

Towson University
Department of Economics
Working Paper Series



Working Paper No. 2014-05

Can Removing the Tax Cap Save Social Security?

by Shantanu Bagchi

May, 2016

© 2014 by Author. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Can Removing the Tax Cap Save Social Security?*

Shantanu Bagchi[†]

May 20, 2016

Abstract

The maximum amount of earnings in a calendar year that can be taxed by U.S. 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 under a realistic longevity improvement, removing this cap leads to Social Security benefits declining by less than 3%, compared to almost 15% 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 also makes Social Security more progressive, which has positive insurance effects for households with unfavorable earnings histories, but the higher marginal tax rates impose larger distortions on households that are no longer subject to the cap, 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

The maximum amount of earnings in a calendar year that can be taxed by Social Security is currently set at \$118,500. 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).

Reforming the taxable maximum of Social Security has also become an important part of the 2016 U.S. presidential election campaigns. Senator Bernie Sanders (I-VT), currently running 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.¹

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.²

This literature, however, has been silent on the question of whether or not the cap on the amount of earnings subject to the the Social Security tax 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 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 potentially undo the initial expansion in the tax base. Therefore, the ability of the tax cap in improving Social Security’s fiscal situation will depend on the quantitative importance of these effects.

In this paper, I quantitatively examine if removing the cap on the amount of taxable earnings can solve Social Security’s budgetary problems. To do this, I construct an overlapping-generations macroeconomic model with incomplete markets, an unfunded public pension system that closely mimics Social Security, and households that experience two types of risk: mortality risk and labor income risk. 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

¹See Sanders (2015) for further details on Senate Bill S.731.

²See, 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).

labor income tax schedule similar to the U.S. Factor markets in the model are competitive, firms maximize profit, and the government provides public goods and Social Security. I calibrate this model to match some key features of the U.S. economy, such as overall capital accumulation, pattern of labor supply over the life cycle (both with respect to labor force participation and hours per week), the earnings distribution relative to the 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 consequences of removing the cap. 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 Social Security taxes and benefits based on current law, the longevity improvement causes the benefits to decline by almost 15% from their baseline level. However, when the cap is removed from the amount of earnings subject to the payroll tax as well as the amount of earnings creditable towards the benefits, I find that the average decline in benefits is only 2.8%. In equilibrium, subjecting all earnings to the payroll tax increases Social Security’s revenues by 15%, 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, which reduces labor supply, capital stock, and output, and reverses almost one-third of the initial expansion in Social Security’s revenues due to the elimination of the cap.

I also find that the fiscal advantages to Social Security are similar 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. As in the previous case, removal of the cap increases Social Security’s revenues by 15% under the longevity improvement, but retaining the cap on the creditable earnings leads to these revenues being spent in paying benefits to *all* 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, 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. When the cap is removed from the amount of earnings subject to the payroll tax and also the amount of earnings creditable towards benefits, Social Security becomes more progressive and has small insurance effects on the households with relatively unfavorable earnings histories. However, the higher marginal tax rates negatively affect labor supply and saving, especially for households with relatively high labor income, and also output. I find that these distortionary losses outweigh the insurance effects, as a result of which overall welfare declines. I find evidence of a small welfare improvement when the cap is removed only from the amount of earnings subject to the tax, relative to when it is removed altogether. Removing the cap on taxes increases the marginal tax rates for households with higher labor income, but retaining it on the amount of creditable earnings makes Social Security considerably more progressive. I find that the associated insurance effects in this

case are stronger and are large enough to offset the distortionary losses, relative to when the cap is removed altogether.

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 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 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 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 incorporate an empirically reasonable longevity improvement into the baseline model in Section 5, and examine its consequences in Section 6. Finally, I describe the quantitative results of the two experiments in Section 7, and I conclude in Section 8.

