Towson University
Department of Economics
Working Paper Series

Working Paper No. 2014-05

Can Removing the Tax Cap Save Social Security?

by Shantanu Bagchi

September, 2014

Can Removing the Tax Cap Save Social Security?*

Shantanu Bagchi[†]

September 7, 2014

Abstract

The maximum amount of earnings in a calendar year that can be taxed by Social Security in the U.S. is currently capped at \$106,800. In this paper, I use a general-equilibrium overlapping-generations model to examine if removing this cap can solve Social Security's budgetary problems. I find that in general, removal of the cap increases Social Security revenues, but by only a small percentage, and most of these extra revenues go towards paying benefits to high-income retirees no longer subject to the cap. Even when the cap is removed only from taxes but retained on the amount of earnings creditable towards Social Security benefits, the fiscal advantages are quite small.

JEL Classifications: E21, E62, H55

Keywords: Social Security; tax cap; mortality risk; productivity shock; partial insurance; general equilibrium

---

[*] I would like to thank Frank Caliendo, Jim Feigenbaum, and Scott Findley for useful comments and suggestions.

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

1 Introduction

The maximum amount of earnings in a calendar year that can be taxed by Social Security in the U.S. is currently capped at \$106,800. This cap, which is adjusted annually in proportion to wage growth, has been a salient feature of the U.S. Social Security program since its inception. However, the projected insolvency of Social Security, coupled with growing income inequality (Heathcote et al., 2010), has recently brought this particular institutional feature to the forefront of the national political debate. It has often been proposed that this cap be completely eliminated, or at the very least, be increased from its current level, to bring in additional revenues for Social Security from the wealthy, who are the beneficiaries of the cap, and for whom Social Security benefits are a relatively small fraction of retirement income (Reno and Lavery, 2005; Simpson and Bowles, 2010).

From an economic standpoint, any proposal to reform Social Security requires a careful evaluation of the potential advantages of changing a particular institutional feature of Social Security, with the potential costs associated with such a change. This particular methodological approach has engendered a large amount of research over the last few decades, wherein studies have looked at changing the contribution rate, the retirement age, the link between Social Security contributions and benefits, and a phased transition to a fully funded system, to name a few. Studies have carefully considered both the fiscal costs and benefits associated with modifying these institutional features, as well as their welfare consequences.[1]

This literature, however, has been relatively silent on the question of whether or not the cap on Social Security taxes can be an effective policy tool in improving the program's fiscal situation. The answer to this question is not clear because of two reasons. First, Social Security benefits in the U.S. are calculated based on a measure of average earnings through the work life, and earnings only up to the cap are creditable towards the benefits. If this historical link between the cap on taxes and earnings creditable towards the benefits is retained, then removing the cap will also expand the benefit payments and ultimately total Social Security expenditures. Second, even if the cap is removed only from taxes but retained on the benefits, the extra revenues will have a notable impact on Social Security's budget only if the households initially subject to the cap contribute a sizable fraction of total revenues. The ability of the cap in improving Social Security's fiscal situation will, therefore, depend on the quantitative importance of these mechanisms.

In this paper, I examine if removing the annual cap on taxes in the U.S. can solve Social Security's budgetary problems. To do this, I begin by constructing an overlapping-generations macroeconomic model with heterogeneous households, mortality risk, incomplete markets, and an annual cap on the amount of earnings that can be subject to the Social Security tax. In the model, Social Security provides partial insurance against mortality risk and unfavorable shocks to labor income.[2] Firms in the model maximize profit, the government provides public goods and Social Security, and the markets for all goods and services clear. I calibrate the model to the current U.S. economy, and then incorporate an empirically reasonable improvement in life expectancy. Finally, I compute the consequences of removing the annual cap on Social Security taxes in this environment. In the computations, I allow for all the household-level and macroeconomic adjustments to the longevity improvement, as well as to the institutional changes in Social Security.

---

[1] 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 (2013).

[2] While it is well-known that Social Security can provide partial insurance against shocks to labor income, Caliendo et al. (2013) have recently shown that Social Security fails to replace missing annuity markets (i.e. provide insurance against mortality risk), if one accounts for how Social Security affects the accidental bequests that households leave (and receive) in equilibrium. I avoid this complication by assuming that the government taxes away all accidental bequests and uses them to purchase public goods. Due to this assumption, households in the current model leave, but do not receive accidental bequests in equilibrium.

Intuitively, the idea of removing the cap 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 will likely come with little additional distortions, this policy change will fundamentally alter the pattern of redistribution implicit in the Social Security program. In an environment where Social Security partially replaces missing insurance markets, any change in this implicit redistribution will have varying welfare effects on the different household types. In this paper, I also evaluate the overall welfare effects of removing the annual cap on Social Security taxes, as well as the associated distributional consequences.

In general, I find that the fiscal advantages of removing the annual cap on Social Security taxes are quite small. As expected, with the Social Security taxes and benefits based on current law, the longevity improvement leads to a decline in the equilibrium level of benefits: Social Security expenditures increase at a faster rate than Social Security revenues. However, I find that on the average, benefits decline by roughly the same percentage, even when the annual cap on Social Security taxes is removed. While subjecting all earnings to the Social Security tax does generate extra revenues, these gains turn out to be quite small in magnitude relative to the overall Social Security budget. Due to the removal of the cap, all earnings are counted towards Social Security benefits, as a result of which most of these extra revenues are spent in paying higher benefits to retirees for whom the cap expires. Moreover, the tax payments from the households initially subject to the cap are only a small fraction of total Social Security revenues. Accounting for all of these effects, the average decline in Social Security benefits is almost identical to the case when the cap is held fixed at its current level in the U.S.

My computations also predict that the fiscal advantages to Social Security are only marginal, even if the annual cap is removed only from the amount of earnings subject to the Social Security tax, but retained on how much of those earnings can be counted towards future benefits. Removal of the cap from taxes increases Social Security revenues, while retaining the cap on benefits limits Social Security expenditures. In this case, I find that the longevity improvement still causes Social Security benefits to decline from the baseline level, but by only a slightly smaller percentage, relative to when the cap is held fixed at its current level in the U.S.

I also find that subjecting all earnings to the Social Security tax and counting them towards future Social Security benefits has a negative impact on overall welfare. Removal of the cap from taxes increases the average Social Security tax rates for high-income households, especially during their peak productivity years, and counting all earnings towards benefits makes Social Security less progressive than the baseline. The higher tax rates negatively affect the welfare of the high-income households, and the reduced progressivity negatively affects the welfare of the low- and medium-income households.

However, I find evidence of a slight welfare improvement when the cap is removed from the earnings subject to the tax, but retained on the earnings that can be counted towards Social Security benefits. Removing the cap on taxes increases the average tax rates for the high-income households, but retaining the cap on benefits makes the Social Security program more progressive than the baseline. I find that the high-income households are worse-off due to this policy, but their welfare losses are more than offset by the gains to the low- and medium-income households.

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 in the U.S. With benefits being paid out as defined by current law, significant payroll tax increases may be required to balance the Social Security budget in the long run (De Nardi et al., 1999). Kitao (2013) 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 from 66 to 73, 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 supplied over the life cycle. Conesa and Garriga (2009) argue that some of these distortions can be minimized with an age-dependent labor income tax structure, removal of the compulsory retirement age, and 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 only play a limited role in solving Social Security's budgetary problems.

The rest of the paper is organized as follows: Section 2 introduces the model and Section 3 describes the baseline calibration. In Section 4, I incorporate an empirically reasonable longevity improvement into the baseline model, and also examine its consequences on Social Security benefits based on current law. I define the computational experiments in Section 5, and discuss the quantitative results in Section 6. Finally, I conclude in Section 7.

2 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, and a fraction f_i of newborns in a cohort receive a permanent productivity shock φ_i, where $\sum_i f_i = 1$. Maximum lifespan is T, and households face an unconditional probability $Q(s)$ of surviving to age s. Therefore, total population at date t

$$P(t) = \sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s) \tag{1}$$

grows at rate n over time, where $N(t-s)$ is the size of the cohort born at date $t-s$. There is also labor-augmenting technological progress at the rate of g per annum in the model economy.

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.[3] 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%.

At each date, surviving households earn labor income if they work, and they also collect Social Security benefits starting at an exogenously specified retirement age T_{r}. 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 government funds the public goods purchases using the proceeds from the estate tax and also a proportional tax on labor income, and it funds Social Security through a payroll tax on labor income. Finally, the amount of labor income that can be taxed and counted towards future Social Security benefits is capped at a maximum level at every age.

---

[3] 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).

2.1 Preferences

Period utility depends on both consumption (c) and the fraction of total time endowment enjoyed in leisure (l). It has the standard CIES form

$$u(c,l) = \begin{cases} \frac{\left(c^{\eta}l^{1-\eta}\right)^{1-\sigma}}{1-\sigma} & \text{if} \quad \sigma \neq 1 \\ \ln\left(c^{\eta}l^{1-\eta}\right) & \text{if} \quad \sigma = 1 \end{cases} \tag{2}$$

where η is the share of consumption, and σ is the inverse of intertemporal elasticity. Expected lifetime utility from the perspective of a household of type i born at date t is

$$U_i = \sum_{s=0}^{T} \beta^s Q(s)\, u\left(c_i(t+s,t), l_i(t+s,t)\right) \tag{3}$$

where β is the discount factor. Also, since I define leisure as a fraction of the total time endowment, $0 \leq l_i(t+s,t) \leq 1$.

