is a persistent shock that follows a first-order autoregressive process

$$z_t = \rho z_{t-1} + \upsilon_t \tag{21}$$

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

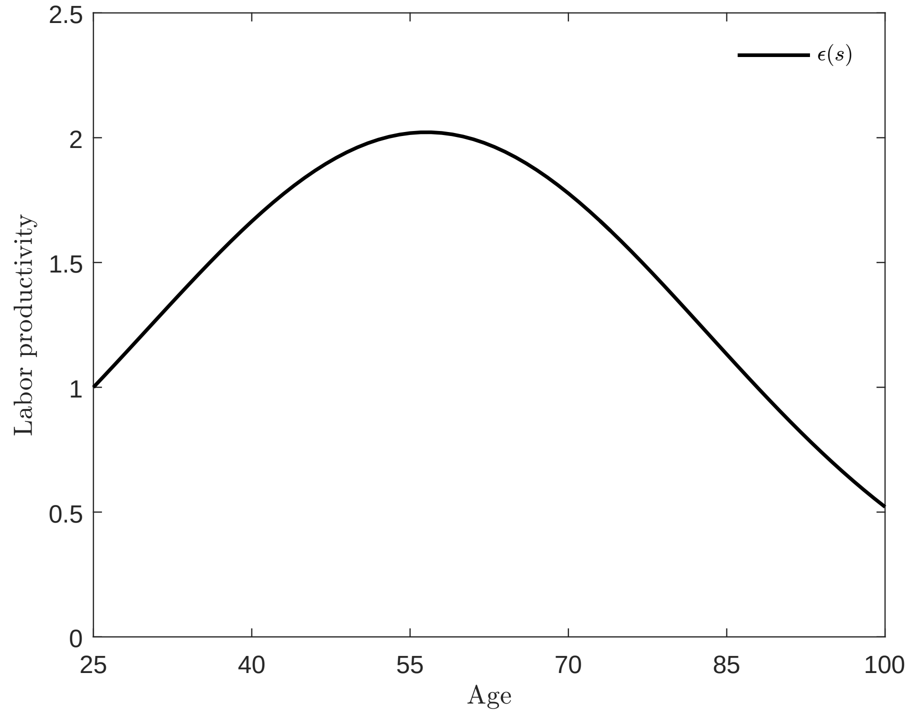

Figure 2: The age-dependent component of labor productivity from Kitao (2014).

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 2.

I use estimates from Heathcote et al. (2010) to calibrate the stochastic component. I set the persistence parameter to $\rho = 0.973$, the variances of the permanent fixed effect and the transitory shock to $\sigma_p^2 = 0.124$ and $\sigma_\nu^2 = 0.04$ respectively, and variance of the white-noise disturbance to $\sigma_v^2 = 0.015$. I use Gaussian quadrature to approximate the distribution of the permanent fixed effect using a three-point discrete distribution, and I approximate the joint distribution of the persistent and the transitory shocks using a five-state first-order discrete Markov process following Tauchen and Hussey (1991).

3.4 Taxes

To calibrate the labor income tax function, I follow Karabarbounis (2012) and Heathcote et al. (2014) and assume that

$$T_y(y) = y - (1 - \tau_y)y^{1-\tau_1}, \tag{22}$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (22) 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, and the parameter τ_1 controls the progressivity of the tax code. Following Heathcote et al. (2014), I set the value of this parameter to $\tau_1 = 0.151$. Heathcote et al. (2014) estimate this value using data from the 2000, 2002, 2004, and 2006 waves of the PSID, and also NBER's TAXSIM program, accounting for federal and state income taxes plus public transfers. I plot the average and marginal tax rates that emerge from the estimated tax function in Figure 3.

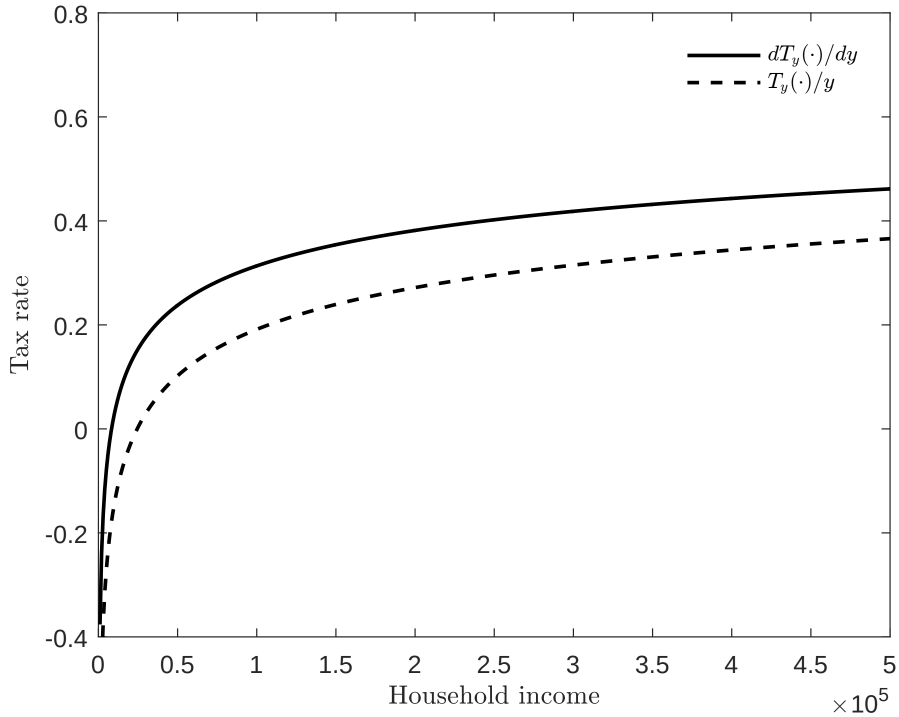

Figure 3: The average and marginal labor income tax rates from Heathcote et al. (2014).

Finally, capital income tax in the U.S. consists of taxes on interest income, and also on capital gains. However, there are no capital gains in current model because it has only one asset and there is no aggregate uncertainty. Therefore, to parameterize the capital income tax rate, I follow the literature and set $\tau_k = 30\%$ (De Nardi et al., 1999; İmrohoroğlu and Kitao, 2012; Kitao, 2014).

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 $\delta = 0.06$.

3.6 Structural parameters

Once all the observable parameters have been assigned empirically reasonable values, I jointly calibrate the remaining unobservable structural parameters of the model, i.e. the preference parameters σ, β, and η, the age-dependent labor force participation cost θ_P, and also the labor income tax parameter τ_y, to match key 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. I adopt both of these empirical facts as targets.[9]

[9] The labor force participation and the hours per week per worker targets are based on PSID data as noted in Kitao (2014).

σ	β	η	κ_1	κ_2	κ_3	τ_y
4	0.9826	0.438	6.12×10^{-8}	3.43×10^{-7}	2.98	0.103

Table 1: Structural parameter values under the baseline calibration.

