Towson University
Department of Economics
**Working Paper Series**

Working Paper No. 2026-04

# Borrowing Constraints and their Implications for Social Security

By Shantanu Bagchi

March 2026

# Borrowing Constraints and their Implications for Social Security*

Shantanu Bagchi†

February 13, 2026

**Abstract**

This paper uses a stylized overlapping-generations model to examine the effect of borrowing constraints on the economic implications of Social Security. In this framework, Social Security provides partial insurance against income risk that is uninsured due to incomplete markets. I find that when borrowing consistent with life cycle behavior is allowed in this framework, the micro- and macroeconomic effects of a downsizing in Social Security are considerably smaller than when borrowing is prohibited. I also find that the key mechanism behind this result is labor supply: with endogenous borrowing, households are able to exploit increasing labor productivity in early life to better self-insure against income risk.

JEL Classifications: E21, G51, H55

Keywords: Social Security; borrowing constraint; incomplete markets; partial insurance; labor supply

*This project received financial support from the Towson University College of Business and Economics Summer Research Grant Program.

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

# 1 Introduction

Economists generally think of unfunded public pension systems as underwriting implicit financial contracts that are typically not offered in private insurance markets. These systems partially insure older individuals against risks that markets do not insure well – risks that may be demographic or economic in nature. For example, unfunded pensions establish an implicit contract between current workers and retirees, which takes advantage of population growth as a mechanism to share demographic risk across generations. Most public pension programs in the industrialized world, including Social Security in the United States (U.S.), also provide benefits that are a progressive function of earnings. The curvature of this benefit function determines the extent to which these pensions provide insurance against unfavorable labor-market events in early life, such as the inability to secure a high-paying job, or unemployment due to an economic recession. In this sense, these pensions establish an implicit contract between current workers (or future retirees) with varying labor-market outcomes.[1]

Since the early 1980's, a large literature has evolved that examines quantitatively how Social Security's implicit financial contracts affect household-level and aggregate economic outcomes. The broad approach in this literature has been to set up consistent macroeconomic models with households, firms, and a government as its building blocks, and then to carefully incorporate appropriate market structures within which these building blocks interact. One of the key features of the market structures typically considered in this literature is "incomplete markets", which roughly amounts to a lack of financial instruments to hedge against all possible future risks, leading to constrained and suboptimal allocations.

One specific way incomplete markets have been modeled in this literature is through idiosyncratic earnings risk and hard no-borrowing constraints. In other words, households are assumed to face future labor income risk, but they are not allowed to hedge against them through borrowing, so the only way for them to "self-insure" is to engage in "precautionary" saving. This mechanism is a crucial determinant of the household-level and aggregate economic implications of Social Security's implicit financial contracts. In an incomplete–markets setup, Social Security's mandatory savings reduce households' disposable income, which interferes with the precautionary saving motive. Not surprisingly, this literature largely finds that once this mechanism is accounted for, Social Security's net effects on household welfare and the overall economy are negative.

The goal of this project is to examine the quantitative importance of hard no-borrowing constraints in a macroeconomic and welfare analysis of unfunded public pensions – U.S. Social Security in particular. Specifically, my objective is to compare the effects of an identical change in the size of Social Security across two alternative models: a canonical one with hard no-borrowing constraints, and a hypothetical one in which the hard no-borrowing constraint is relaxed to allow borrowing consistent with life-cycle permanent income behavior. A comparison of the effects from the two models will help illustrate the importance of the assumption.

I begin by constructing a stylized heterogeneous–agent overlapping–generations model that consists of rational utility-maximizing households, profit-maximizing firms, incomplete markets, and a government that provides public goods and Social Security. Households in the model experience permanent and idiosyncratic earnings risk, and they also face a hard exogenous no-borrowing constraint that affects their ability to self insure. I calibrate this model to match U.S. macroeconomic targets, and then compute the micro- and macroeconomic effects of a 50% downsizing in the size of Social Security. Next, I modify the stylized model by allowing borrowing consistent with life-cycle

[1]These pensions also function as imperfect annuities, which can be viewed as financial contracts between an individual's current and future "selves".

behavior, and then perform an identical computational experiment on this modified model. Finally, I compare the results across the two models to tease out the effect of the borrowing constraint.