2.2 Income

A household of type i born at date t earns after-tax wage income at age s if labor supplied $(1 - l_i(t+s,t))$ is positive. After-tax wage income is given by

$$y_i^{at}(t+s,t) = (1 - \tau_{SS}(D_i(t+s,t)) - \tau_y)y_i(t+s,t), \tag{4}$$

where $y_i(t+s,t) = (1 - l_i(t+s,t))w(t+s)e(s)\varphi_i$ is before-tax wage income, $\tau_{SS}(D_i(t+s,t))$ is the tax function for Social Security, and τ_y is the proportional income tax rate. The wage rate at date $t+s$ is $w(t+s) = w(t)(1+g)^s$, the efficiency endowment at age s is $e(s)$, and φ_i is the permanent productivity shock. Finally, $D_i(t+s,t)$ is the ratio of before-tax wage income to the earnings cap $\bar{y}(t+s,t)$ at age s, or

$$D_i(t+s,t) = \frac{(1 - l_i(t+s,t))w(t+s)e(s)\varphi_i}{\bar{y}(t+s,t)}. \tag{5}$$

The average tax rate for Social Security, given by the tax function $\tau_{SS}(D_i(t+s,t))$, is constant as long as $D_i(t+s,t) \leq 1$, but starts to decline monotonically when $D_i(t+s,t) > 1$. This is because the marginal Social Security tax rate on any income above the cap is equal to zero.

The permanent productivity shock φ_i affects the after-tax wage income both directly and indirectly: households with a favorable productivity shock are also likely to supply more labor. Finally, households that survive past the retirement age T_{r} also receive Social Security benefits $b_i(t+s)$, which are positively linked to their past work-life income.[4]

It is worth noting that in the current model, the shocks to labor productivity are ex-ante, or are realized before the agents enter the model. An alternative specification used in the literature assumes that the shocks are ex-post, or are realized after the agents enter the model.[5] The key difference between these two specifications is that in the ex-ante case, households accumulate wealth

---

[4] Note that because the current model does not entertain the decision to claim Social Security benefits, it most likely underestimates the distortionary effect of Social Security on labor supply. In the U.S., households can start collecting Social Security benefits as early as age 62, and can delay collection to as late as age 70. Based on how early or late the actual collection date is from the full-retirement age, households receive an adjustment in their benefits. These adjustments, if actuarially unfair, are an additional source of distortion to labor supply, which the current model ignores.

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

due to purely life-cycle motives, whereas in the ex-post 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 research 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 turns out that even in a model with borrowing constraints, the manner in which the labor productivity shocks are modeled has little effect on the macroeconomic and distributional effects of Social Security (Huggett and Ventura, 1995, 1999).

2.3 Retirement Benefits

The Social Security benefit at age $s > T_{\mathrm{r}}$ for a household with the productivity shock φ_i is $b_i(t+s)$, which is progressively linked to past work-life income. The amount of period earnings that can be counted towards Social Security benefits are also capped at $\bar{y}(t+s,t)$. The average amount of work-life income creditable towards Social Security for a household with the productivity shock φ_i is calculated as

$$\frac{1}{T_{\mathrm{r}}}\sum_{s=0}^{T_{\mathrm{r}}} \min\left\{(1-l_i(t+s,t))w(t+s)e(s)\varphi_i,\ \bar{y}(t+s,t)\right\}. \tag{6}$$

This measure of work-life income creditable towards Social Security is identical to the Average Indexed Monthly Earnings ($AIME$) measured by the Social Security Administration (SSA). In the model, the government calculates the retirement benefit based on this measure of Social Security contributions during the work life, but it also adjusts the level of benefits so that the budget for Social Security is balanced.

2.4 A Household's Optimization Problem

A household of type i born at date t faces the following optimization problem

$$\max_{c_i,l_i} \sum_{s=0}^{T} \beta^s Q(s)\, u\left(c_i(t+s,t), l_i(t+s,t)\right) \tag{7}$$

subject to

$$c_i(t+s,t) + k_i(t+s+1,t) = (1+r)k_i(t+s,t) + y_i(t+s,t) \tag{8}$$

$$y_i(t+s,t) = (1-\tau_{SS}(D_i(t+s,t)) - \tau_y)(1-l_i(t+s,t))w(t+s)e(s)\varphi_i + \Theta(s-T_{\mathrm{r}})b_i(t+s) \tag{9}$$

$$0 \leq l_i(t+s,t) \leq 1 \tag{10}$$

$$k_i(t,t) = k_i(t+T+1,t) = 0 \tag{11}$$

where

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

is a step function.

2.5 Technology and Factor Prices

Output is produced using a Cobb-Douglas production function with inputs capital, labor and a stock of technology $A(t)$

$$Y(t) \;=\; K(t)^{\alpha}\,(A(t)L(t))^{1-\alpha} \tag{12}$$

where $A(t) = A(0)(1+g)^t$, α is the share of capital in total income and $A(0)$ is the initial stock of technology. Firms face perfectly competitive factor markets, which implies

$$r = MP_K - \delta = \alpha\left[\frac{K(t)}{A(t)L(t)}\right]^{\alpha-1} - \delta \tag{13}$$

$$w(t) = MP_L = A(t)(1-\alpha)\left[\frac{K(t)}{A(t)L(t)}\right]^{\alpha} \tag{14}$$

where δ is the depreciation rate of physical capital, and $w(t)$ is the wage rate at time t. In the steady-state, the wage rate grows at rate g per annum, and the rate of return r is constant.

2.6 Aggregation

Aggregate capital stock and labor supply are given by

$$K(t) = \sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s)\,k_i(t, t-s-1) \tag{15}$$

$$L(t) = \sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s)\,\{1 - l_i(t, t-s)\}\,e(s)\varphi_i. \tag{16}$$

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

$$\begin{aligned} B(t) = (1+r)&\left[\sum_i f_i \sum_{s=0}^{T} \{N(t-s)Q(s) - N(t-s-1)Q(s+1)\}\,k_i(t, t-s-1)\right] \\ &- \sum_i f_i \sum_{s=0}^{T} (N(t-s+1) - N(t-s))\,Q(s)k_i(t+1, t-s). \end{aligned} \tag{17}$$

The budget-balancing condition for Social Security is given by

$$\begin{aligned} &\sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s)\tau_{SS}(D_i(t, t-s))\,(1 - l_i(t, t-s))\,w(t)e(s)\varphi_i \\ &\qquad = \sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s)\Theta(s - T_r)b_i(t). \end{aligned} \tag{18}$$

Finally, the government also adjusts the income tax rate (τ_y) such that the revenues from the estate tax and the income tax are sufficient to finance its expenditures

$$B(t) + \sum_i f_i \sum_{s=0}^{T} N(t-s)Q(s)\tau_y\,(1 - l_i(t, t-s))\,w(t)e(s)\varphi_i = G(t), \tag{19}$$

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

2.7 Competitive Equilibrium

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

1. cross-sectional consumption allocations $\{c_i(t, t-s)\}_{s=0}^{T}$, asset allocations $\{k_i(t, t-s)\}_{s=0}^{T}$, and labor supply allocations $\{1 - l_i(t, t-s)\}_{s=0}^{T}$,
2. aggregate capital stock $K(t)$ and labor $A(t)L(t)$,
3. rate of return r and wage rate $w(t)$, and
4. Social Security benefits $b_i(t)$

that

1. solves the households' optimization problems,
2. equilibrates the factor markets,
3. balances the Social Security budget, and
4. satisfies the steady-state requirements $c_i(t, t-s) = c_i(t+s, t)(1+g)^{-s}$, $k_i(t, t-s) = k_i(t+s, t)(1+g)^{-s}$, and $b_i(t+s) = b_i(t)(1+g)^{s}$.

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.

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

In computing this equilibrium, I focus only on steady-state analysis and set $t = 0$, and I also normalize the initial newborn cohort size and technology to $N(0) = A(0) = 1$.

3 Calibration

I use empirical evidence from various sources to assign values to the model's parameters. I set the population growth rate to $n = 1\%$, and the rate of technological progress to $g = 1.56\%$, which is the trend growth rate of per-capita income in the postwar U.S. economy (Bullard and Feigenbaum, 2007). I assume that households enter the model at actual age 25, which corresponds to the model age of zero. To get the survival probabilities, I use Feigenbaum's (2008a) sextic fit to the mortality data in Arias (2004), which is given by

$$\begin{aligned} \ln Q(s) \;=\; & -0.01943039 + \left(-3.055 \times 10^{-4}\right) s + \left(5.998 \times 10^{-6}\right) s^2 \\ & + \left(-3.279 \times 10^{-6}\right) s^3 + \left(-3.055 \times 10^{-8}\right) s^4 + \left(3.188 \times 10^{-9}\right) s^5 \\ & + \left(-5.199 \times 10^{-11}\right) s^6 \end{aligned} \tag{21}$$

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 $\bar{T} = 75$. The resulting survivor function is plotted in Figure 1. Under these survival probabilities, the model life expectancy at birth turns out to be about 79 years.