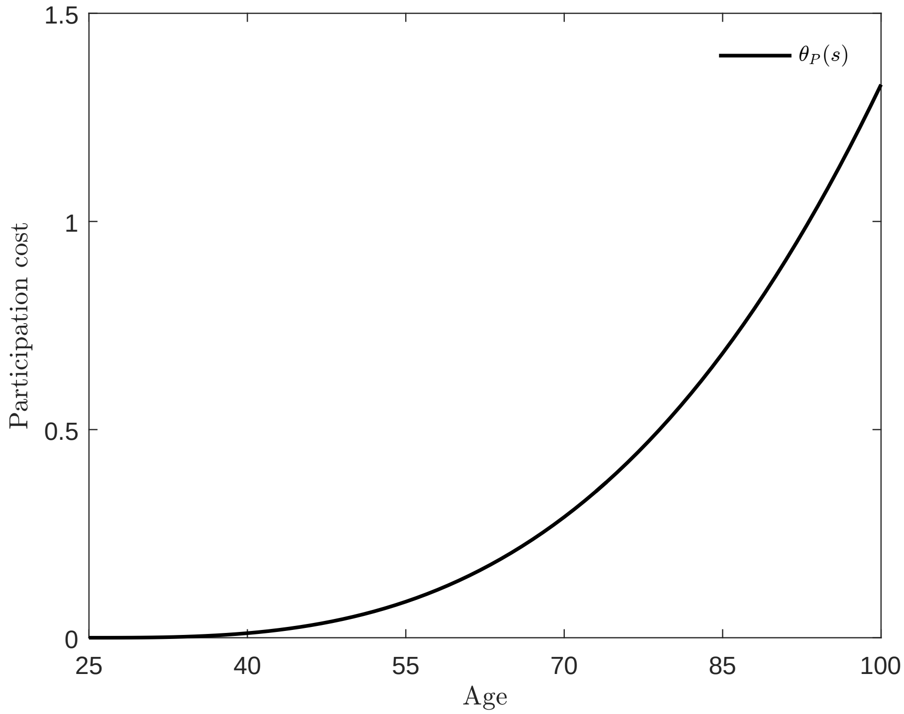

Figure 4: Age-dependent labor force participation cost $\theta_P(s)$.

Following Kitao (2014), I assume that the labor force participation cost increases with age as

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

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 on participation), so I calibrate it to match the hours per week target.

Finally, I calibrate τ_y such that the model yields a ratio of government expenditures to GDP of around 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.

4 Baseline results

The structural 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 4). The model-generated values for the targets under

	Target	Model
Capital-output ratio	3.0	3.03
Avg. hours of market work per week per worker (25-55)	40	38.6
Share of govt. expenditures in GDP	0.2	0.24
Social Security's tax base as a fraction of total earnings	–	0.82

Table 2: Model performance under the baseline calibration.

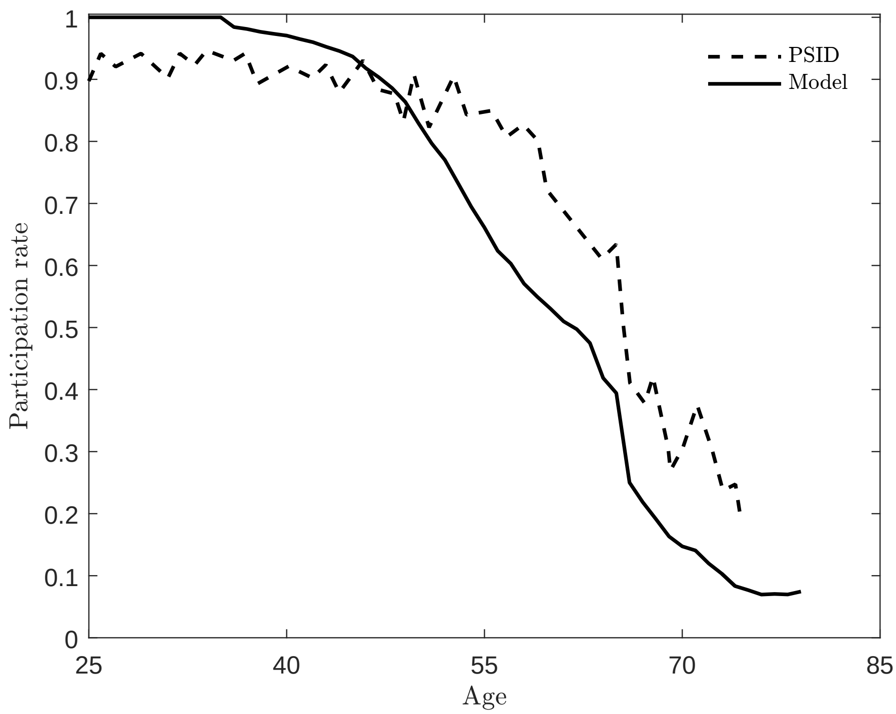

Figure 5: Cross-sectional labor force participation rates under the baseline calibration.

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 5 and 6.

It is clear from Figures 5 and 6 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 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 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

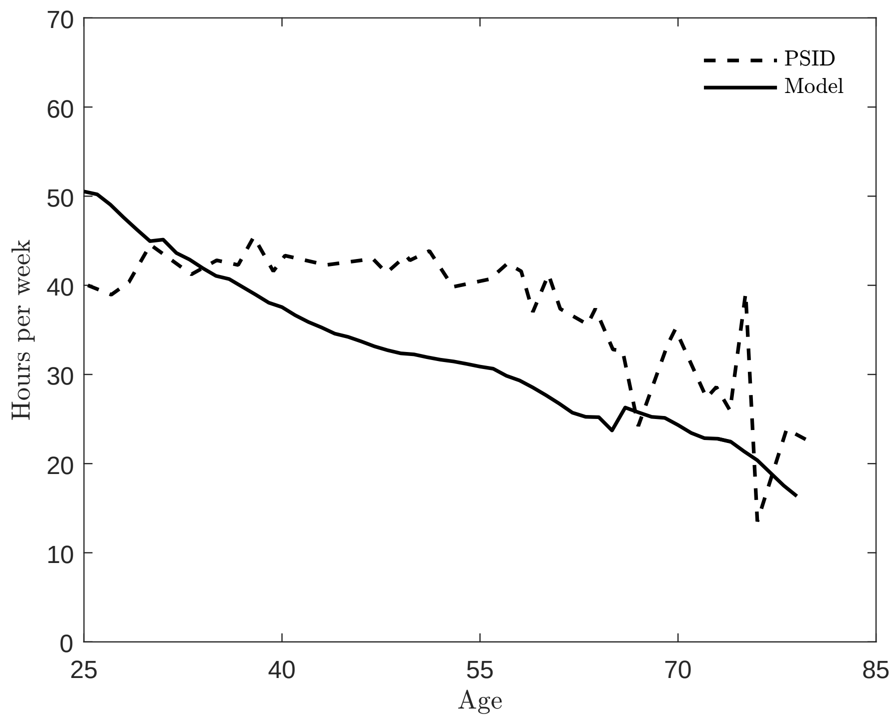

Figure 6: Cross-sectional mean of labor hours per week (conditional on participation) under the baseline calibration.

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 risks, explains why the current model underestimates asset holdings at later ages, as seen in Figure 7 (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. 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 8, I report the fraction of workers 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, as mentioned earlier, 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 mortality risks. However, the equilibrium interest rate in the baseline calibration is 5.6%, which is considerably higher than the population growth rate of 1%. Therefore, as an

[10]Social Security's Annual Statistical Supplement, Table 4B.1.

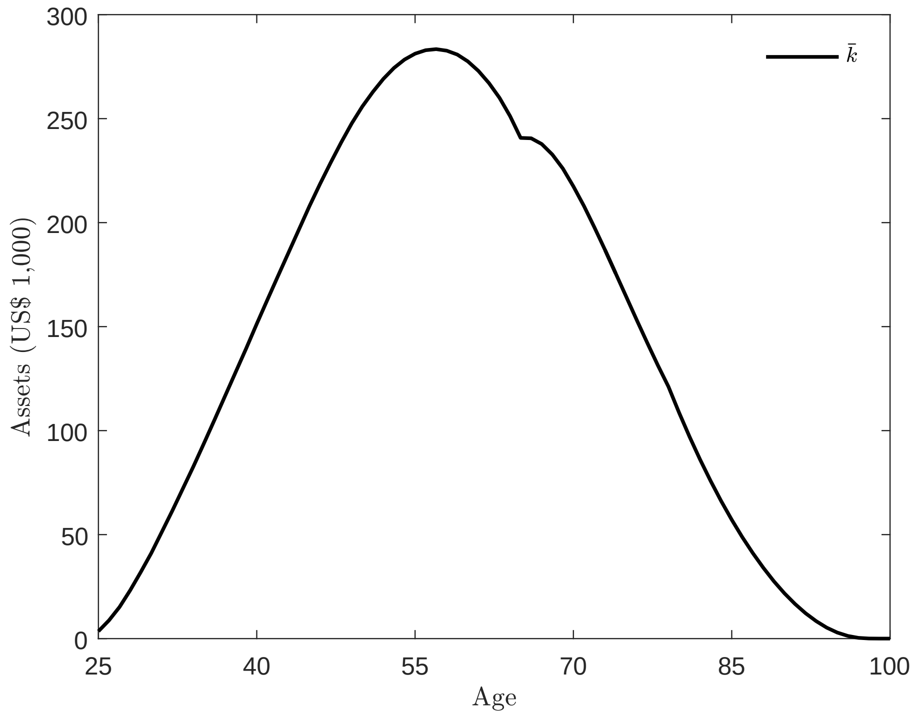

Figure 7: Cross-sectional mean of asset holdings under the baseline calibration.