I find that when borrowing is allowed in this model framework, the micro- and macroeconomic effects of a downsizing in Social Security are considerably smaller than when borrowing is prohibited. I also find that the key mechanism behind this result is labor supply: with endogenous borrowing, households are able to exploit increasing labor productivity in early life to better self-insure against income risk. I demonstrate this by performing the same computational experiment on two alternative versions of the original model with fixed labor supply: one with a hard exogenous no-borrowing constraint, and another with borrowing consistent with life-cycle behavior. A comparison of the results shows that the micro- and macroeconomic effects of the Social Security cut are quite similar across the two alternative models, confirming the importance of the household-level labor supply decision.

The rest of the paper is organized as follows. I review the related literature in Section 2, describe the stylized model in Section 3, and I detail the calibration and results in Sections 4 and 5. I explain how I relax the exogenous hard no-borrowing constraint in Section 6, and I report the primary quantitative results in Section 7. Finally, I examine the importance of labor supply in Section 8, and I conclude in Section 9.

## 2 Literature review

This question will contribute to the following strands of the literature. In one of the earliest papers on this topic, Hubbard and Judd (1987) examine the impact of Social Security on national saving and individual welfare with realistic capital market imperfections – missing private annuity markets and restrictions on borrowing against anticipated future wages. They find that in this framework, Social Security increases lifetime welfare and reduces national saving if borrowing restrictions are absent. However, the increase in individual welfare is reduced, and in some cases eliminated, when borrowing constraints are taken into consideration. An important drawback of Hubbard and Judd (1987), however, is the fact that their model framework does not account for idiosyncratic earnings risk. This likely means that they underestimate households' precautionary saving motives, which is an important determinant of Social Security's welfare effects.

In a more recent study, Rojas and Urrutia (2008) consider this question using a framework that does a more careful modeling of idiosyncratic risk within an incomplete market environment. Specifically, they incorporate financial intermediaries that account for the fact that from a rational borrower's perspective, the value of paying back debt is no less than the value of defaulting, which is associated with some level of exclusion from financial markets in the future. As a result, relaxing borrowing constraints in their model does not lead to default in equilibrium. Using this model, Rojas and Urrutia (2008) find that with exogenous borrowing constraints, the elimination of social security increases the capital to output ratio and the savings rate by 27%, but with endogenous borrowing, the increase is smaller – about 25%. In other words, they find that relaxing the borrowing constraint leads to a weakening of the mechanisms that cause social security to reduce welfare in a calibrated general equilibrium environment. However, in their analysis, Rojas and Urrutia (2008) ignore a key margin of household behavior – labor supply – that could further weaken those mechanisms.

Finally, two papers that account for labor supply along with endogenous borrowing constraints in a welfare analysis of Social Security are Nishiyama and Smetters (2007) and (2008). Both papers consider different Social Security reform scenarios with incomplete markets and endogenous borrowing constraints during the work life, and find that the long-run welfare effects of downsizing or eliminating Social Security are actually negative in the presence of uninsurable wage shocks.

And they also find that this result is robust with respect to alternative redistribution schemes, existence of annuity markets, and open/closed economy assumptions. Overall, in their models, Nishiyama and Smetters (2007) and (2008) find completely different welfare effects from Social Security reform. However, because their focus is on the alternative policy experiments, it is not clear exactly how much of their results are being driven by endogenous vs exogenous borrowing constraints. The goal of my project is to address this gap in the literature.

# 3 The model

The unit of the model is a life-cycle permanent-income household that smooths consumption and labor supply by accumulating a risk-free asset: physical capital. This household experiences two types of risk over the course of the life cycle: labor income risk (both long- and short-term), and mortality risk. Both types of risk and uninsurable in the initial, i.e. households do not have access to markets to privately insure against these risks.

In the initial setup, households face a hard exogenous no-borrowing constraint. In other words, they are not allowed to hold negative assets, and this constraint is later relaxed when I quantitatively explore its importance. However, it is worth noting that even when the exogenous no-borrowing constraint is relaxed, life-cycle behavior itself places an endogenous constraint on households' borrowing. Specifically, the highest borrowing allowed in this scenario is equal in magnitude to the present value of the lowest future labor income stream at the maximum possible work hours.

At each date, a surviving household earns labor income (conditional on labor supply), and also receives Social Security benefits conditional on surviving past 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.