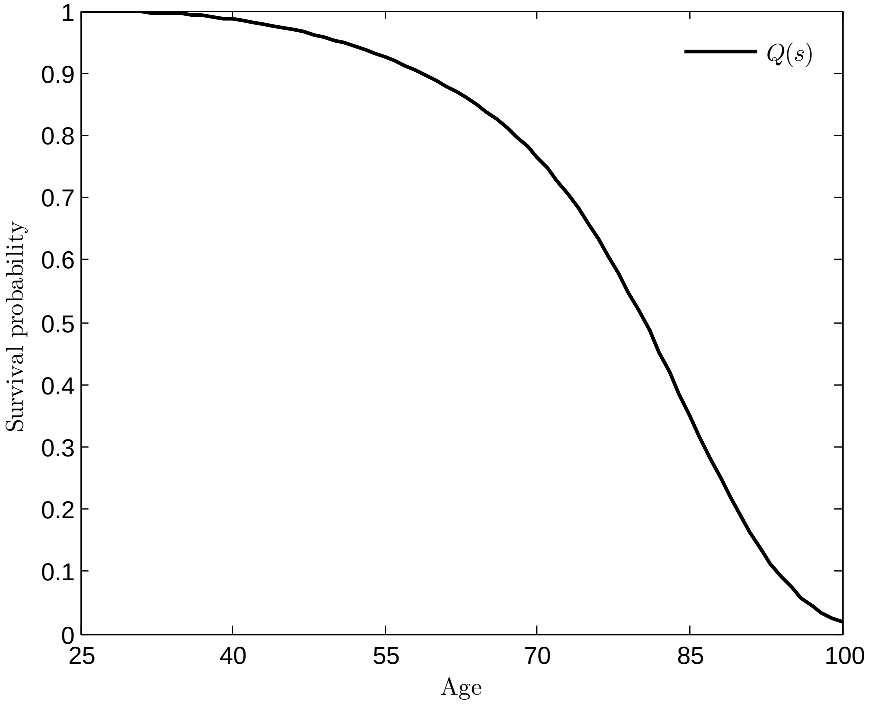

Figure 1: Survival probabilities from Feigenbaum's (2008a) sextic fit to the mortality data in Arias (2004).

I set the retirement age in the model to $T_{\mathrm{r}} = 41$, which corresponds to the current actual full retirement age of 66 in the U.S. Following İmrohoroğlu et al. (1995) and Conesa and Garriga (2008), I parameterize the efficiency endowment profile $e(s)$ using data from Hansen (1993). However, it is well known that efficiency 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 efficiency data in Hansen (1993) only for ages 25-65, which gives

$$\begin{aligned} \ln e(s) &= -3.273 \times 10^{-5} + \left(3.7484 \times 10^{-2}\right) s + \left(-1.7541 \times 10^{-3}\right) s^2 \\ &\quad + \left(3.4625 \times 10^{-5}\right) s^3 + \left(-2.7949 \times 10^{-7}\right) s^4 \end{aligned} \tag{22}$$

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

$$\ln e(s) = -f_0 - f_1 s - 0.01 s^2 \tag{23}$$

and parameterize f_0 and f_1 such that $e(s)$ is continuous and once differentiable at age $s = 40$.[6] Note that the coefficient of 0.01 on the squared term in (23) ensures that households do not continue to work beyond age 70.[7] The resulting efficiency endowment profile is plotted in Figure 2.

The next step is to calibrate the Social Security tax function in the model $\tau_{SS}(D)$ to the corresponding institutional feature in the U.S. The combined tax rate for the Old-Age and Survivors Insurance (OASI) part of Social Security in the U.S. is 10.6%, and this rate is applied to all labor income up to the earnings cap.[8] This cap is adjusted annually relative to the average wage in

---

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

[7] Historically, the employment-to-population ratios for ages 70 and above in the U.S. have been less than 10%.

[8] In reality, this rate is evenly split between the employer and the employee, but the standard hypothesis is that due to pressures in the labor market, the employee ends up bearing the full burden of the tax.

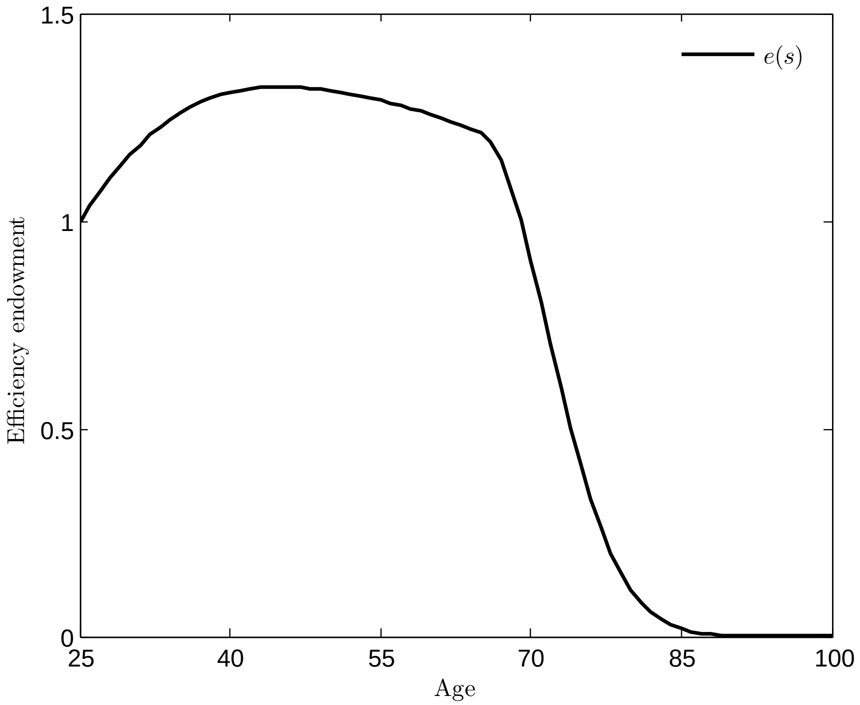

Figure 2: Efficiency endowment profile fitted to data from Hansen (1993).

the U.S. For example, the earnings cap was set at $76,200 in the year 2000, but was adjusted to $106,800 in the year 2010. During the same period, the national average wage index increased from $32,155 to $41,674.[9] Huggett and Ventura (1999) calculate that the ratio of the earnings cap to the average wage index has averaged at about 2.47 in the U.S., using which I set the cap in the model to 2.47 times the average wage.

Below the earnings cap, the average tax rate for the OASI component of Social Security in the U.S. is 10.6%, but it starts to decline linearly above the cap. For example, in the year 2000, the average OASI tax rate for an individual with an annual income of $200,000 was

$$\frac{0.106 \times \$76,200}{\$200,000} \approx 4.0\%,$$

and for an individual with an annual income of $400,000 was

$$\frac{0.106 \times \$76,200}{\$400,000} \approx 2.0\%.$$

I calibrate the tax function in the model $\tau_{SS}(D)$ such that it matches the relationship between the average OASI tax rate and the ratio of labor income to the earnings cap in the U.S. To ensure that the function $\tau_{SS}(D)$ is smooth across the cap for computational purposes, I approximate it using cubic spline interpolation. In Figure 3, I plot the estimated function $\tau_{SS}(D)$ along with the actual relationship between the average OASI tax rate and the ratio of labor income to the earnings cap in the U.S. Note that the tax rate is plotted for a maximum level of income of $400,000, which is equal to 3.75 times the current cap of $106,800.

---

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

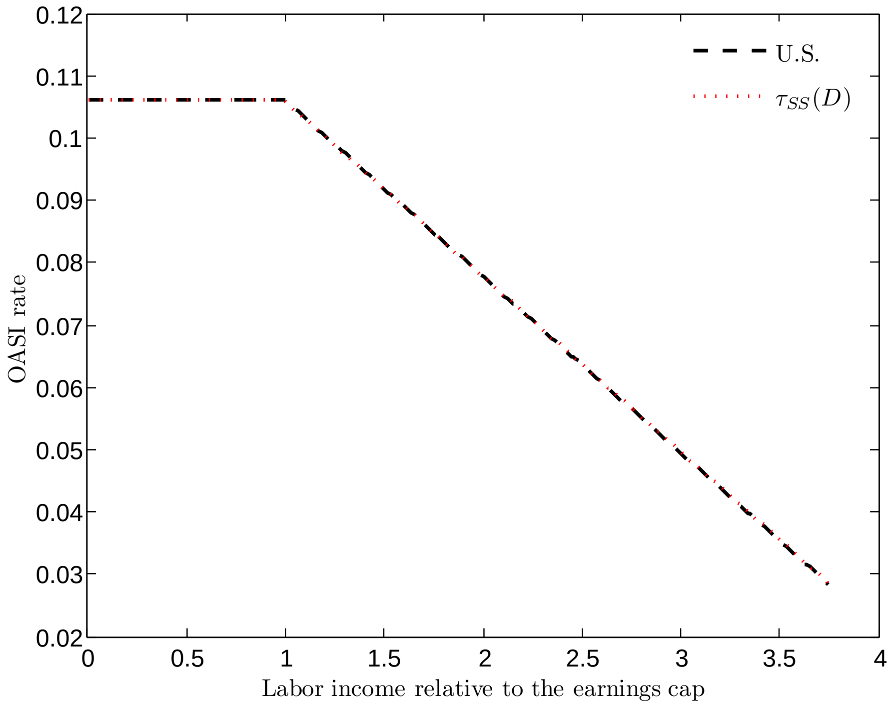

Figure 3: The average OASI tax rate in the U.S. and the estimated tax function $\tau_{SS}(D)$.