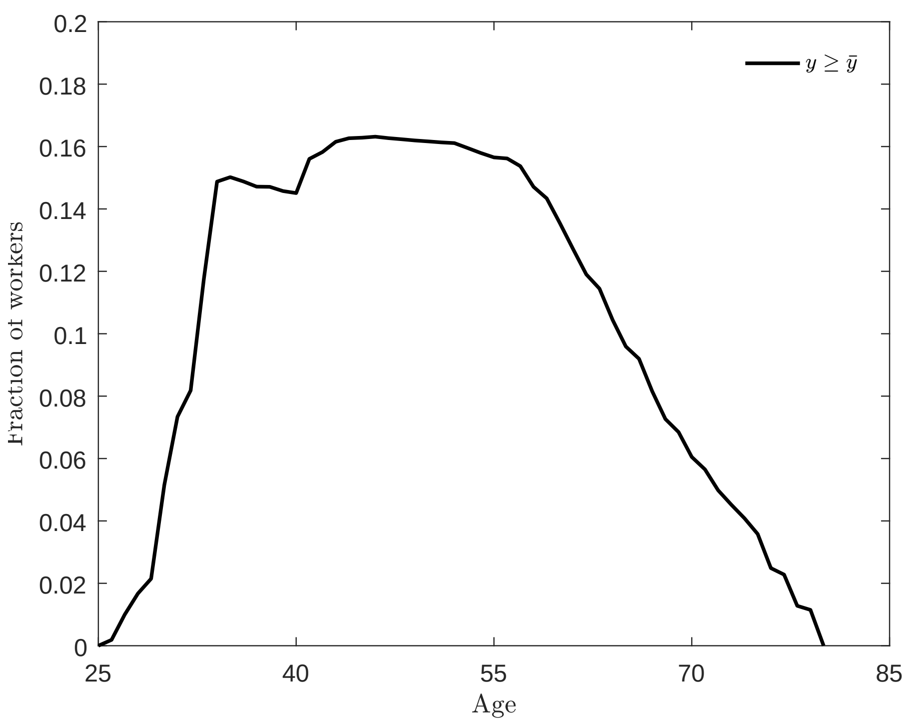

Figure 8: Fraction of workers subject to the tax cap in the baseline calibration.

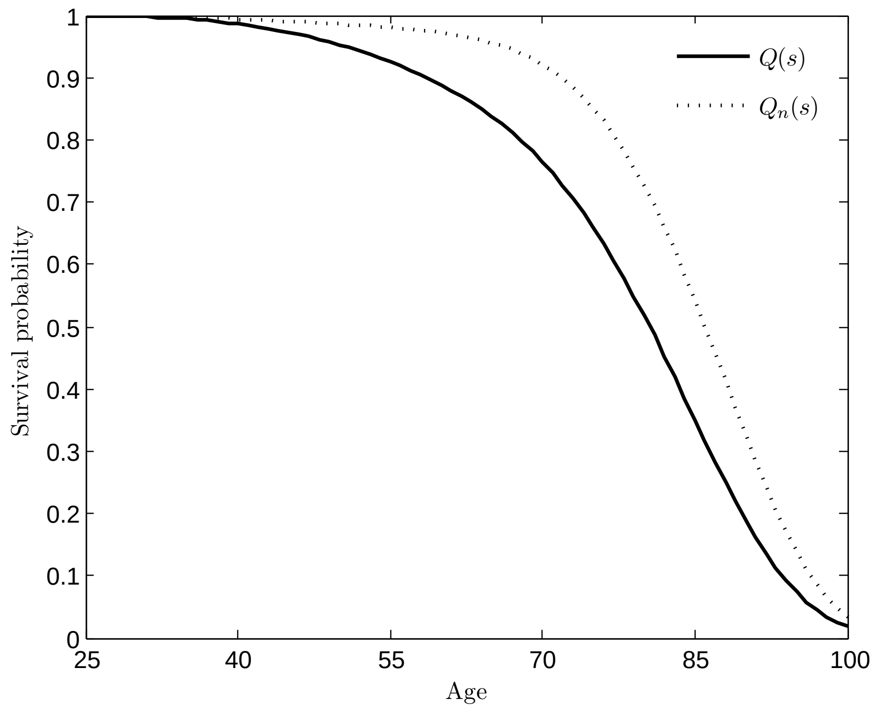

Figure 9: Baseline (Q) and the projected (Q_n) survival probabilities.

intergenerational transfer, Social Security does not have a welfare-improving role in the baseline calibration.

5 The longevity improvement

Based on the historical life tables, 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). A straightforward way to incorporate such a longevity improvement into the baseline model is to reduce the baseline age-specific death rates $d(s)$ based on the following formula:

$$d_n(s) = d(s) - \lambda s^{\gamma}, \tag{23}$$

where λ and γ are positive constants. I set these parameters to $\lambda = 10^{-6}$ and $\gamma = 1.8509$, under which the life expectancy in the model is 85 years, matching the Social Security Trustees Report's average period life expectancy projection for the year 2085 under the intermediate assumption. I compare the survivor function resulting from this longevity improvement to the baseline in Figure 9.

It is well known that under these projected demographics, Social Security in its current form is insolvent in the long run. However, several studies contend that the SSA's actuarial projections overestimate the magnitude of the insolvency, and also underestimate the tax increases that will be needed to balance Social Security's budget. Bagchi (2016) argues that the SSA's actuarial projections overestimate the Social Security crisis because they do not account for how key household-level and macroeconomic variables will respond to the longevity improvements. Bagchi (2016) shows that once these responses are accounted for, the budget-balancing decline in Social Security benefits is

	Baseline	Case 0
Labor	41.9	45.9
Avg. hours of market work per week per worker (25-55)	38.6	39.5
Capital	230.4	284.1
Output	76.1	86.8
Interest rate	5.53%	4.71%
Wage	1.18	1.23

Table 3: Select macroeconomic variables under Case 0, compared to that under the baseline calibration.

Productivity fixed effect (p)	0.54	1.00	1.84	Average
% change	−1.4	−10.1	−30.7	−12.1

Table 4: Change in expected Social Security benefits from the baseline under Case 0.

only two-thirds of the SSA's estimates. On the other hand, De Nardi et al. (1999) find that balancing Social Security's budget will require an additional 17.1 percentage point tax increase on the top of the SSA's actuarial projections. This is because the SSA's projections do not account for the negative impact that these higher taxes will have on the tax base. De Nardi et al. (1999) argue that households will respond to the higher Social Security taxes by working and saving less, which will shrink the tax base from its current level. Because the current model is an equilibrium model of the economy, it accounts for all the relevant household-level and macroeconomic adjustments to changes in longevity as well as Social Security.

To determine if removing the tax cap can be a viable policy tool in solving Social Security's future budgetary problems, I first compute a benchmark scenario in which all the institutional features of Social Security, including the cap, are held at their status-quo (baseline) levels. Only the survival probabilities under this experiment are changed to Q_n to reflect the longevity improvement. From this experiment, I calculate the change in benefits needed to balance Social Security's budget in an environment where all future benefit payments are based on current law (Case 0).