2 The model

The unit of the current model is a permanent-income household that smooths consumption and labor supply over the life cycle by accumulating a risk-free asset: physical capital. Over the course of the life cycle, this household experiences two types of risk: labor income risk and mortality risk, but does not have access to markets where it can 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

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

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

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

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

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

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%.³

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 a household's 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 or the PIA, is then calculated as a piecewise linear function of the AIME. Finally, the benefits are scaled up or down proportionally so that Social Security's budget is balanced.⁴

³How these accidental bequests are handled within the model has important consequences for its quantitative predictions. A common assumption in the literature is that these accidental bequests are evenly distributed back to the surviving population. However, it has been recently shown that with this assumption, Social Security fails to provide any insurance against mortality risk. Caliendo et al. (2014) demonstrate that if one accounts for how Social Security affects the accidental bequest that households leave (and also receive) in equilibrium, then higher mandatory saving through Social Security crowds out these accidental bequests, and therefore has zero effect on life-cycle wealth. Moreover, with this assumption, the accidental bequests create an additional layer of redistribution in the model that does not exist in reality. Because a higher life expectancy increases saving, it also increases accidental bequests and therefore has a pure income effect on *all* households in equilibrium.

⁴While Social Security has a trust fund and does not satisfy the definition of a Pay-As-You-Go program in the narrow sense in reality, 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), İmrohoroglu 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.

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

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

subject to

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

$$y^{at}(s, \varphi) = h(s)w(s)e(s, \varphi) - T_y(h(s)w(s)e(s, \varphi)) - T_{ss}(h(s)w(s)e(s, \varphi)) \quad (6)$$

$$k' \geq 0 \quad (7)$$

$$AIME' = AIME \quad \forall s \geq T_c \quad (8)$$

$$0 \leq h \leq 1, \quad (9)$$

$$(10)$$

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

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

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

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

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

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

2.6 Aggregation

The population structure in the model is as follows: at each instant a new cohort is born and the oldest cohort dies, and cohort size grows at the rate of n over time. Let us denote the measure of

households at age s with state x as $\mu_s(x)$. Then, the aggregate capital stock and labor supply at any instant t are given by

$$K(t) = \sum_{s=0}^T N(t-s)Q(s) \sum_x k(s+1; x) \mu_s(x) \quad (14)$$

$$L(t) = \sum_{s=0}^T N(t-s)Q(s) \sum_x h(s; x) e(s, x) \mu_s(x), \quad (15)$$

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

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

$$\begin{aligned} \text{Beq}(t) = & (1 + (1 - \tau_k)r) \left[\sum_{s=0}^T N(t-s)Q(s) \sum_x \{k(s+1; x) - k(s; x)\} \mu_s(x) \right] \\ & - n \sum_{s=0}^T N(t-s)Q(s) \sum_x k(s+1; x) \mu_s(x), \end{aligned} \quad (16)$$

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

$$\sum_{s=0}^T N(t-s)Q(s) \sum_x T_{ss} (h(s; x) w(s) e(s, x)) \mu_s(x) = \sum_{s=0}^T N(t-s)Q(s) \sum_x \Theta(s - T_c) b(x) \mu_s(x). \quad (17)$$

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

$$\text{Beq}(t) + \tau_k r K(t) + \sum_{s=0}^T N(t-s)Q(s) \sum_x T_y(y(s; x)) \mu_s(x) = G(t), \quad (18)$$

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

2.7 Competitive equilibrium

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

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

that

1. solves the households' optimization problems,

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

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

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

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

3 Calibration

3.1 Demographics

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

First, I set the payroll tax rate for Social Security to $\tau_{ss} = 0.106$, which is the combined tax rate for the Old-Age and Survivors Insurance (OASI) part of Social Security. This rate is applied to all labor income up to the maximum taxable earnings. The maximum taxable earnings is adjusted annually relative to the average wage in the U.S. For example, the cap 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 \$44,888.⁵ Huggett and Ventura (1999) calculate that the ratio of the maximum taxable earnings to the average wage index has averaged at about 2.47 in the U.S., using which I set the maximum taxable earnings in the model to $\bar{y} = 2.47$.

Second, to compute the Social Security benefit amount (also known as the Primary Insurance Amount or PIA), I incorporate the U.S. benefit-earnings rule into the model. The benefit-earnings rule in the U.S. is a concave (piecewise linear) function of work-life income. The Social Security Administration (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

⁵See <http://www.ssa.gov/oact/cola/awiseries.html> for more details.