The Social Security retirement benefit in the U.S. is a concave (piecewise linear) function of average past work-life income. The SSA measures what is known as the Average Indexed Monthly Earnings ($AIME$) for every covered individual, and then replaces a fraction of the $AIME$. In the model, I approximate the $AIME$ using the average past work-life income variable defined in equation (6).

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 earnings replaced by Social Security. For example, in the year 2000, the OASI benefit in the U.S. 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% and 124% 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. Finally, I set the maximum creditable earnings equal to the earnings cap, which is 2.47 times the average wage.

The historically observed value of capital’s share in total income in U.S. ranges between 30-40%, so I set $\alpha = 0.35$, and following Stokey and Rebelo (1995), I set the depreciation rate to $\delta = 0.06$. To calibrate the permanent productivity shock and its distribution, I follow Zhao (2014) and assume that $\ln \varphi \sim N\left(0, \sigma_{\varphi}^2\right)$, and then set $\sigma_{\varphi}^2 = 0.65$ to match 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).

Once all the observable parameters have been assigned empirically reasonable values, I calibrate the unobservable preference parameters σ (IEIS), β (discount factor) and η (share of consumption in period utility), and the income tax rate (τ_y) such that the baseline equilibrium jointly matches the following targets:

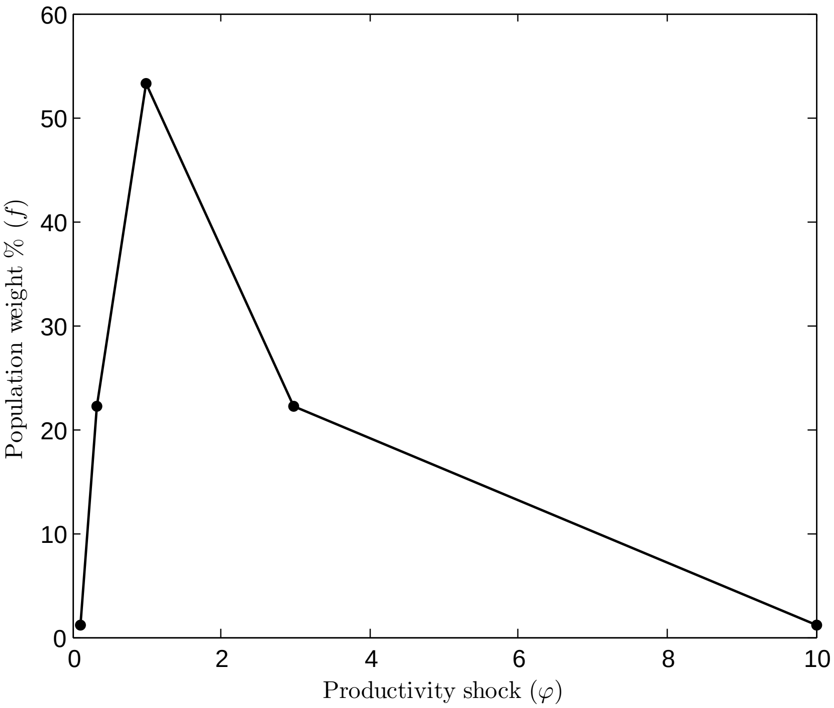

Figure 4: The distribution of the permanent productivity shock.

σ	β	η	τ_y
4	0.9918	0.3038	0.186

Table 1: Unobservable parameter values under the baseline calibration.

- an equilibrium capital-output ratio of 2.5,
- an average of 34 hours per week spent on market work between ages 25-55, and
- a share of government expenditures in GDP of 15%.

It is useful to note here that the target capital-output ratio here is slightly lower than 3.0, which the most commonly used value for the U.S. capital-output ratio in the macro-public finance literature. However, Carroll (1998) points out that households at the top of the wealth distribution account for a large share of overall wealth in the U.S., and their saving behavior is not well explained by the standard Life Cycle/Permanent Income Hypothesis. I adjust the target capital-output ratio downward to account for this fact.

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 $\sigma^c = 1 + \eta(\sigma - 1) = 1.91$, 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 means of consumption and labor hours are reported in Figures 5 and 6.[10] Note that in calculating labor hours per week, I net out eight hours per night as sleep time.

---

[10] The cross-sectional mean of a variable $x_i(t, t-s)$ is calculated using the formula $\bar{x}(t, t-s) = \sum_i f_i x_i(t, t-s)$.

	Target	Model
Capital-output ratio	2.5	2.51
Avg. hours of market work per week between ages 25-55	34	34.2
Share of govt. expenditures in GDP	0.15	0.151

Table 2: Model performance under the baseline calibration.

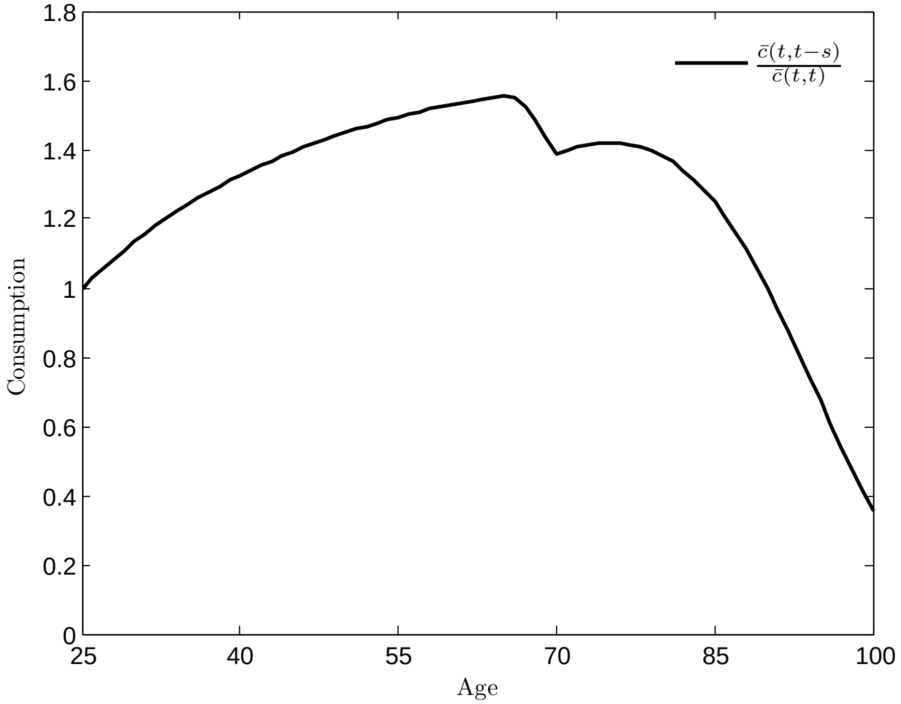

Figure 5: Cross-sectional mean of consumption (normalized by consumption at age 25) under the baseline calibration.

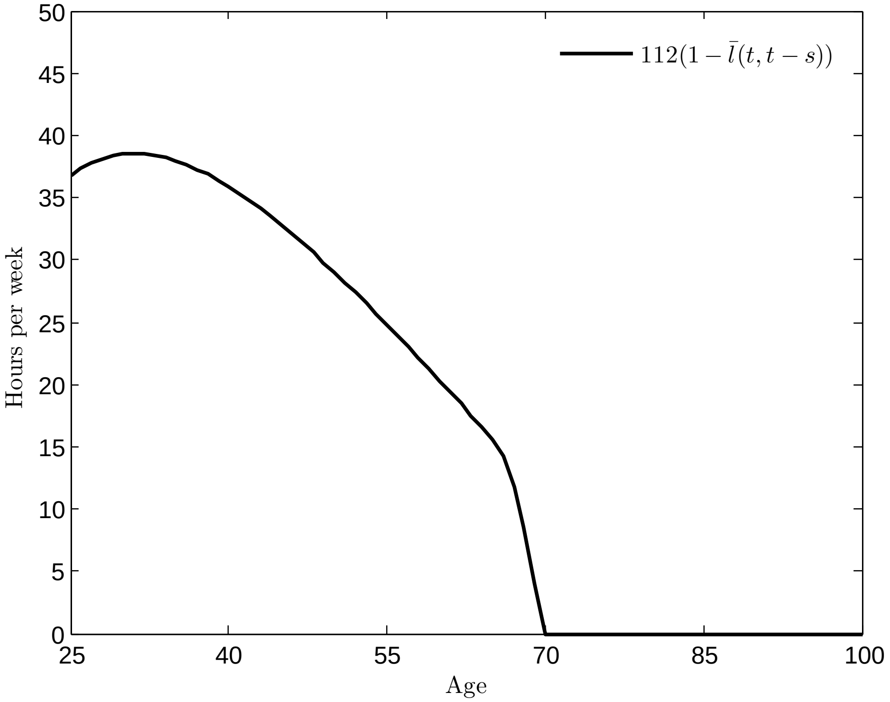

Figure 6: Cross-sectional mean of labor hours per week under the baseline calibration.