The macroeconomic results for Case 0 are reported in Table 3. The table shows that the longevity improvement increases both labor supply and saving. Households respond by working more, both in terms of labor force participation and the hours per week (see Figures 10 and 11), and they also save more to smooth consumption over a longer expected lifespan. This increases labor supply by 9.4%, capital stock by 23%, and GDP by 14% from the initial baseline.[11]

The change in benefits needed to balance Social Security's budget under this experiment are reported in Table 4. The table shows that on the average, benefits decline by about 12% from their baseline level, which is significantly smaller that the actuarial estimates of the SSA. This should not be surprising, as the current model accounts for all the household-level and macroeconomic adjustments to the longevity improvement: households respond by supplying more labor and also by saving more, which leads to a natural expansion in Social Security's tax base (Bagchi, 2016).

[11]It is important to note here that the longevity improvement actually affects household behavior through two different channels. First, through pure life-cycle motives, a higher life expectancy gives households the incentive to supply more labor and save more. Second, a higher life expectancy also reduces Social Security benefits in the equilibrium, which has a negative income effect and encourages labor supply and saving even more. The above results are the combined effect of these two channels.

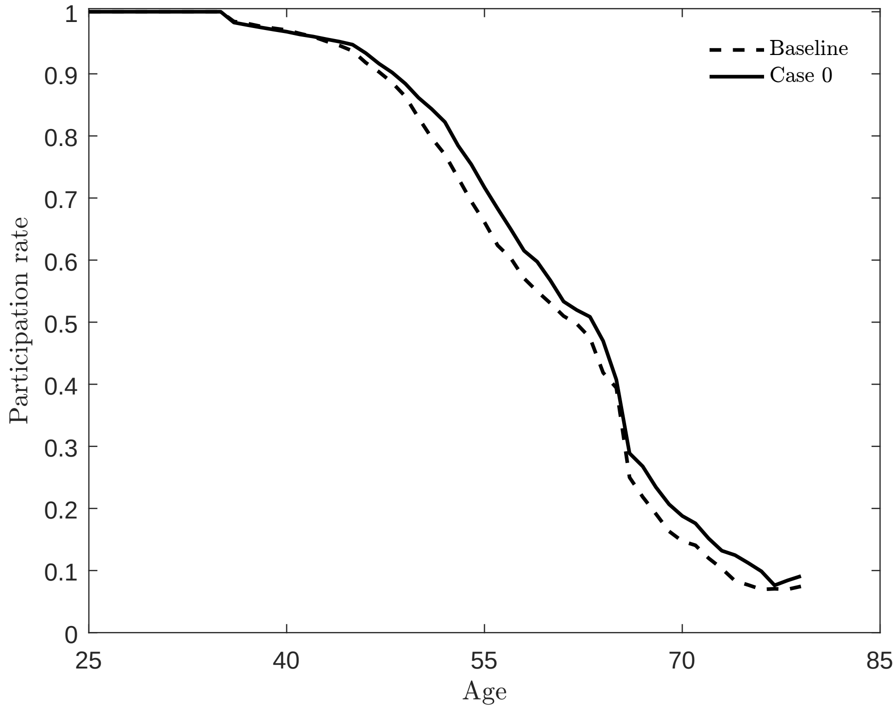

Figure 10: Cross-sectional labor force participation rates under Case 0, along with the baseline calibration.

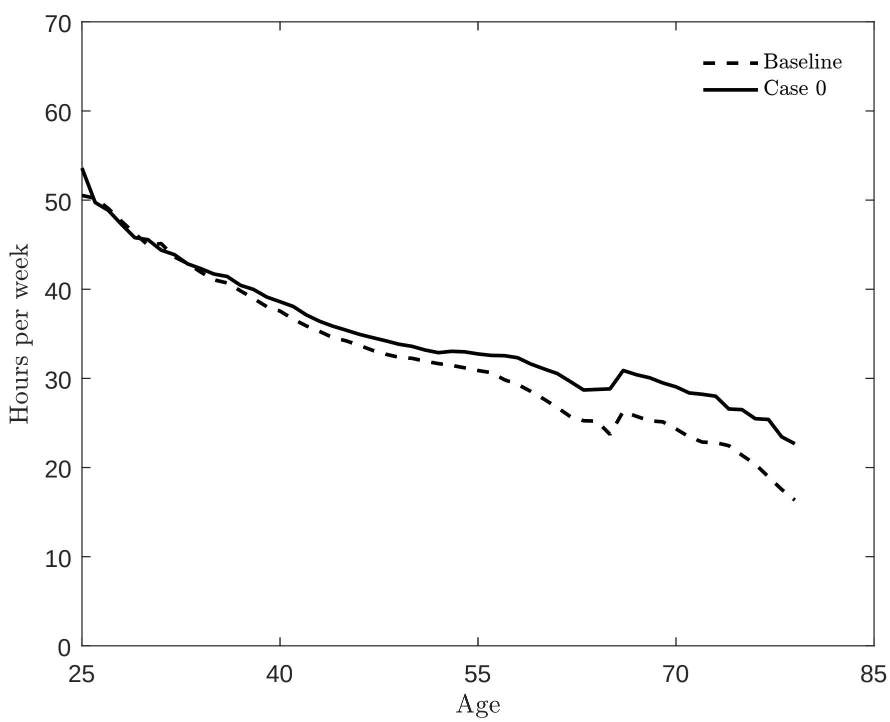

Figure 11: Cross-sectional mean of labor hours per week (conditional on participation) under Case 0, along with the baseline calibration.

6 Two experiments with the cap

As explained earlier, Social Security benefits in the U.S. are calculated as a function of average work-life income, and earnings only up to the cap are counted towards the benefits. Therefore, the maximum taxable earnings also sets a de-facto limit on the amount of benefit payments from Social Security. Given this fact, there are two possible ways in which the tax cap can be used as a policy tool in improving Social Security's fiscal situation. The first policy option is to subject all earnings to the Social Security tax, and to allow all earnings to be counted towards future Social Security benefits. This option would retain the historical link between the cap on taxes and benefits paid out by Social Security in the U.S.

The second policy option is to remove the cap only from the amount of earnings subject to the Social Security tax, but to retain it on the amount of earnings that can be counted towards future Social Security benefits. This policy change would expand Social Security revenues, but retaining the cap on creditable earnings would limit the amount of benefit payments. However, this option would break the historical link between the cap on taxes and benefits in the U.S. social security program.[12]

To examine the consequences of these two policy changes, I define the following two experiments. In the first experiment, I compute a new equilibrium of the model with the improved survival probabilities, while subjecting all earnings to the Social Security tax rate of $\tau_{ss} = 0.106$, and also counting all earnings toward benefits in the Social Security benefit-earnings formula (Case 1). In the second experiment, I subject all earnings to the Social Security tax rate of $\tau_{ss} = 0.106$, but only count earnings up to the current cap of 2.47 times the average earnings, in the benefit-earnings formula (Case 2). In both the computations, I account for all the household-level and macroeconomic adjustments to the longevity improvement, as well as to the policy changes.

While the change in benefits needed to balance Social Security's budget is a sufficient metric for the fiscal consequences of these policy changes, I define the following two metrics to measure the welfare consequences of these two experiments. First, to understand the overall welfare consequences, I define

$$W = \sum_{s=0}^{T} \beta^s Q_n(s) \int u\left(c(s;x), 1-h(s;x), P(s;x)\right) d\mu_s(x) \tag{24}$$

which is the ex-ante expected lifetime utility, given the longevity improvement. Second, to understand the distributional consequences of these policy changes, I define a consumption equivalence ψ for each realization of the permanent productivity fixed effect (p) that solves

$$E\left[\sum_{s=0}^{T} \beta^s Q_n(s)\, u\left((1+\psi)\, c^C(s), 1-h^C(s), P^C(s)\right)\right] = E\left[\sum_{s=0}^{T} \beta^s Q_n(s)\, u\left(c^{NC}(s), 1-h^{NC}(s), P^{NC}(s)\right)\right], \tag{25}$$

where C denotes current Social Security law, and NC denotes a hypothetical Social Security law without the cap. Intuitively, this consumption equivalence captures the welfare gains (or losses) in units of consumption, as a function of the productivity fixed effect, under each of the two experiments. Taken together, these two metrics provide an overall, as well as a disaggregated picture, of the welfare consequences of removing the cap.