Figure 1: Benefit formula in the U.S.

year 2000, the OASI benefit was 90% of the AIME for the first \$531, 32% of the next \$2,671, and 15% of the remaining up to the maximum taxable earnings. As shown by Huggett and Ventura (1999), these dollar amounts come out to be roughly 20%, 124%, and 247% of the average wage in the economy. These percentage amounts are referred to as the “bend points” of the benefit rule, and I take them directly to the model. Note that the progressivity in the benefit 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

$$\log e(s, \varphi) = \epsilon(s) + \varphi, \quad (20)$$

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

$$\varphi_t = p + z_t + \nu_t \quad (21)$$

$$z_t = \rho z_{t-1} + v_t, \quad (22)$$

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




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

I parameterize $\epsilon(s)$ using the estimates from Kitao (2014), who uses work hour and wage data from the 2007 PSID to derive this age-dependent component of productivity as a residual of wages, after accounting for hours worked and also the part-time wage penalty. The resulting $\epsilon(s)$, normalized with respect to productivity at age 25, is plotted in Figure 2.

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

3.4 Taxes

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

$$T_y(y) = y - (1 - \tau_y)y^{1-\tau_1}, \quad (23)$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (23) reduces to a proportional tax function with a marginal rate of τ_y . With this income tax function, after-tax labor income is log-linear in before-tax labor income. To estimate the parameters of this tax function, I take the 2012 tax rate schedule for a single filer in the U.S., compute the after-tax income for each level of before-tax income, and then regress the log of after-tax income on the log of before-tax income. This yields the following estimate for the parameter τ_1 , which controls the progressivity of the tax code:

$$\hat{\tau}_1 = 0.06411. \quad (24)$$




Figure 3: The average tax rates in the model, compared to those from the U.S. tax schedule.

I plot the average tax rates that emerge from the estimated tax function along with those from the U.S. tax schedule in Figure 3. Note that because these are the average rates, they are slightly lower than the marginal tax rates in the U.S. tax schedule. The top marginal tax rate in the U.S. tax schedule is 39.6%, but the top average rate is only around 28% because only a small fraction of income is subject to the top marginal rate.

3.5 Technology

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

3.6 Unobservable parameters

Once all the observable parameters have been assigned empirically reasonable values, I jointly calibrate the remaining unobservable parameters of the model, i.e. the preference parameters σ , β , and η , the age-dependent labor force participation cost θ_P , and also the income tax parameters τ_y and τ_k , to match certain macroeconomic targets.

First, so that overall wealth accumulation in the model matches the U.S. economy, I fix the IES to $\sigma = 4$ and then calibrate the discount factor (β) to get an equilibrium capital-output ratio of 3.0. Second, two salient features of cross-sectional labor supply data in the U.S. are (i) a rapid decline in the labor force participation rate from about 90% to almost 30% between ages 55 to 70, and (ii) an average of 40 hours per week per worker spent on market work between ages 25 to 55. I adopt both of these empirical facts as targets.⁶

⁶The labor force participation and the hours per week per worker targets are based on PSID data as noted in

σ	β	η	κ_1	κ_2	κ_3	$\tau_y (= \tau_k)$
4	0.9616	0.368	6.12×10^{-8}	3.43×10^{-7}	2.98	0.165

Table 1: Unobservable parameter values under the baseline calibration.

	Target	Model
Capital-output ratio	3.0	3.02
Avg. hours of market work per week per worker (25-55)	40	40.5
Share of govt. expenditures in GDP	0.2	0.207
Social Security’s tax base as a fraction total earnings	-	0.82

Table 2: Model performance under the baseline calibration.

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

Finally, I set $\tau_y = \tau_k$ and calibrate them such that the model matches a ratio of government expenditures to GDP of 20% in equilibrium. This step ensures that the scale of tax revenues relative to GDP in the model is consistent with that in the U.S. economy.

4 Baseline results

The unobservable parameter values under which the baseline equilibrium reasonably matches the above targets are reported in Table 1. Note that with leisure in period utility, the relevant inverse elasticity for consumption is $\sigma^c = 1 + \eta(\sigma - 1) = 2.1$, 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 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 Figure 5 that the model does a reasonable job of matching observed labor force participation behavior in the U.S. However, while it replicates the rapid decline in participation after age 50 quite well, the model in general underestimates the participation rates at the later ages. This, in turn, has important consequences for the intensive margin of labor supply, or the hours per week per worker, at these ages. Figure 6 shows that while the model reasonably matches the labor hours (conditional on participation) prior to age 60, it overestimates the hours thereafter. In fact, the rapid decline in work hours does not occur until about age 80 in the model, considerably later than what is seen in the data.