Whether or not the cap on Social Security taxes binds for a particular household depends on three key factors: the productivity shock it experiences, its life-cycle pattern of labor hours, and finally the interaction of the labor hours with the life-cycle endowment profile. Unconditionally, the cap is more likely to bind for a household with a favorable productivity shock. Conditional on a particular realization of the shock, the cap is more likely to bind at the ages during which before-tax wage income is near or at its peak. I report in Figure 7 the before-tax wage income profile in the baseline calibration for each household type, along with the cap.

It is clear from Figure 7 that in the baseline calibration, the only households for whom the cap binds are the ones with the most favorable realization of the productivity shock. For these households, the cap binds through most of the work life, during which before-tax income initially increases, reaches its peak, and then declines. During these ages, the Social Security tax rate for these households declines from 10.6%, as the marginal tax rate on their wage income above the cap is equal to zero (see Figure 8).

The minimum value of the Social Security tax rate in the baseline calibration is 9.04%, which corresponds to the point where the before-tax income of a household with the most favorable productivity shock in the model reaches its peak. In the U.S., this tax rate would roughly correspond to a middle-aged individual whose annual before-tax income is roughly 1.8 times the cap. For the year 2010, this amount in dollars is roughly $192,000, which falls in the 95th income percentile based on the estimates from the Current Population Survey (CPS) and the Panel Study of Income Dynamics (Gouskova et al., 2010).

Finally, due to the concave Social Security benefit-earnings rule, a household with an unfavorable productivity shock in the baseline calibration receives a less-than-proportional adjustment in their retirement benefits. To measure the pattern of redistribution implicit in the Social Security program, I compute the ratio of the share of total benefits received by retirees of each type, to the

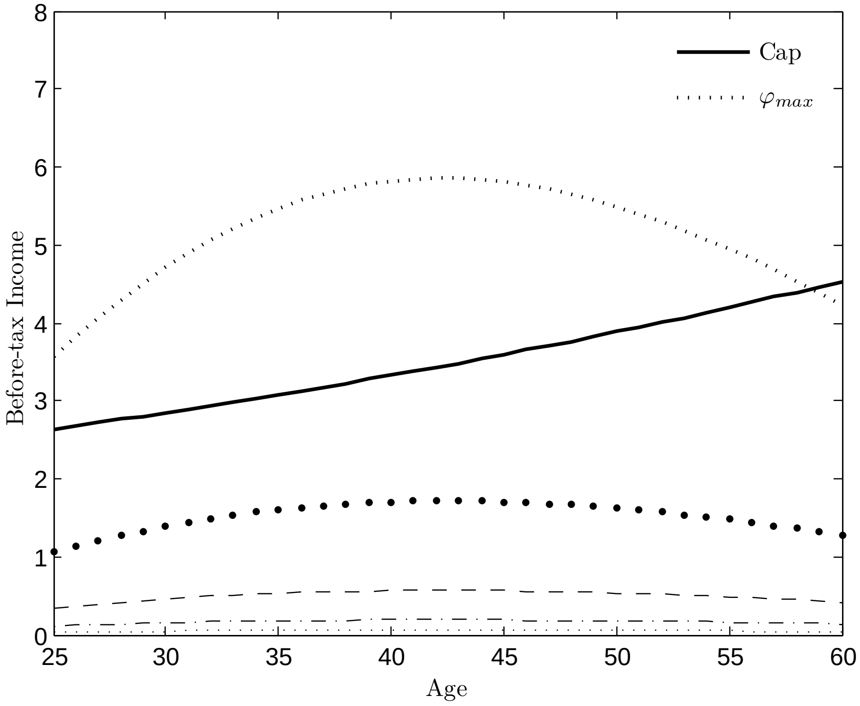

Figure 7: Before-tax wage income in the baseline calibration. The solid line is the cap.

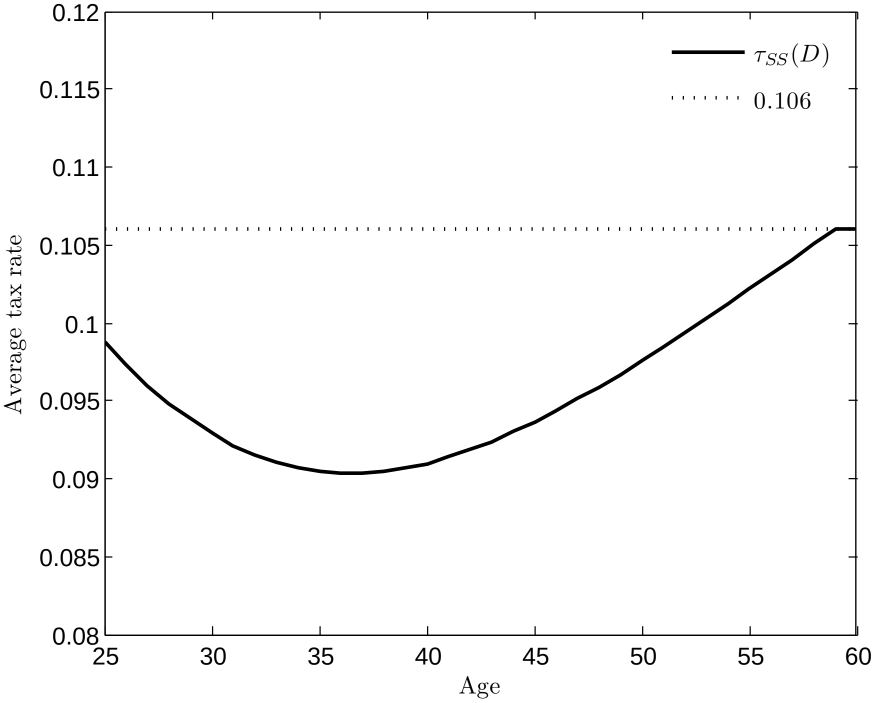

Figure 8: Average Social Security tax rate in the baseline calibration for the households with the most favorable productivity shock.

φ	0.1	0.34	1.0	2.98	10.0
Ratio	1.48	1.48	1.31	0.79	0.43

Table 3: The ratio of the share of total Social Security benefits received to the share of total taxes paid for each household type in the baseline calibration.

share of total taxes paid by workers of each type. I report this ratio in Table 3. For a Social Security program that is perfectly proportional and has zero implicit redistribution, this ratio would be equal to unity for each household type. It is clear from the table that in the baseline calibration, there are positive fiscal transfers occurring from the workers with favorable productivity shocks to the retirees with unfavorable shocks. This is because the calculated ratio is larger than unity for households at the bottom-end of the productivity distribution, and is smaller than unity at the top-end.

4 The Longevity Improvement

Based on the life tables, improvements in longevity experienced in the U.S. have followed a specific pattern: old-age survivorship has increased at a faster rate in the later half of the twentieth century, making the population survival curve more rectangular (Arias, 2004). A convenient way to incorporate an improvement in longevity into the baseline model that matches this observed fact is to reduce the baseline age-specific death rates $h(s)$ based on the following formula:

$$h_n(s) = h(s) - \theta s^{\nu}, \tag{24}$$

where θ and ν are positive constants. I set these parameters to $\theta = 10^{-6}$ and $\nu = 1.8509$, under which the life expectancy in the model is 85 years, which matches the 2011 Social Security Trustees Report's average period life expectancy projection for the year 2085 under the intermediate assumption. The survivor function resulting from this longevity improvement is compared to the baseline in Figure 9.

In order to understand the consequences of removing the annual cap on Social Security taxes in the U.S., it is first useful to compute the results of an experiment in which all the institutional features of Social Security, including the cap, are held fixed at their baseline levels, but only the survival probabilities are changed to reflect the improvement in longevity. The goal of this experiment is to identify the likely effect of the longevity improvement on Social Security in an environment where the future benefit payments are based on current law.

It is well known that because of the improvements in longevity, Social Security in its current form in the U.S. is insolvent in the long run. According to the projections of the SSA, 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. However, actuaries of the SSA also estimate that increases in the payroll tax rate of the order of 2-5 percentage points will be sufficient to balance the budget in the long run.

A number of studies contend that the SSA's actuarial projections overestimate the extent of the insolvency in Social Security, and underestimate the tax increases that will be needed to balance Social Security's budget in the long run. Bagchi (2012) argues that because the SSA's projections do not account for how key household-level and macroeconomic variables will respond to the longevity improvements, the actuarial estimates do not satisfy the Lucas Critique.[11] Once these responses are

---

[11] See Ljungqvist (2008) for further details on the Lucas Critique.