[12]This is the provision included in senate bill S. 731, *The Social Security Expansion Act*, introduced by Senator Bernie Sanders.

	Case 0	Case 1
Labor	45.9	43.8
Avg. hours of market work per week per worker (25-55)	39.5	38.8
Capital	284.1	257.5
GDP	86.8	81.4
Ratio of Social Security expenditures to GDP	5.7%	6.9%
Interest rate	4.71%	5.03%
Wage	1.23	1.21

Table 5: Select macroeconomic variables under Case 1, compared to those under Case 0.

Productivity fixed effect (p)	0.54	1.00	1.84	Average
Case 0 (% change)	-1.4	-10.1	-30.7	-12.1
Case 1 (% change)	-9.6	-6.9	14.9	-3.7

Table 6: Change in expected Social Security benefits from the baseline under Case 1, and also under Case 0.

7 Results from the experiments

7.1 Removing the cap both from taxes and benefits

I first consider the policy option in which all earnings are subject to Social Security's payroll tax and also counted towards benefits in the Social Security benefit-earnings formula. As discussed earlier, this policy preserves the historical link between the caps on taxes and benefits in U.S. Social Security.

I report the macroeconomic consequences of this experiment in Table 5. The table shows that removing the cap both from the amount of earnings subject to the Social Security tax, as well as the amount of earnings counted towards benefits, has a negative impact on aggregate labor supply, capital stock, and GDP, relative to the benchmark case when the cap is held fixed at the baseline level (Case 0). This is because the marginal tax rates under this experiment are higher for households formerly subject to the cap, giving them an incentive to retire earlier, reduce their weekly hours, and also save less. Overall labor supply increases by only 4.6% from the baseline level under Case 1, compared to by 9.4% under Case 0. Similarly, capital increases by only around 12% from the baseline under Case 1, compared to by more than 23% under Case 0. Relative to the current status-quo under the longevity improvement (Case 0), this implies a 4.4% reduction in labor supply, an 9.3% reduction in capital stock, and a 6.2% reduction in GDP. Finally, Social Security expenditures increase from 5.7% of GDP under the current status-quo, to almost 7% of GDP under this experiment.

I compare in Table 6 the change in benefits needed to balance Social Security's budget under Case 1, to that under Case 0. It is clear from the table that on the average, the decline in Social Security benefits under Case 1 is considerably smaller: only 3.7%, compared to around 12% under Case 0. Moreover, while Social Security benefits decline for *all* households under Case 0, expected benefits for households with $p = 1.84$ increase by roughly 15% from the baseline under Case 1.

To understand why this experiment leads to a smaller decline in benefits on the average, but actually higher benefits for households with $p = 1.84$, first consider the effect of removing the tax cap on Social Security's overall revenues. With the factor prices and labor supply fixed at their

Productivity fixed effect (p)	0.54	1.00	1.84
Case 0	0.438 ($23,430)	0.345 ($38,678)	0.266 ($50,011)
Case 1	0.408 ($24,609)	0.331 ($44,519)	0.300 ($78,643)

Table 7: Average replacement rates and AIME values under Case 1, and also under Case 0.

levels under Case 0, eliminating the cap on taxes would increase Social Security's revenues by about 22%. However, the current model accounts for all the household-level and macroeconomic adjustments to this policy change, and as Table 5 shows, overall labor supply is negatively affected in this case. Because Social Security's tax base depends on overall labor supply, Social Security's revenues increase only by about 14% when this negative effect is accounted for.

Second, recall that because Case 1 preserves the historical link between the cap on earnings subject to the Social Security tax and also on the earnings creditable toward Social Security benefits, eliminating the cap on taxes in this experiment also implies eliminating the cap on Social Security benefits. Therefore, much of the extra revenues generated in this experiment are actually spent in paying benefits to households with relatively favorable earnings histories that are no longer subject to the cap. As seen in Table 6, benefits for households with $p = 1.84$ actually *increase* from the baseline by roughly 15% under Case 1, compared to declining by about 30% under Case 0, and benefits for households with $p = 0.54$ decline by almost 10% under Case 1, compared to only a modest decline of 1.4% under Case 0.

Intuitively, the effect of Case 1 on the redistribution implicit in Social Security is less clear. Social Security's current tax structure is regressive because due to the cap, the average tax rate is a decreasing function of earnings. Eliminating the cap on taxes, therefore, makes Social Security's tax structure strictly proportional and more progressive than the baseline. However, on the benefit side, Social Security's current benefit structure is progressive. This is because the amount of earnings creditable toward benefits is capped, and the replacement rate is a decreasing function of past earnings. Eliminating the cap from the creditable earnings, therefore, makes Social Security's benefit structure less progressive than the baseline. To understand the net redistributive effect of this policy change, I report in Table 7 the average replacement rate for each value of the productivity fixed effect (p) under Cases 0 and 1. I calculate this average replacement rate as the ratio of expected Social Security benefits to the AIME, and I also report the corresponding average AIME values in parentheses.

It is clear from the table that Social Security remains progressive, in an overall sense, under both cases: the average replacement rates are a decreasing function of past earnings (the AIME). However, it is also clear that compared to Case 1, replacement rates decline for households with relatively unfavorable earnings histories, and increase for those with relatively favorable earnings histories. Therefore, even though the U.S. benefit-earnings rule is quite flat in this range (slope of 15%), counting all earnings towards benefits in the current model has a larger negative effect on the implicit redistribution in Social Security. As a result, this policy change makes Social Security less progressive relative to that under Case 0.

There are two mechanisms through which Case 1 affects household-level and overall welfare. First, given that Social Security in the current model provides partial insurance against labor income and mortality risks, the reduction in the implicit redistribution observed in Table 7 causes welfare losses to households with relatively unfavorable earnings histories. Second, removing the cap from the amount of earnings subject to the payroll tax, increases the average Social Security tax rates for households with higher labor income. These higher tax rates impose larger distortions on the labor supply and saving decisions of these households. Accounting for all of these mechanisms,

	W
Case 0	-65.89
Case 1	-66.73

Table 8: Ex-ante expected utilities under Cases 0 and 1.

Productivity fixed effect (p)	0.54	1.00	1.84
%	-1.27	-0.98	0.30

Table 9: The consumption equivalence (ψ) under Case 1.

I find the net effect of Case 1 on overall welfare to be negative (see Table 8).

The fact that an expansion of Social Security reduces overall welfare, in general, is not surprising. It is well known that in a general-equilibrium model with endogenous labor, the distortionary effects of a Social Security expansion on labor supply and saving are often large enough to outweigh the welfare gains from its insurance effects (Nishiyama and Smetters, 2008; Bagchi, 2015). Moreover, as seen from Table 7, eliminating the cap simultaneously from Social Security's taxes and benefits also reduces the implicit redistribution in Social Security, thereby further weakening its insurance effects. I report in Table 9 the consumption equivalence (ψ) for each value of the productivity fixed effect in this case. As expected, the table shows welfare losses for households with relatively unfavorable earnings histories (i.e. $p = 0.54$ and 1.0), and a slight welfare gain for households with relatively favorable earnings histories ($p = 1.84$) under Case 1. This is due to Social Security's lower implicit redistribution under this experiment. However, the magnitude of the welfare losses are significantly larger in this case, generating the decline in overall welfare identified in Table 8.