## 3.1 Preferences

Households derive utility both from consumption and leisure, and also an end-of-life "warm glow" bequest. A household's labor supply decision at each instant consists of two components: the extensive margin or the participation decision ($P$), and the intensive margin or the hours of work ($h$), conditional on participation. The period utility function is given by

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

where $\eta$ is the share of consumption, $\sigma$ is the inverse of the intertemporal elasticity of substitution (IES), $\theta_P$ is the age-dependent cost of labor force participation (measured in unit of time), and $P$ is the labor force participation status: $P = 1$ if $h > 0$, and $P = 0$ otherwise. Also, I normalize the period time endowment to unity, i.e. $0 \leq h \leq 1$. Finally, the end-of-life "warm glow" bequest utility function is given by

$$v\left(B_{WG}\right) = \begin{cases} \zeta \, \frac{B_{WG}^{1-\sigma}}{1-\sigma} & \text{if} \quad \sigma \neq 1 \\ \zeta \, \ln\left(B_{WG}\right) & \text{if} \quad \sigma = 1 \end{cases} \tag{2}$$

where $\zeta$ controls the strength of the bequest motive. Expected lifetime utility from the perspective of a newborn household is given by

$$U = E\left[\sum_{t=0}^{T} \beta^t Q_t \, u\left(c_t, 1-h_t, P_t\right) + \beta^T \, v\left(B_{WG}\right)\right], \tag{3}$$

where $\beta$ is the discount factor, and $Q_t$ is the unconditional probability of surviving up to age $t$.

## 3.2 Income

Conditional on labor force participation, a household earns before-tax income of $y_t = h_t w e_t(\varphi) + (R-1)k_t$ at age $t$, where $w$ is the wage rate, $e_t(\varphi)$ is a labor productivity endowment that depends on age and a stochastic productivity shock $\varphi$, $k_t$ is asset holdings, and $R$ is the risk-free gross interest rate. This income is subject to two separate taxes: a progressive income tax $T_y(\cdot)$ that applies to both capital and labor income, and a payroll tax for Social Security $T_{ss}(\cdot)$ that applies only to labor income. After-tax income at age $t$ is therefore given by

$$y_t^{at}(\varphi) = y_t(\varphi) - T_y(y_t(\varphi)) - T_{ss}\left(y_t(\varphi) - (R-1)k_t\right). \tag{4}$$

It is worth noting here that because of the uninsurable nature of the mortality risk, deceased households at every age leave behind their assets as accidental bequests. These accidental bequests are distinct from the end-of-life "warm glow" bequests and they do not provide any utility. 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%. Also, when the hard exogenous no-borrowing constraint is relaxed, any borrowing by a household at age $t$ must accrue interest at a gross rate of $R/Q_t$ to account for the fact that some households will die with outstanding debt.[2]

## 3.3 Social Security

Households receive Social Security benefits after the full retirement age ($T_r$), and the amount of benefits paid to a particular household depends on its earnings history. For each household, the government calculates an average of past earnings (up to the maximum taxable earnings), referred to as the Average Indexed Monthly Earnings (AIME). The replacement rate is then calculated as a piecewise linear function of the AIME. Finally, the government scales benefits up or down proportionally so that Social Security's budget is balanced.[3]

[2] How these accidental bequests are handled within the model has important consequences for its quantitative predictions. A common assumption in the literature is that these accidental bequests are evenly distributed back to the surviving population. However, Caliendo et al. (2013) 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.

[3] While in reality, Social Security has a trust fund and does not satisfy the definition of a Pay-As-You-Go program in the narrow sense, it is a common practice in the literature to ignore the trust fund and model Social Security's budget as balanced every period (See, for example, studies such as Huggett and Ventura (1999), Conesa and Krueger (1999), İmrohoroğlu et al. (2003), Jeske (2003), Conesa and Garriga (2009), and Zhao (2014), among others). This is due to disagreement on whether or not the trust fund assets are "real", i.e. whether or not they have increased national saving. In fact, Smetters (2003) finds that the trust funds assets have actually increased the level of debt held by the public, or reduced national saving.

### 3.4 A household's optimization problem

Let us define the state vector of an age$-t$ household as $x = \{k, \varphi, AIME\}$, where $k$ denotes the beginning-of-period assets, $\varphi$ the stochastic productivity shock, and AIME the earnings history measure that determines Social Security benefits. Conditional on a particular realization of the state, the household chooses consumption, assets holdings for the next period, and labor supply. Note that next-period assets at age $T$ (i.e. the last period) effectively become the end-of-life "warm glow" bequests.