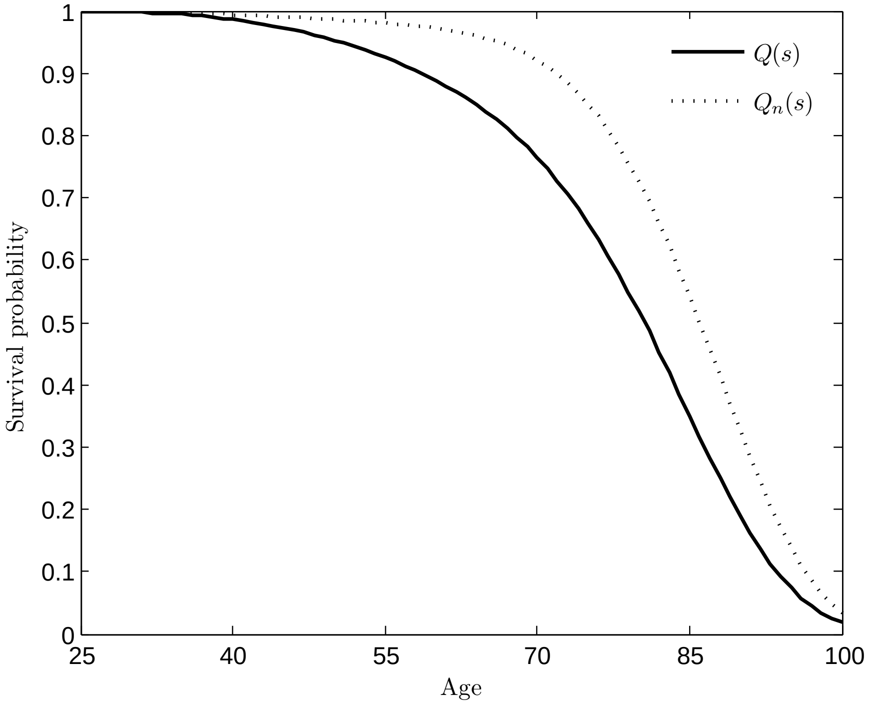

Figure 9: Baseline and the projected survival probabilities.

φ	0.1	0.34	1.0	2.98	10.0	Average
% change	-14.3	-14.3	-17.4	-15.9	-18.6	-16.4

Table 4: Effect of the longevity improvement on equilibrium Social Security benefits under current law.

accounted for, Bagchi (2012) shows that longevity improvements lead to a decline in Social Security benefits that is only two-thirds of what the SSA estimates. Similarly, De Nardi et al. (1999) find that balancing the Social Security budget will require an additional 17.1 percentage points tax increase on the top of the SSA's actuarial projections, simply because the SSA's projections do not account for the distortionary effect that the higher taxes will have on the tax base. Households will respond to the higher Social Security taxes by working and saving less, because of which the tax base will shrink from its current level.

Because the current model is an equilibrium model of the economy, it accounts for all the household-level and macroeconomic adjustments to demographic or policy-induced changes in the economy. In Table 4, I report the percentage change in Social Security benefits for each realization of the productivity shock, in a new equilibrium with the improved longevity, and the benefits based on current law.

Table 4 shows that the effect of improved longevity on equilibrium Social Security benefits in the current model are very similar to those identified in Bagchi (2012). On the average, Social Security benefits decline by slightly over 16% from their baseline level, but because benefits are also linked to the past work-life income, households with different productivity shocks experience different declines in the level of their benefits. It is worth noting that the percentage declines identified in Table 4 are significantly smaller that the actuarial estimates of the SSA, which are of the order of 30%. This is because the current model accounts for the household-level and macroeconomic

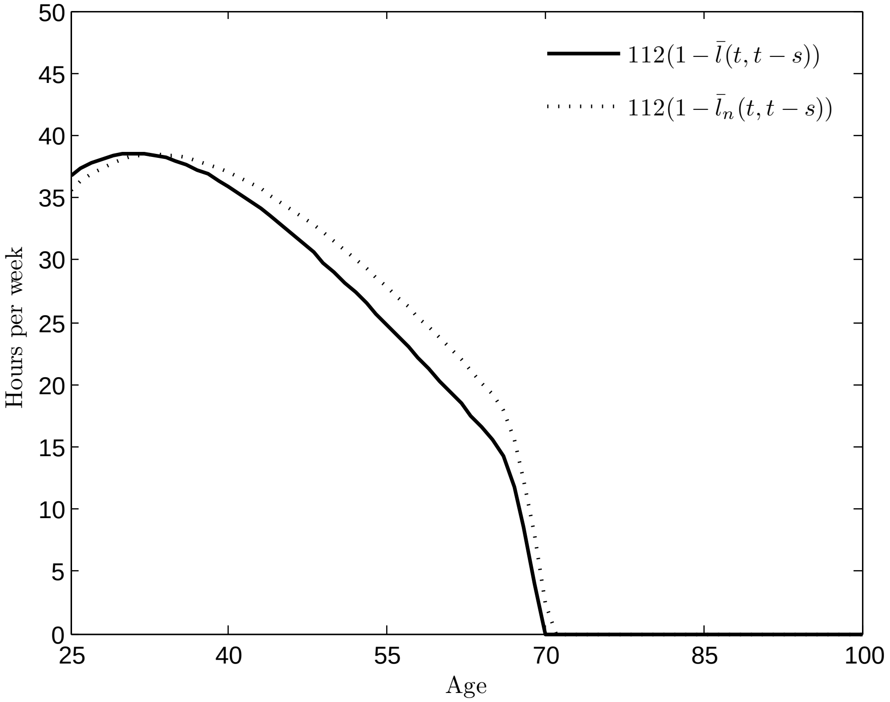

Figure 10: Cross-sectional mean of labor hours per week under the longevity improvement, along with the baseline calibration.

adjustments to the longevity improvement. Households respond to the increased life expectancy by supplying more labor and also saving more, which leads to a natural expansion of the Social Security tax base.

In Figure 10, I plot the cross-sectional mean of labor hours per week under the longevity improvement, along with those under the baseline calibration. It is clear from the figure that the improvement in longevity leads to an overall increase in the cross-sectional hours per week, and also delays retirement. The average hours per week between ages 25-55 increases from 34.1 in the baseline calibration, to 35.2 with the improved longevity. Even though there is a slight decline in the weekly hours prior to age 33, the increased hours for the rest of the work life are more than sufficient to compensate for it. The increased hours per week interact with the life-cycle endowment profile, and ultimately lead to an increase of 8.3% in aggregate labor supply from the baseline level.

Intuitively, the effect of a longevity improvement on household saving is less clear, especially because households increase their hours per week. With fixed labor, the only way to smooth consumption over a longer expected lifespan is to save more, but this need not necessarily be true with endogenous labor supply. To examine what happens to saving due to the longevity improvement in the current model, I report in Figure 11 the cross-sectional mean of asset holdings under the longevity improvement, along with those under the baseline calibration. It is clear from Figure 11 that the improvement in longevity also causes households to save more than the baseline level. This leads to an increase of 21% in the aggregate capital stock, and a roughly 4% increase in the equilibrium wage rate. This increase, along with the higher aggregate labor supply, leads to a natural expansion of the Social Security tax base that the SSA's actuarial estimates do not account for. Because of this reason, the percentage decline in the Social Security benefits predicted by the current model are considerably lower than the SSA's estimates.

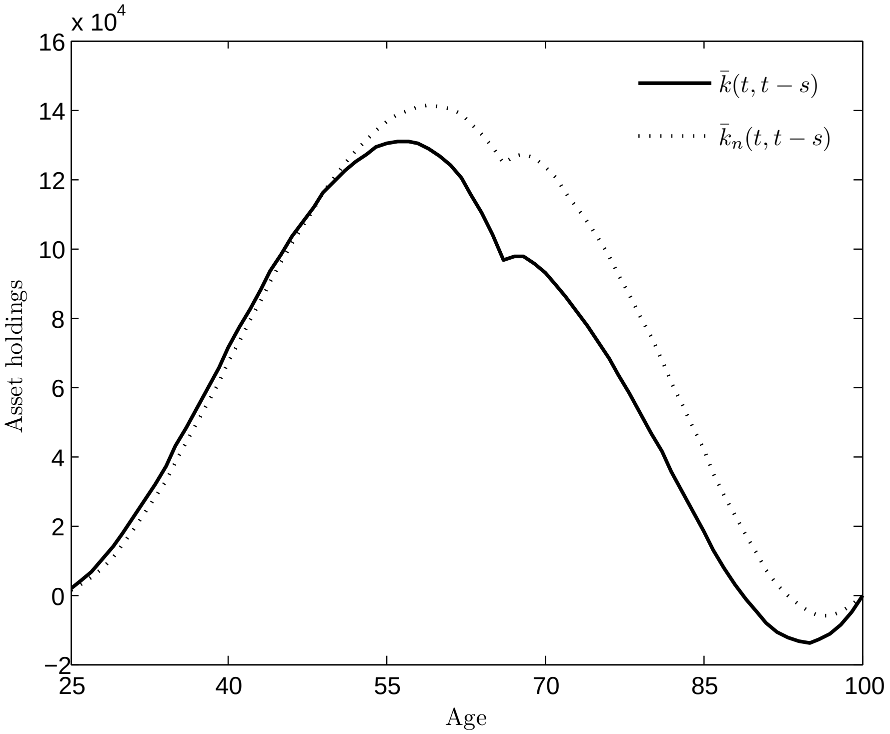

Figure 11: Cross-sectional mean of asset holdings under the longevity improvement, along with those under the baseline calibration (in US$).