To summarize, I find that subjecting all earnings to the payroll tax, and also counting them towards benefits, has a positive impact on Social Security's fiscal status, but a negative impact on labor supply, capital stock, and output. With this policy change, benefits need to decline by less than 4%, on the average, to keep Social Security solvent under the longevity improvement. This is significantly smaller than the 12% decline needed when the taxes and benefits continue to be based on current law. This policy change also reduces the redistribution implicit in Social Security, which negatively affects the welfare of households with relatively unfavorable earnings histories. Finally, the expansion of Social Security under this experiment causes distortionary welfare losses that are large enough to yield a decline in overall welfare.

7.2 Removing the cap only from taxes

I now consider the second policy option in which all earnings are subject to the Social Security tax, but earnings only up to the level of the current cap are counted towards Social Security benefits. Note that this experiment breaks the historical link between the cap on Social Security taxes and benefits in the U.S. I report the macroeconomic consequences of this experiment in Table 10. It is clear from the table that the macroeconomic effects of Case 2 are virtually indistinguishable from those of Case 1: labor supply, capital stock, and GDP, all decline compared to the current status-quo, and the relative size of Social Security increases. Therefore, eliminating the cap from the amount of earnings subject to Social Security's payroll tax, but retaining it on the amount of earnings creditable towards benefits, has the same effect on macroeconomic aggregates as does eliminating the cap altogether.

However, the effect of this policy change on Social Security benefits is quite different from that under Case 1. In Table 11, I report the change in benefits needed to balance Social Security's

	Case 0	Case 1	Case 2
Labor	45.9	43.8	43.8
Avg. hours of market work per week per worker (25-55)	39.5	38.8	38.8
Capital	284.1	257.5	257.9
GDP	81.4	81.5	
Ratio of Social Security expenditures to GDP	5.7%	6.9%	6.9%
Interest rate	4.71%	5.03%	5.07%
Wage	1.23	1.21	1.21

Table 10: Select macroeconomic variables under Case 2, compared to those under Cases 0 and 1.

Productivity fixed effect (p)	0.54	1.00	1.84	Average
Case 0 (% change)	−1.4	−10.1	−30.7	−12.1
Case 1 (% change)	−9.6	−6.9	14.9	−3.7
Case 2 (% change)	12.6	2.6	−20.9	0.4

Table 11: Change in expected Social Security benefits from the baseline under Case 1 and 2, and also under Case 0.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 0	0.438 ($23,430)	0.345 ($38,678)	0.266 ($50,011)
Case 1	0.408 ($24,609)	0.331 ($44,519)	0.300 ($78,643)
Case 2	0.516 ($24,093)	0.409 ($39,466)	0.304 ($52,388)

Table 12: Average replacement rates and AIME values under Cases 1 and 2, compared to that under Case 0.

budget under Case 2, along with those under Cases 0 and 1, for each value of the productivity fixed effect (p). Two facts are clear from the table. First, the average budget-balancing change in Social Security benefits in this case is only 0.4%, which actually indicates a marginal increase from the baseline. This is in spite of the fact that under this experiment, Social Security's revenues increase by the exact same percentage as under Case 1: 14% relative to the current status-quo. Second, unlike Case 1, benefits decline only for households with relatively favorable earnings histories under Case 2. Households with relatively unfavorable earnings histories see their benefits increase under this experiment. This is because earnings only up to the cap are now counted towards benefits, and the extra revenues collected from applying the payroll tax to all earnings is now spent overwhelmingly in paying benefits to low- and medium-income retirees, rather than to those for whom the cap expires.

The effect of this policy change on Social Security's implicit redistribution is also different from that under Case 1. As before, applying the payroll tax to all earnings makes Social Security's tax structure strictly proportional, which is more progressive than the current tax structure where the tax applies to earnings only up to the cap. However, under this experiment, the cap is retained on the amount of earnings creditable towards benefits, which keeps the progressivity of Social Security's current benefit structure unchanged. Therefore, one would expect this policy change to increase the redistribution implicit in Social Security. To verify this, I report the the average replacement rates under this experiment (along with their corresponding AIME values) in Table 12. The table shows a clear increase in Social Security's implicit redistribution. The average replacement rates under this experiment are considerably higher for households with relatively unfavorable earnings

	W
Case 0	−65.89
Case 1	−66.73
Case 2	−66.49

Table 13: Ex-ante expected utilities under Cases 0, 1, and 2.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 1 (%)	−1.27	−0.98	0.30
Case 2 (%)	−0.56	−0.68	−1.36

Table 14: The consumption equivalence (ψ) under Case 2, and also under Case 1.

histories (i.e. with a lower AIME value), but almost unchanged for households with favorable earnings histories.

Because this experiment makes Social Security more progressive relative to what is implied by its current tax and benefit structure, the welfare consequences of Case 2 are also slightly different. On the one hand, removing the cap on earnings subject to the tax while retaining it on the earnings creditable towards benefits makes Social Security more progressive and has positive insurance effects. On the other hand, the higher marginal tax rates impose larger distortions on the labor supply and saving decisions of the households no longer subject to the cap. Accounting for both of these effects, I find that while the overall welfare implications of Case 2 are slightly better than Case 1, utility is still lower than Case 0, i.e. the current status-quo (see Table 13). Evidently, the welfare gains associated with the insurance effects of Social Security are larger in this case, but not large enough to compensate for the distortionary losses due to the expansion of Social Security.

Finally, I report in Table 14 the distributional consequences of Case 2 as the consumption equivalence (ψ) for each realization of the productivity fixed effect (p). The table shows smaller welfare losses for households with $p = 0.54$ and 1.00 under Case 2. This is due to the fact that Social Security is more progressive under this experiment. However, as expected, households with $p = 1.84$ suffer a larger welfare loss under this experiment because of the increased redistribution, and also because their marginal tax rates increase. Overall, the welfare results remain consistent with the general fact that in a general-equilibrium economy with endogenous labor, Social Security's insurance effects are not strong enough to offset its distortionary welfare losses.

To summarize, my computations suggest that while the macroeconomic effects of Case 2 are roughly comparable to those under Case 1, the effect on Social Security benefits is quite different. While Social Security's revenues increase by the same percentage under this experiment, removing the cap only from the amount of earnings subject to the payroll tax actually warrants a 0.4% increase in benefits, on the average, to balance Social Security's budget. This is compared to a decline of 3.7% when the cap is removed altogether, and a decline of 12% when Social Security's tax and benefit structure is held fixed at the current status-quo. Under this experiment, benefits decline for retirees with relatively favorable earnings histories, and they increase for retirees with unfavorable earnings histories. The effect of this experiment on labor supply, capital stock, and GDP are almost identical to when the cap is removed altogether, but the welfare consequences are slightly better.