At a given age $t$, this optimization problem can be recursively represented as

$$V_t(x) = \max_{c,k',h} \left\{ u(c, 1-h, P) + \beta \frac{Q_{t+1}}{Q_t} \; E\left[V_{t+1}(x')\right] \right\} \tag{5}$$

subject to

$$c + k' = Rk + y_t^{at}(\varphi) + \Theta(t - T_r)\, b(AIME) \tag{6}$$

$$y_t^{at}(\varphi) = h_t w e_t(\varphi) + (R-1)k_t - T_y\left(h_t w e_t(\varphi) + (R-1)k_t\right) - T_{ss}\left(h_t w e_t(\varphi)\right) \tag{7}$$

$$0 \leq h_t \leq 1 \tag{8}$$

$$k' \geq 0 \tag{9}$$

$$AIME' = \begin{cases} \left[AIME \times (t-1) + \min\left\{h_t w e_t(\varphi), \bar{y}\right\}\right]/t & t < T_r \\ AIME & t \geq T_r \end{cases}, \tag{10}$$

where

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

is a step function, and $\bar{y}$ is the maximum earnings counted toward Social Security. Households are born with zero assets and die with "warm-glow" bequests $B_{WG}$, i.e. $k_0 = 0$ and $k_{T+1} = B_{WG}$. Prior to age $T_r$, the earnings history measure AIME evolves based on the realized labor productivity shocks and the endogenous labor supply decisions. Also, note that the exogenous no-borrowing constraint (9) is relaxed in a later version of the model.

### 3.5 Technology and factor prices

Firms produce output using a Cobb-Douglas production function with inputs capital and labor

$$Y \;=\; K^{\alpha} L^{1-\alpha}, \tag{11}$$

where $\alpha$ is the share of capital in total income. Factor markets are perfectly competitive, which implies that factor prices are equal to their net marginal products

$$R = MP_K - \delta + 1 = \alpha \left[\frac{K}{L}\right]^{\alpha-1} - \delta + 1 \tag{12}$$

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

where $\delta$ is the depreciation rate of physical capital.

### 3.6 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 $t$ with state $x$ as $\mu_t(x)$. Then, the aggregate capital stock and labor supply in the steady state are given by

$$K' = \sum_{t=0}^{T} N_t Q_t \int k_{t+1}(x) d\mu_t(x) \tag{14}$$

$$L = \sum_{s=0}^{T} N_t Q_t \int h_t(x) e_t(x) d\mu_t(x), \tag{15}$$

where $N_t = \frac{N_0}{(1+n)^t}$ is the size of the age-$t$ cohort.

The budget-balancing condition for Social Security is given by

$$\sum_{t=0}^{T} N_t Q_t \int T_{ss}\left(h_t(x) w e_t(x)\right) d\mu_t(x) = \sum_{t=0}^{T} N_t Q_t \int \Theta(t - T_r) b(x) d\mu_t(x). \tag{16}$$

Finally, the government adjusts its level of public goods expenditures ($G$) to match the total tax revenues from labor and capital income, and also the accidental bequests ($B_{ACC}$)

$$B_{ACC} + \sum_{t=0}^{T} N_t Q_t \int T_y\left(y_t(x)\right) d\mu_t(x) = G. \tag{17}$$

### 3.7 Competitive equilibrium

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

1. cross-sectional consumption allocations $\{c_t(x)\}_{t=0}^{T}$ and labor supply allocations $\{h_t(x)\}_{t=0}^{T}$,
2. an aggregate capital stock $K$ and labor $L$,
3. a gross rate of return $R$ and a wage rate $w$,
4. Social Security benefits $b(x)$ and government expenditures $G$, and
5. a measure of households $\mu_t(x)\ \forall\ t$,

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 $\mu_{t+1}(x) = R_\mu\left[\mu_t(x)\right]$, where $R_\mu(\cdot)$ is a one-period transition operator on the measure distribution.

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

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

I consider only steady-state equilibria of this model, i.e. when all aggregate quantities grow at the rate of population growth and all per-capita quantities are constant. I also normalize the initial newborn cohort size to $N_0 = 1$.

# 4 Calibration

## 4.1 Demographics

I assume that households enter the model at the actual age of 25, which corresponds to the model age of zero. I use the average age-specific death rates from the 2001 U.S. Life Tables in Arias (2004) to calibrate the model survival probabilities. 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.