Even with the improved survival probabilities, the cap on Social Security taxes binds only for the households with the most favorable productivity shock. However, because the longevity improvement leads to increased labor hours per week, especially during the peak productivity years, it causes before-tax income to exceed the cap by a larger extent relative to the baseline. Due to this reason, the average Social Security tax rate for a household with the most favorable productivity shock drops to a new minimum of 8.96% during the work life (see Figure 12).

5 Two Experiments With the Cap

As explained earlier, Social Security benefits in the U.S. are calculated as a function of average earnings through the work life, and earnings only up to the cap are counted towards the benefits. Therefore, the cap also serves as a de-facto limit on the amount of benefit payments from Social Security. Given this fact, there are two possible ways in which the cap can be used as a policy tool in improving the 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 retaining 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 contributions would limit the amount of benefit payments. This option would break the historic link between the cap on taxes and benefits in the U.S. Social Security program.

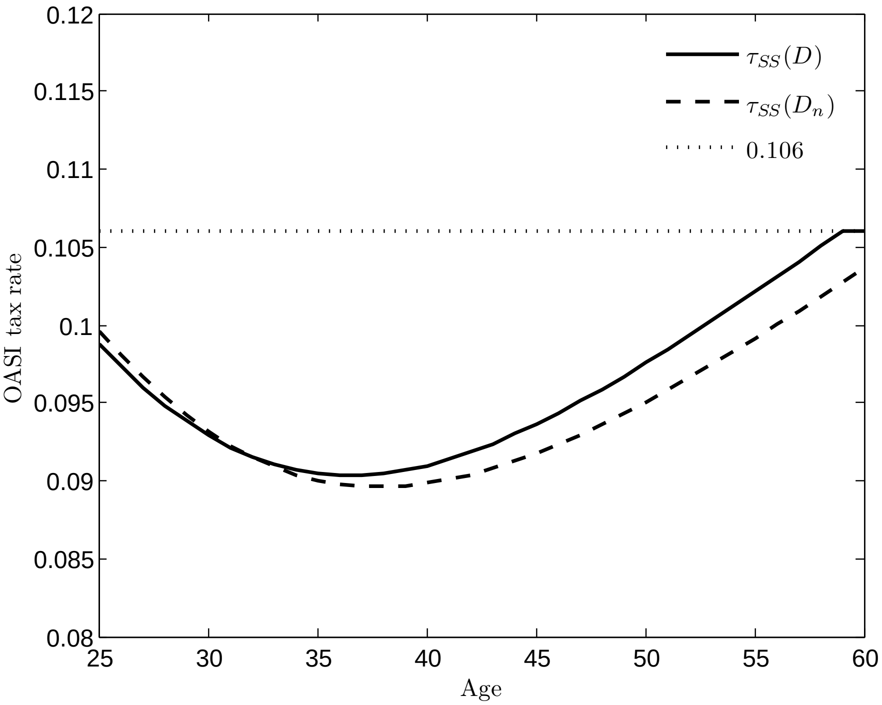

Figure 12: Average Social Security tax rates with the longevity improvement for a household with the most favorable productivity shock, along with those in the baseline calibration.

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 income to the Social Security tax rate of 10.6%, and also counting all income in the Social Security benefit-earnings formula. In the second experiment, I subject all income to the Social Security tax rate of 10.6%, but only count income up to the current cap of 2.47 times the average wage in the benefit-earnings formula. 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 percentage decline in equilibrium Social Security benefits is a convenient 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

$$W = \sum_i f_i \sum_{s=0}^{T} \beta^s Q_n(s) u\left(c_i(t+s,t), l_i(t+s,t)\right), \tag{25}$$

which is nothing but the aggregate ex-ante expected lifetime utility, weighted by the shares of the respective household types in the population. Second, to understand the distributional consequences of these policy changes, I define a consumption equivalence ψ_i for each household type that solves

$$\sum_{s=0}^{T} Q_n(s)\beta^s u\left((1+\psi_i)\, c_i^C(t+s,t), l_i^C(t+s,t)\right) = \sum_{s=0}^{T} Q_n(s)\beta^s u\left(c_i^{NC}(t+s,t), l_i^{NC}(t+s,t)\right), \tag{26}$$

where C denotes current Social Security law with a cap on the amount of earnings subject to the

φ	0.1	0.34	1.0	2.98	10.0	Average
% change	-14.3	-14.3	-17.4	-15.9	-1.13	-16.2

Table 5: Effect of the longevity improvement on equilibrium Social Security benefits, with the cap removed both from the annual tax payments and earnings counted towards benefits.

φ	0.1	0.34	1.0	2.98	10.0
With longevity improvement, cap at 2.47	1.52	1.52	1.30	0.79	0.43
With longevity improvement, no cap	1.52	1.52	1.31	0.80	0.46

Table 6: The ratio of the share of total Social Security benefits received to the share of total taxes paid for each household type.

Social Security tax, and NC denotes a hypothetical Social Security law without the cap. Intuitively, this captures the welfare gains (or losses) for a type$-i$ household in units of consumption. Taken together, these two metrics provide an overall, as well as a disaggregated picture of the welfare consequences of removing the annual cap on Social Security taxes.

6 Quantitative Results

Let us first consider the policy option in which all earnings are subject to the Social Security tax, and they are also counted towards future Social Security benefits. As mentioned before, this policy change preserves the historical link between the caps on taxes and benefits in U.S. Social Security. I report the percentage change in Social Security benefits for each household type for this experiment in Table 5.

It is clear from the table that with this policy change, the fiscal consequences of the longevity improvement on Social Security are virtually indistinguishable from the case when the cap is held fixed at its current level in the U.S. On the average, equilibrium Social Security benefits decline by 16.2% in this experiment, compared to the 16.4% decline when the Social Security tax applies to all earnings up to the current annual cap. Under this experiment, the average tax rate on labor income for the households with the most favorable productivity shock increases by roughly 0.5-1.6 percentage points through their work life. But the tax revenues collected from these households constitute only 8.3% of total Social Security revenues, because of which elimination of the cap leads to only a 0.8% increase in total Social Security revenues.

Because eliminating the cap on taxes in this experiment also implies eliminating the cap on the amount of earnings that can be counted towards Social Security benefits, a large part of the extra revenues generated in this experiment are actually spent in paying benefits to the retirees for whom the cap expires. This can be seen by comparing Table 4 with Table 5, which demonstrates that benefits for the highest-income households now decline by only about 1%, which is significantly smaller than the 18.6% decline with the current annual cap applying to Social Security taxes and benefits.

Removing the cap simultaneously from taxes as well as benefits makes the Social Security program slightly less progressive, relative to when the cap is held fixed at its current level. To demonstrate this, I compare the ratio of the share of total benefits received to the share of total taxes paid for each household type under this policy change, to when the cap is held at its current level in the U.S., in Table 6. It is clear from the table that the fiscal transfers occurring from the high-productivity workers to the low-productivity retirees are weaker under this policy change. This is indicative of a decline in the degree of progressivity in the Social Security program.

	W
With the cap fixed at the current level	-85.91
With the cap removed from both taxes and benefits	-85.93

Table 7: Aggregate ex-ante expected utilities under the longevity improvements, with the cap fixed at its current level, and removing it from both taxes and benefits.

φ	0.1	0.34	1.0	2.98	10.0
%	-0.026	-0.026	-0.018	-0.009	-1.30

Table 8: The consumption equivalence ψ, with the cap removed from the annual tax payments as well as benefits.

φ	0.1	0.34	1.0	2.98	10.0	Average
% change	-13.8	-13.8	-16.8	-15.4	-17.8	-15.8

Table 9: Effect of the longevity improvement on equilibrium Social Security benefits, with the cap removed from the taxes only.

There are two separate channels through which this policy change affects household-level, and also overall welfare. First, given that Social Security in the current model provides partial insurance against mortality risk and an unfavorable productivity shock, this reduction in progressivity has negative consequences on the welfare of low- and medium-income households. Second, removing the cap on Social Security taxes increases the average tax rates for the most productive households through their work life. These higher tax rates impose slightly larger distortions on the consumption, saving, and labor supply decisions of these households, which negatively affect their welfare. Taken together, I find that these two effects have a negative impact on overall welfare (see Table 7).