It is worthwhile to note here that a crucial determinant of the labor supply responses to the elimination of the tax cap is a household's labor supply elasticity. If labor supply is highly elastic, then the effects of this policy change on Social Security's budget and the macroeconomy will be larger, compared to when labor supply is relatively inelastic. There is a large literature that has

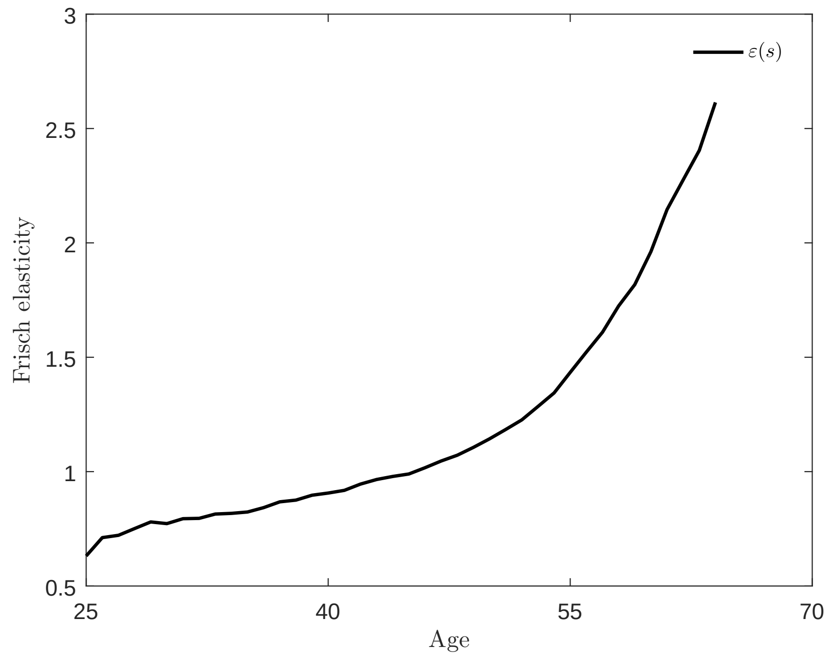

Figure 12: Cross-sectional mean Frisch elasticity of labor supply under Case 0.

examined the sensitivity of labor supply decisions to various tax-and-transfer programs. In general, this literature finds that while the aggregate labor supply consequences are largely insensitive to the value of a household's labor supply elasticity, the distribution of work hours over the life cycle, as well as the fraction of lifetime spent in employment, vary considerably (Rogerson and Wallenius, 2009; İmrohoroğlu and Kitao, 2009). Given this fact, it is useful to get a sense of the size of the household-level labor supply elasticities underlying the results described here.

With the utility function assumed in the current model, the Frisch elasticity of labor hours at age s, conditional on participation, is given by

$$\varepsilon(s) = \frac{1-h(s)}{h(s)} \frac{1+\eta(\sigma-1)}{\sigma}. \tag{26}$$

I report the cross-sectional mean Frisch elasticity between ages 25–65 under the current status-quo with the longevity improvement (i.e. Case 0) in Figure 12. The figure shows that the Frisch elasticities in this case range from 0.6 to around 2.6, with an average of around 1.2. With elasticities of this order, a 67% increase in Social Security's payroll tax rate leads to a 1.4% decline in the average hours worked, and a 4.6% decline in the average fraction of lifetime spent in employment. Rogerson and Wallenius (2009) simulate the labor supply consequences of a similar tax increase, and the above responses are consistent with their estimates in terms of both the extensive and the intensive margins of labor supply. Therefore, it appears that the household-level labor supply elasticities in the current model are well within the range typically observed in the data.

8 Alternative reform policies

As mentioned earlier, recent literature on Social Security reform has rigorously examined a large number of policy options to keep Social Security solvent, such as increasing the contribution rate,

	Case 0	Case 1	Case 2	Case 3
Labor	45.9	43.8	43.8	45.0
Avg. hours of market work per week per worker (25-55)	39.5	38.8	38.8	39.2
Capital	284.1	257.5	257.9	263.6
GDP	86.8	81.4	81.5	83.5
Ratio of Social Security expenditures to GDP	5.7%	6.9%	6.9%	6.9%
Interest rate	4.71%	5.03%	5.07%	5.08%
Wage	1.23	1.21	1.21	1.21

Table 15: Select macroeconomic variables under Case 3, compared to those under the other experiments.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 0	0.438 ($23,430)	0.345 ($38,678)	0.266 ($50,011)
Case 1	0.408 ($24,609)	0.331 ($44,519)	0.300 ($78,643)
Case 2	0.516 ($24,093)	0.409 ($39,466)	0.304 ($52,388)
Case 3	0.458 ($23,374)	0.396 ($38,487)	0.402 ($49,190)

Table 16: Average replacement rates and AIME values under Case 3, and also under the other experiments.

increasing the retirement age, and modifying the link between Social Security contributions and benefits (Huang et al., 1997; De Nardi et al., 1999; Kitao, 2014). In this section, I examine how the effects of eliminating Social Security's tax cap (i.e. Cases 1 and 2) compare to the effects of two reform policies commonly examined in this literature: separately increasing Social Security's payroll tax rate, and also its full retirement age.

First, I consider the increase in Social Security's payroll tax rate needed to keep expected benefits unchanged at their current level under the longevity improvement (Case 3). I find that a payroll tax rate of $\tau_{ss} = 0.13$ accomplishes this in the current model, which indicates a roughly 23% increase from the current contribution rate of $\tau_{ss} = 0.106$.[13] I report the macroeconomic consequences of this experiment in Table 15. It is clear from the table that increasing the payroll tax rate to keep Social Security's benefits unchanged under the longevity improvement has slightly better implications for aggregate labor, capital stock, and national income. Under this experiment, labor supply, capital, and GDP increase by 7.3%, 14%, and 9.7% from their respective baseline levels. These are larger than the corresponding increases noted both under Cases 1 and 2. Relative to the current status-quo under the longevity improvement (i.e. Case 0), this constitutes a 1.9%, 7.2%, and 2.3% reduction in labor supply, capital, and GDP, respectively. Evidently, the payroll tax increase, while keeping the current taxable maximum fixed, has a smaller impact on key macroeconomic aggregates, as compared to keeping the current payroll tax fixed and eliminating the taxable maximum.

The effect of this experiment on Social Security's implicit redistribution can be seen in Table 16. The table shows that Case 3 has a negative impact on the overall progressivity of Social Security: the average replacement rates are a considerably flatter function of past earnings (the AIME values), compared to those under Cases 1 and 2. This should not be surprising, as a higher payroll tax rate with the current taxable maximum makes Social Security's tax structure even more regressive than the baseline. The associated reduction in Social Security's insurance effects, along with the higher distortions from the payroll tax increase, causes significant welfare losses. In fact,

[13] Note that this tax increase is smaller than that identified in Kitao (2014), because while I consider only a longevity improvement in the current paper, Kitao (2014) considers a comparable longevity improvement along with a 50% reduction in the population growth rate.

	W
Case 0	-65.89
Case 1	-66.73
Case 2	-66.49
Case 3	-68.77

Table 17: Ex-ante expected utilities under Cases 0, 1, 2, and 3.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 1 (%)	-1.27	-0.98	0.30
Case 2 (%)	-0.56	-0.68	-1.36
Case 3 (%)	-3.46	-3.27	-2.27

Table 18: The consumption equivalence (ψ) under Cases 1, 2, and 3.

	Case 0	Case 1	Case 2	Case 3	Case 4
Labor	45.9	43.8	43.8	45.0	45.6
Avg. hours of market work per week per worker (25-55)	39.5	38.8	38.8	39.2	39.4
Capital	284.1	257.5	257.9	263.6	278.2
GDP	86.8	81.4	81.5	83.5	85.8
Ratio of Social Security expenditures to GDP	5.7%	6.9%	6.9%	6.9%	5.7%
Interest rate	4.71%	5.03%	5.07%	5.08%	4.82%
Wage	1.23	1.21	1.21	1.21	1.22

Table 19: Select macroeconomic variables under Case 4, compared to those under the other experiments.

the welfare implications of a higher payroll tax are considerably worse than what I obtain under Cases 1 and 2, at both the overall and the individual level (see Tables 17 and 18).

I now examine the second policy also widely considered in the literature: increasing Social Security's full retirement age, holding the payroll tax rate and the cap on the amount of taxable earnings fixed at their current status-quo (Case 4). I find that under this experiment, a full retirement age of 69 is needed to keep expected Social Security benefits roughly unchanged.[14] I report the macroeconomic consequences of this experiment in Table 19. As the table shows, increasing the full retirement age from 66 to 69 under the longevity improvement causes labor supply, capital, and GDP to increase by 8.8%, 21%, and 13% from their respective baseline levels. These changes are larger in magnitude than those obtained with the payroll tax increase, which in turn, are larger than those obtained when the tax cap is eliminated. Relative to the current status-quo (i.e. Case 0), this constitutes a 0.6%, 2.1%, and 0.2% reduction in labor supply, capital, and GDP, respectively. Therefore, postponing the retirement age by three years has only a minor effect on key macroeconomic aggregates under the longevity improvement, compared to all the other policies considered here.