The current model underestimates the old-age labor force participation rates (and overestimates the hours per worker), most likely because it abstracts from the bequest motive: an important determinant of labor supply and saving behavior. Households in the current model smooth consumption

Kitao (2014).




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




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




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

across the work life and retirement (the life-cycle motive), and also across the stochastic realizations of the idiosyncratic productivity shock (the precautionary motive). However, both life-cycle and precautionary motives are less important at later ages, especially because the idiosyncratic productivity shock is highly persistent. Therefore, the absence of the bequest motive causes older households to reduce labor force participation in the model, relative to what is observed in the data. This fact, along with the absence of health risks, also explains why the current model underestimates overall asset holdings, especially 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. From the perspective of a household, whether or not the cap on Social Security taxes binds depends on three key factors: the stochastic labor productivity shock, its implications for the household's life-cycle pattern of labor supply, and finally the interaction of labor supply with the life-cycle endowment profile. Unconditionally, the cap is more likely to bind for households with a favorable productivity shock, and conditional on a particular realization of the shock, the cap is more likely to bind when before-tax labor income is near or at its peak in the life cycle. In Figure 8, I report the fraction of workers subject to 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. Moreover, Social Security's tax base is about 82% of total earnings in the baseline calibration. Based on Social Security's Annual Statistical Supplement, this ratio is currently 83%, so the baseline model matches the U.S. earnings distribution relative to the taxable maximum quite well.⁷

⁷Social 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.

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 risk. However, the equilibrium interest rate in the baseline calibration is 5.6%, which is considerably higher than the population growth rate of 1%. As a result, intergenerational transfers do 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 $h(s)$ based on the following formula:

$$h_n(s) = h(s) - \gamma s^\nu, \quad (25)$$

where γ and ν are positive constants. I set these parameters to $\gamma = 10^{-6}$ and $\nu = 1.8509$, under which the life expectancy in the model is 85 years, matching the 2011 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 due to these longevity improvements, the worker-to-retiree ratio in the U.S. will decline from its current level of roughly 3 to around 2 by 2085, which will make Social Security in its current form 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

	Baseline	Case 0
Labor	36.4	39.7
Avg. hours of market work per week per worker (25-55)	40.5	41.1
Capital	199.2	238.2
Output	66.0	74.3
Interest rate	5.60%	4.89%
Wage	1.18	1.22

Table 3: Model performance under Case 0, compared to that under the baseline calibration.

pay only 77% of scheduled benefits in 2036, and only 74% of scheduled benefits in 2085. Actuaries of the SSA also estimate that increases in the payroll tax rate of the order of two to five percentage points will be sufficient to balance Social Security’s budget in the long run.

Several studies contend that the SSA’s actuarial projections overestimate the extent of the insolvency in Social Security, and also underestimate the tax increases that will be needed to balance Social Security’s budget in the long run. Bagchi (2014) 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 projections overestimate the Social Security crisis. Once these responses are accounted for, Bagchi (2014) shows that longevity improvements lead to a decline in Social Security benefits that is only two-thirds of the SSA’s estimates. Similarly, De Nardi et al. (1999) find that balancing Social Security’s 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 negative impact that these higher taxes will have on the tax base. They find that households respond to the higher Social Security taxes by working and saving less, because of which the tax base shrinks from its current level. Because the current model is an equilibrium model of the economy, it accounts for all of these relevant household-level and macroeconomic adjustments to demographic or policy-induced changes in the economy.

To determine if removing the cap on the amount of taxable earnings 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, but only the survival probabilities are changed to Q_n to reflect the longevity improvement. This experiment will shed some light on the effect of a higher longevity on 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 with all the institutional features of Social Security held at their current level, 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, and they also save more to smooth consumption over a longer expected lifespan. This increases labor supply by 8.9%, capital stock by 19.6%, and output by 12.5% from the baseline.⁸

In Figures 10 and 11, I plot the cross-sectional labor force participation rates and the mean of labor hours per week (conditional on participation) under Case 0, along with those under the

⁸It is important to note here that the longevity improvement actually affects household behavior through two different mechanisms. First, through pure life-cycle motives, a higher life expectancy gives households the incentive to supply more labor and save more. Second, a higher in life expectancy also reduces Social Security benefits, which has a negative income effect and encourages labor supply and saving even more. The above results are the combined effect of these two mechanisms.