The distributional consequences of this policy change are shown in Table 8, where I report the value of the consumption equivalence ψ in this experiment for each household type. It is clear from the table that all the household types experience welfare losses from this policy change: the consumption equivalence is always negative. Also, households with the most favorable productivity shock suffer the largest welfare losses, as the higher tax rates impose additional distortions on only these households.

Let us 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 in this case, I break the historical link between the cap on Social Security taxes and benefits in the U.S. In Table 9, I report the percentage change in Social Security benefits due to the longevity improvement for each household type for this experiment. Table 9 shows that on the average, equilibrium Social Security benefits decline by 15.8% in this experiment, which is slightly smaller than the 16.4% decline when the current cap applies to both Social Security taxes and benefits. Moreover, comparing Table 9 with Table 4 shows that this policy change has little effect on how the decline in Social Security benefits is distributed across the different household types.

The findings from Table 9 suggest that even if the annual cap is removed only from the taxes paid, this policy change does not seem to offer significant fiscal advantages for Social Security. Similar to the first experiment, the average tax rate on labor income for the households with the most favorable productivity shock increases by roughly 0.5-1.6 percentage points through their work

φ	0.1	0.34	1.0	2.98	10.0
With longevity improvement, cap at 2.47	1.52	1.52	1.30	0.79	0.43
With longevity improvement, no cap on taxes only	1.53	1.53	1.31	0.8	0.38

Table 10: The ratio of the share of total Social Security benefits received to the share of total taxes paid for each household type.

	W
With longevity improvement, cap at 2.47	-85.91
With longevity improvement, no cap on taxes only	-85.89

Table 11: Aggregate ex-ante expected utilities under the longevity improvements, with the cap fixed at its current level, and removing it from taxes only.

φ	0.1	0.34	1.0	2.98	10.0
%	0.018	0.018	0.021	0.016	-1.654

Table 12: The consumption equivalence ψ, with the cap removed from the annual tax payments only.

life, but these households still constitute 8.3% of the total Social Security revenues. Because of this reason, removing the cap only from taxes leads to only a 0.8% increase in total Social Security revenues. However, because earnings only up to the cap are counted towards Social Security benefits under this experiment, these extra revenues are now spent in paying benefits to all retirees, rather than mostly those for whom the cap expires. As a result, the percentage decline in benefits in this case is slightly smaller than when the cap is held fixed at its current level in the U.S.

The welfare consequences of this policy, however, are quite different from the first experiment. As before, there are two channels through which this policy change affects household welfare, but they now work in opposite directions. On the one hand, removing the cap on earnings subject to the tax while retaining it on the earnings that can be counted towards benefits makes Social Security more progressive. To demonstrate this, I report in Table 10 the ratio of the share of total benefits received to the share of total taxes paid for each household type, with the cap held at its current level in the U.S., and also when it is removed from taxes only. It is clear from the table that this policy change increases the fiscal transfers occurring from the high-productivity workers to the low-productivity retirees through Social Security, which generates welfare gains through the associated insurance effects.

On the other hand, removing the cap on Social Security taxes imposes larger distortions on the consumption, saving, and labor supply decisions of the households with the most favorable productivity shock. Accounting for both of these channels, I find that the welfare gains associated with the insurance effects of Social Security are large enough to more than compensate for the distortionary losses: removing the annual cap on taxes while retaining it on the benefits leads to a slight improvement in overall welfare. I compare in Table 11 the aggregate ex-ante expected utility under the longevity improvement with the cap held at its current level, to that when it is removed from taxes only.

To see how this welfare improvement is distributed, I report the value of the consumption equivalence ψ in this experiment for each household type in Table 12. The table shows that as expected, all but the highest-income households experience a welfare gain from this policy change: the consumption equivalence is positive for all but the households with the most favorable productivity shock. However, it is clear from Table 11 that in the aggregate, the welfare gains to the low- and medium-income households are more than sufficient to compensate for the welfare losses to the high-income households.

It is also worth noting from Table 12 that in this experiment, households with the most favorable productivity shock fare worse relative to the first experiment (see Table 8). This is because of the stronger fiscal transfers occurring from the high-productivity workers to the low-productivity retirees when all earnings are subject to the Social Security tax, but earnings only up to the current cap are counted towards Social Security benefits.

To summarize, the computations suggest that the fiscal advantages of removing the annual cap from the amount of earnings subject to the Social Security tax are relatively small. With the longevity improvement, eliminating the annual cap from Social Security taxes and counting all earnings towards Social Security benefits has almost no effect on Social Security's budget, and this policy change also leads to a decline in overall welfare. Even when the cap is removed only from taxes but retained on the amount of earnings that are counted towards benefits, the longevity improvement leads to only a slightly smaller decline in equilibrium Social Security benefits, and a slight improvement in overall welfare.

7 Conclusions

The amount of earnings that can be annually taxed by Social Security is currently capped at \$106,800. This cap has recently drawn a lot 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 examine quantitatively if the annual cap on taxes can be an effective policy tool in solving Social Security's budgetary problems. To evaluate this question, I use a calibrated general-equilibrium overlapping-generations model with heterogeneous households, mortality risk, and incomplete insurance markets.

In general, the computational results suggest that the fiscal advantages of relaxing the annual cap on Social Security taxes are relatively small. I find that under an empirically reasonable longevity improvement, subjecting all earnings to the Social Security tax and also counting them towards future Social Security benefits has almost no effect on Social Security's budget. This is because most of the extra revenues are spent in paying Social Security benefits to retirees initially subject to the cap, and also because the fraction of total Social Security revenues collected from the households that are initially subject to the cap is quite small. This policy makes Social Security less progressive, and also increases the distortions on the high-productivity workers. Even when the cap is removed only from taxes but retained on the amount of earnings that can be counted towards Social Security benefits, I find that the fiscal advantages are quite small. However, this policy makes Social Security more progressive, which generates welfare gains that more than offset the welfare losses due to the increased distortions on the high-productivity workers.

There is a large literature that has considered modifications in the various institutional features of Social Security to keep the program solvent in the long run, ranging from changes in the payroll tax rate, the eligibility age, to a complete privatization of the existing Pay-As-You-Go structure. The current paper complements this literature by evaluating whether of not the annual cap on Social Security taxes and contributions can play an important role in this debate. The results in this paper suggest that the ability of this particular institutional feature in solving Social Security's long-run budgetary problems may be limited.

References

Altig, D., A. J. Auerbach, L. J. Kotlikoff, K. A. Smetters, and J. Walliser (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. and L. Kotlikoff (1987). *Dynamic Fiscal Policy.* Cambridge University Press.

Bagchi, S. (2012). Is the Social Security Crisis Really as Bad as We Think? *Working Paper*.

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

Caliendo, F. N., N. L. Guo, and R. Hosseini (2013). Social Security is NOT a Substitute for Annuities. *Working Paper*.

Carroll, C. D. (1998). Why do the rich save so much? Working Paper 6549, National Bureau of Economic Research.

Conesa, J. C. and C. Garriga (2008). Optimal Fiscal Policy In The Design Of Social Security Reforms. *International Economic Review 49*(1), 291–318.

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

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

Dushi, I. and A. Webb (2004). Household Annuitization Decisions: Simulations and Empirical Analysis. *Journal of Pension Economics and Finance 3*(2), 109–143.

Feigenbaum, J. (2008a). Can mortality risk explain the consumption hump? *Journal of Macroeconomics 30*(3), 844 – 872.

Feigenbaum, J. (2008b). Information shocks and precautionary saving. *Journal of Economic Dynamics and Control 32*(12), 3917–3938.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the Life Cycle. *Econometrica 70*(1), 47–89.

Gouskova, E., P. Andreski, and R. F. Schoeni (2010). Comparing Estimates of Family Income in the Panel Study of Income Dynamics and the March Current Population Survey, 1968-2007. Technical Series Paper 10-1, Survey Research Center, Institute for Social Research, University of Michigan.

Hansen, G. D. (1993). The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked. *Journal of Applied Econometrics 8*(1), 71–80.

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

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

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

Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics 38*(3), 469–494.

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

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

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

Kitao, S. (2013). Sustainable social security: Four options. *Review of Economic Dynamics*.

Ljungqvist, L. (2008). Lucas critique. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Lockwood, L. M. (2012). Bequest motives and the annuity puzzle. *Review of Economic Dynamics 15*(2), 226–243.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown (1999). New Evidence on the Money's Worth of Individual Annuities. *The American Economic Review 89*(5), 1299–1318.

Nagatani, K. (1972). Life Cycle Saving: Theory and Fact. *The American Economic Review 62*(3), 344–353.

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

Pashchenko, S. (2013). Accounting for non-annuitization. *Journal of Public Economics 98*(0), 53–67.

Reno, V. P. and J. Lavery (2005). Options to Balance Social Security Funds Over the Next 75 Years. Working Paper 18, National Academy of Social Insurance.

Simpson, A. and E. Bowles (2010). The Moment of Truth: Report of the National Commission on Fiscal Responsibility and Reform.

Skinner, J. (1988). Risky income, life cycle consumption, and precautionary savings. *Journal of Monetary Economics 22*(2), 237–255.

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

Turra, C. and O. 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*.

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