I report the effect of this experiment on Social Security's implicit redistribution in Table 20. The table shows that even though postponing the full retirement age has no direct effect on Social Security's tax and benefit structure, Case 4 slightly increases the redistribution implicit in Social

[14]Because the time scale of the model is annual, the retirement age is allowed to increase only in annual increments. A full retirement age of 69 yields a small surplus in Social Security's revenues, as a result of which the budget-balancing benefits are slightly higher than the baseline. On the other hand, a full retirement age of 68 leaves a small deficit, which yields budget-balancing benefits that are slightly lower than the baseline.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 0	0.438 ($23,430)	0.345 ($38,678)	0.266 ($50,011)
Case 1	0.408 ($24,609)	0.331 ($44,519)	0.300 ($78,643)
Case 2	0.516 ($24,093)	0.409 ($39,466)	0.304 ($52,388)
Case 3	0.458 ($23,374)	0.396 ($38,487)	0.402 ($49,190)
Case 4	0.516 ($22,001)	0.407 ($36,248)	0.314 ($46,910)

Table 20: Average replacement rates and AIME values under Case 4, and also under the other experiments.

	W
Case 0	−65.89
Case 1	−66.73
Case 2	−66.49
Case 3	−68.77
Case 4	−66.62

Table 21: Ex-ante expected utilities under the alternative policy experiments.

Productivity fixed effect (p)	0.54	1.00	1.84
Case 1 (%)	−1.27	−0.98	0.30
Case 2 (%)	−0.56	−0.68	−1.36
Case 3 (%)	−3.46	−3.27	−2.27
Case 4 (%)	−1.03	−0.80	−0.43

Table 22: The consumption equivalence (ψ) under the alternative policy experiments.

Security, relative to the current status-quo (i.e. Case 0). The average replacement rates in this case are decreasing in past earnings (the AIME values), and they are also very similar to those under Case 2. Postponing the retirement age delays households' exit from the labor market, but because of the declining age-dependent component of labor productivity, this leads to a reduction in their AIME values relative to Case 0. These lower AIME values interact with Social Security's progressive benefit structure and yield an overall level of redistribution that is comparable to that obtained under Case 2.

The welfare implications of postponing the full retirement age from 66 to 69 under the longevity improvement are reported in Tables 21 and 22. Table 21 shows that Case 4 has the smallest negative effect on overall welfare, compared to all the other policies considered here. This is due the fact that under this experiment, Social Security expenditures remain roughly unchanged at the status-quo level of 5.7% of GDP, compared to increasing to almost 7% under all the other experiments. However, the disaggregated welfare effects in Table 22 show that for households with relatively unfavorable earnings histories, postponing the full retirement age delivers welfare effects that are slightly better than when the tax cap is eliminated altogether (Case 1), but slightly worse than when the cap is eliminated only from the amount of earnings subject to the payroll tax (Case 2).

To summarize, I find that increasing Social Security's payroll tax rate or the full retirement age has slightly better macroeconomic outcomes than eliminating Social Security's tax cap. Increasing the payroll tax from its current level of 10.6% to 13% leads to a smaller decline in labor supply, capital stock, and national income, from their levels under the current status-quo with the longevity improvement, but has a large negative effect on welfare. On the other hand, increasing the current full retirement age from 66 to 69 has almost negligible effects on key macroeconomic aggregates, and only a small negative effect on overall welfare.

9 Conclusions

The amount of earnings that can be annually taxed by Social Security is currently capped at \$118,500. This cap has recently drawn a lot of attention from politicians and policymakers as a potential institutional feature that can be used to solve Social Security's long-run insolvency. In this paper, I quantitatively examine if this cap can be an effective policy tool in solving Social Security's budgetary problems. To evaluate this question, I set up a calibrated general-equilibrium overlapping-generations model with mortality and labor income risk, and incomplete insurance markets. Then, I use this model to compute the change in benefits needed to keep Social Security's budget balanced under projected future U.S. demographics.

In general, my computational results suggest that eliminating the cap on the amount of earnings subject to the payroll tax can partially solve Social Security's future budgetary problems. I find that subjecting all earnings to the payroll tax and also counting them towards future benefits requires expected benefits to decline by less than 4%, on the average, to keep Social Security solvent. This is significantly smaller than the 12% budget-balancing decline in benefits needed when the cap is held fixed at its current level. This policy change increases Social Security's revenues by only 14%, because households that are no longer subject to the cap respond by reducing their labor supply and saving. Under this policy, most of the additional tax revenues are spent in paying benefits to wealthy retirees for whom the cap expires, which reduces the extent of redistribution implicit in Social Security.

I find that the fiscal advantages to Social Security are larger when the cap is eliminated only from taxes, but retained on the amount of earnings creditable towards Social Security benefits. In this case, the average budget-balancing change in expected benefits is actually positive: a 0.4% increase from the baseline. As before, Social Security's revenues increase by 14% under this experiment, but now these revenues are used overwhelmingly in paying benefits to low- and medium-income retirees. Both policies lead to an overall reduction in labor supply, capital stock, output, and overall welfare relative to the current status-quo, but the welfare losses are somewhat smaller when the cap is retained on the amount of earnings creditable towards Social Security benefits. Finally, I find that the macroeconomic effects of eliminating the tax cap are somewhat larger than when the current payroll tax rate and the full retirement age are separately increased to balance Social Security's budget.

There is a large literature that has considered modifications to the various institutional features of Social Security to keep the program solvent in the long run. These modifications have ranged from changing the payroll tax rate and the eligibility age, to a complete privatization of the existing Pay-As-You-Go structure. The current paper complements this literature by evaluating whether or not the cap on the amount of earnings subject to Social Security taxes can play an important role in this debate. The results in this paper suggest that this particular institutional feature may be a good candidate for partially solving Social Security's long-run budgetary problems.

References

Altig, D., Auerbach, A. J., Kotlikoff, L. J., Smetters, K. A., Walliser, J., 2001. Simulating Fundamental Tax Reform in the United States. The American Economic Review 91 (3), 574–595.

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

Auerbach, A., Kotlikoff, L., 1987. Dynamic Fiscal Policy. Cambridge University Press.

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

Bagchi, S., 2016. Is The Social Security Crisis Really as Bad as We Think? Macroeconomic Dynamics 20 (3), 737–776.

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

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., 2008. Optimal Fiscal Policy In The Design Of Social Security Reforms. International Economic Review 49 (1), 291–318.

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.

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.

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

Heathcote, J., Storesletten, K., Violante, G. L., 2010. The Macroeconomic Implications of Rising Wage Inequality in the United States. Journal of Political Economy 118 (4), 681–722.

Heathcote, J., Storesletten, K., Violante, G. L., 2014. Optimal tax progressivity: An analytical framework. Working Paper 19899, National Bureau of Economic Research.

Huang, H., İmrohoroğlu, S., Sargent, T. J., 1997. Two Computations to Fund Social Security. Macroeconomic Dynamics 1 (1), 7–44.

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.

İmrohoroğlu, S., Kitao, S., 2009. Labor supply elasticity and social security reform. Journal of Public Economics 93 (7-8), 867–878.

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

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

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

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

Nishiyama, S., Smetters, K., 2005. Does Social Security Privatization Produce Efficiency Gains? Working Paper 11622, National Bureau of Economic Research.

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

Rogerson, R., Wallenius, J., 2009. Micro and macro elasticities in a life cycle model with taxes. Journal of Economic Theory 144 (6), 2277–2292.

Sanders, B., 2015. Social Security Expansion Act. Legislative proposal, United States Senate.

Shrestha, L. B., 2006. Life Expectancy in the United States. CRS Report for Congress RL32792, Congressional Research Service.

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.

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.