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

Permanent productivity shock (p)	0.54	1.00	1.84	Average
% change	-14.6	-14.6	-14.6	-14.6

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

baseline calibration. Both figures show an increase in labor supply over the life cycle, both in terms of participation and also weekly hours, under Case 0. Even though there is a slight decline in participation prior to age 43, delayed retirement (and also increased hours) within older age-cohorts interacts with the life-cycle endowment profile, and ultimately leads to an 8.9% increase in the aggregate labor supply from the baseline level.

The effect of the longevity improvement on Social Security’s fiscal status can be seen in Table 4, where I report how Social Security benefits change under Case 0. The table shows that on the average, benefits decline by 14.6% from their baseline level, which is significantly smaller than the actuarial estimates of the SSA (23-25%).⁹ 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, 2014).

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

⁹See the 2011 Social Security Trustees Report.




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

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 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 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 contributions 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.¹⁰

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 $\tau_{ss} = 10.6\%$, 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 $\tau_{ss} = 10.6\%$, 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 percentage decline in the equilibrium Social Security benefits is a sufficient metric

¹⁰This is the provision included in senate bill S. 731, *The Social Security Expansion Act*, introduced by Senator Bernie Sanders.

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

$$W = \sum_{s=0}^T \beta^s Q_n(s) \sum_x u(c(s; x), 1 - h(s; x), P(s; x)) \mu_s(x) \quad (26)$$

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 shock (p) that solves

$$E \left[\sum_{s=0}^T \beta^s Q_n(s) u((1 + \psi) c^C(s), 1 - h^C(s), P^C(s)) \right] = E \left[\sum_{s=0}^T \beta^s Q_n(s) u(c^{NC}(s), 1 - h^{NC}(s), P^{NC}(s)) \right], \quad (27)$$

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 shock, under each of our two experiments with the cap. Taken together, these two metrics provide an overall, as well as a disaggregated picture of the welfare consequences of removing the cap.

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 the Social Security tax, and they are 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 labor supply, capital stock, and output, relative to the benchmark case when the cap is held fixed at the baseline level (Case 0). This is because the marginal tax rates are now 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 4.4% from the baseline under Case 1, compared to 8.9% under Case 0. Similarly, capital increases by only around 10% from the baseline under Case 1, compared to almost 20% under Case 0. Seen alternatively, removing the cap from Social Security taxes as well as benefits leads to a 4.1% reduction in labor supply, an 8.2% reduction in capital stock, and a 5.6% reduction in output from the current status-quo under the longevity improvement.

I compare in Table 6 how Social Security benefits change for households under Case 1, relative 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 2.8%, compared to almost 15% under Case 0. Moreover, while Social Security benefits decline for *all* households under Case 0, benefits for households with $p = 1.84$ roughly increase by 14% from the baseline under Case 1.

To understand why this experiment leads to a smaller decline in benefits for households with $p = 0.54$ and 1.00, and actually higher benefits for households with $p = 1.84$, first consider the

	Baseline	Case 0	Case 1
Labor	36.4	39.7	38.1
Avg. hours of market work per week per worker (25-55)	40.5	41.1	40.2
Capital	199.2	238.2	218.6
Output	66.0	74.3	70.2
Interest rate	5.60%	4.89%	5.25%
Wage	1.18	1.22	1.20

Table 5: Select macroeconomic variables under Case 1, compared to those under Case 0 and the baseline calibration.

Permanent productivity shock (p)	0.54	1.00	1.84	Average
Case 0 (% change)	-14.6	-14.6	-14.6	-14.6
Case 1 (% change)	-9.5	-5.4	13.9	-2.8

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

Permanent productivity shock (p)	0.54	1.00	1.84
Case 0	0.375	0.327	0.339
Case 1	0.407	0.336	0.299

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

effect of removing the tax cap on Social Security’s overall revenues. Holding the factor prices and labor supply fixed at their 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 15% 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 towards 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 14%, compared to declining by about 15% under Case 0.

Intuitively, the effect of Case 1 on the redistribution implicit in Social Security is less clear, because while eliminating the cap on taxes makes the program more progressive, eliminating the cap on benefits makes it less so. To understand the net effect of this policy change on the implicit redistribution, I compare in Table 7 the average replacement rates under Case 1 to those under Case 0. It is clear from the table that the effect of removing the cap from taxes dominates in this case. Because the replacement rate for benefits is quite low (15%) in this range, counting all earnings towards benefits has only a small effect on the benefits received by households that are no longer subject to the cap. As a result, this policy change makes Social Security more progressive relative to that under Case 0.

There are two mechanisms through which Case 1 affects household-level, and also overall welfare. On the one hand, given that Social Security in the current model provides partial insurance against labor income and mortality risks, the increase in the implicit redistribution (Table 7) can potentially

	W
Case 0	-50.99
Case 1	-51.60

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

Permanent productivity shock (p)	0.54	1.00	1.84
%	-1.03	-1.06	-1.23

Table 9: The consumption equivalence ψ under Case 1.

improve the welfare of households with relatively unfavorable earnings histories. On the other hand, removing the cap on taxes 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. Overall, I find that the impact on welfare is negative (see Table 8).

The fact that eliminating the cap on Social Security taxes and benefits reduces overall welfare should not be surprising. In a general-equilibrium model with endogenous labor, the distortionary effects of Social Security on labor supply are often large enough to outweigh the welfare gains from its insurance effects (Nishiyama and Smetters, 2008; Bagchi, 2015). I report in Table 9 the consumption equivalence (ψ) for each value of the permanent productivity shock under Case 1. The table shows that all households experience welfare losses from this policy change (the consumption equivalence is negative) in spite of the fact that Social Security is more progressive than the baseline under this experiment, although households with $p = 0.54$ and 1.00 experience smaller welfare losses.

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, Social Security benefits need to decline by less than 3% on the average under the longevity improvement, compared to almost 15% when the benefits continue to be based on current law. This policy change also increases the redistribution implicit in Social Security, but the distortionary losses to the households no longer subject to the cap are sufficiently large to yield an overall welfare reduction.

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 output decline, and approximately by the same percentage relative to the current status-quo. Therefore, eliminating the cap from the amount of earnings subject to Social Security taxes but retaining it on the amount of earnings creditable towards benefits has the same effect on macroeconomic aggregates as does eliminating the cap altogether.

In Table 11, I report the change in Social Security benefits from the baseline for each value of the permanent productivity shock under this experiment. The table shows that overall, Social Security benefits decline by the same percentage under Case 2: 2.8% on the average. This suggests

	Baseline	Case 0	Case 1	Case 2
Labor	36.4	39.7	38.1	38.1
Avg. hours of market work per week per worker (25-55)	40.5	41.1	40.2	40.3
Capital	199.2	238.2	218.6	218.7
Output	66.0	74.3	70.2	70.3
Interest rate	5.60%	4.89%	5.25%	5.23%
Wage	1.18	1.22	1.20	1.20

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

Permanent productivity shock (p)	0.54	1.00	1.84	Average
Case 0 (% change)	-14.6	-14.6	-14.6	-14.6
Case 1 (% change)	-9.5	-5.4	13.9	-2.8
Case 2 (% change)	-2.8	-2.8	-2.8	-2.8

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

Permanent productivity shock (p)	0.54	1.00	1.84
Case 0	0.375	0.327	0.339
Case 1	0.407	0.336	0.299
Case 2	0.439	0.382	0.391

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

	W
Case 0	-50.99
Case 1	-51.60
Case 2	-51.55

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

when the cap is removed only from taxes, the fiscal advantages to Social Security are very similar to when it is removed both from taxes and benefits. However, the table also shows that unlike Case 1, benefits decline for *all* households under Case 2, including those with relatively favorable earnings histories. Because earnings only up to the cap are counted towards the benefits under this experiment, the extra Social Security revenues are now spent in paying benefits to *all* retirees, rather than mostly to those for whom the cap expires.

The welfare consequences of this policy change are also slightly different from the first experiment. As before, there are two mechanisms through which this policy change affects household welfare. 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 (see Table 12). On the other hand, removing the cap on Social Security taxes imposes larger distortions on the labor supply and saving decisions of the households with higher labor income. Accounting for both of these mechanisms, 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. when both the cap on taxes as well as benefits are held fixed at the current level (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 higher marginal tax rates.

Permanent productivity shock (p)	0.54	1.00	1.84
Case 1 (%)	-1.03	-1.06	-1.23
Case 2 (%)	-0.86	-0.98	-1.49

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

I report in Table 14 the distributional consequences of Case 2 as the consumption equivalence (ψ) for each realization of the permanent productivity shock. The table shows smaller welfare losses for households with $p = 0.54$ and 1.00 under Case 2, which 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.

To summarize, my computations suggest that the fiscal advantages of removing the cap only from the amount of earnings subject to the Social Security tax are similar to when the cap is removed both from taxes as well as the amount of earnings creditable towards benefits. As before, the longevity improvement causes a significantly smaller decline in Social Security benefits (2.8%), relative to the status-quo (14.6%). However, in this case benefits decline for *all* retirees, as opposed to only retirees with relatively unfavorable earnings histories. The effect of this experiment on labor supply, capital stock, and output is also roughly identical to when the cap is removed altogether.

8 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 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 this 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 mortality and labor income risk, and incomplete insurance markets.

In general, the 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 under an empirically reasonable longevity improvement, subjecting all earnings to the payroll tax and also counting them towards future Social Security benefits causes the benefits to decline by less than 3% on the average, compared to almost 15% when the cap is held fixed at its current level. This policy change increases Social Security’s revenues by only 15%, because households that are no longer subject to the cap respond by reducing their labor supply and saving, and some of these extra revenues are spent in paying Social Security benefits to wealthy retirees for whom the cap expires. I find that the fiscal advantages to Social Security are roughly similar, when the cap is removed only from taxes but retained on the amount of earnings creditable towards Social Security benefits. 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 to the low- and medium-income households are smaller when the cap is retained on the amount of earnings creditable towards Social Security benefits.

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, 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 or not the cap on the amount

of earnings subject to Social Security taxes and contributions 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., 2014. Is the Social Security Crisis Really as Bad as We Think? *Macroeconomic Dynamics* forthcoming.
- Bagchi, S., 2015. Labor Supply and the Optimality of Social Security. *Journal of Economic Dynamics and Control* forthcoming.
- 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., İmrohoroglu, 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.
- Huang, H., İmrohoroglu, 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.
- İmrohoroglu, A., İmrohoroglu, S., Joines, D. H., 2003. Time-Inconsistent Preferences and Social Security. *The Quarterly Journal of Economics* 118 (2), 745–784.
- Jeske, K., 2003. Pension systems and aggregate shocks. *Federal Reserve Bank of Atlanta Economic Review* (Q1), 15–31.
- Karabarbounis, M., 2012. Heterogeneity in labor supply elasticity and optimal taxation. In: 2012 Meeting Papers. No. 655. Society for Economic Dynamics.

- 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.
- Reno, V. P., Lavery, J., 2005. Options to Balance Social Security Funds Over the Next 75 Years. Working Paper 18, National Academy of Social Insurance.
- Sanders, B., 2015. Social Security Expansion Act. Legislative proposal, United States Senate.
- Simpson, A., Bowles, E., 2010. The Moment of Truth: Report of the National Commission on Fiscal Responsibility and Reform.
- Smetters, K., 2003. Is the social security trust fund worth anything? Working Paper 9845, National Bureau of Economic Research.
- Stokey, N. L., Rebelo, S., 1995. Growth Effects of Flat-Rate Taxes. *Journal of Political Economy* 103 (3), 519–550.
- Storesletten, K., Violante, G., Heathcote, J., 2012. Redistributive Taxation in a Partial Insurance Economy. 2012 Meeting Papers 588, Society for Economic Dynamics.
- Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59 (2), 371–396.
- Zhao, K., 2014. Social security and the rise in health spending. *Journal of Monetary Economics* 64 (0), 21–37.