## 4.2 Social Security

To calibrate Social Security in the model, I first set the payroll tax rate to $\tau_{SS} = 0.106$, which is the combined tax rate for the Old-Age and Survivors Insurance (OASI) component.[4] The Social Security Administration (SSA) adjusts the maximum taxable earnings ($\bar{y}$) regularly relative to the average wage in the U.S. For example, the taxable maximum was set at \$76,200 in the year 2000, but was adjusted to \$106,800 in 2010 and \$113,700 in 2013. During the same period, the national average wage index increased from \$32,155 to \$41,674, and finally to \$44,888.[5] Huggett and Ventura (1999) calculate that the ratio of this taxable maximum to the average wage index has averaged at about 2.47 in the U.S. I use this estimate to set the maximum taxable earnings in the model to $\bar{y} = 2.47$.

Second, to compute the Social Security benefit amount (also known as the PIA), I incorporate the U.S. benefit-earnings rule to calculate the replacement rates. The benefit-earnings rule in the U.S. is a concave (piecewise linear) function of past work-life income, the AIME. The SSA calculates the AIME, and then it calculates the replacement rate 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 replacement rate. For example, in the year 2000, the OASI benefit was 90% of the AIME for the first \$531, 32% of the next \$2,671, and 15% of the remaining up to the maximum taxable earnings. As shown by Huggett and Ventura (1999), these dollar amounts come out to be roughly 20%, 124%, and 247% of the average wage in the economy. It is worth noting that the progressivity in the benefit-earnings rule is captured by the fact that the "replacement rate" is decreasing in the AIME (see Figure 1).

Finally, I assume that households in the model receive Social Security benefits starting at age $T_r = 41$, which corresponds to the current full retirement age of 66 in the U.S.

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

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

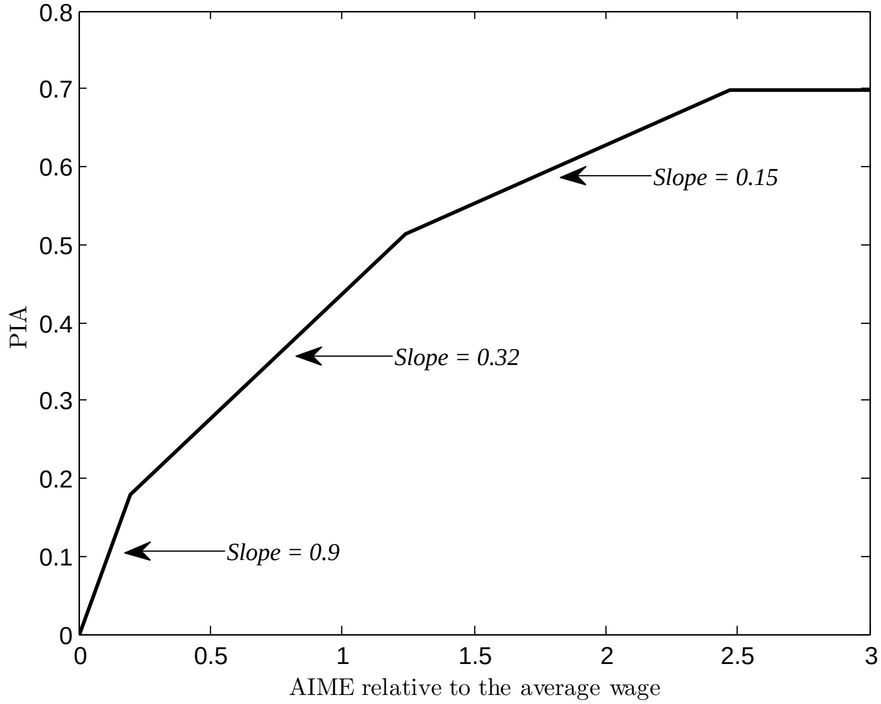


**Figure 1:** Benefit formula in the U.S.

### 4.3 Labor productivity

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

$$\log e_t(\varphi) = \epsilon_t + \varphi_t, \tag{19}$$

where $\epsilon_t$ is a deterministic age-dependent component, and $\varphi_t$ is a stochastic component, which is further decomposed as

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

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

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

with $z_0 = 0$, persistence $\rho$, and a white-noise disturbance $\upsilon_t \sim N(0, \sigma_\upsilon^2)$.

I parameterize $\epsilon_t$ 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 $\epsilon_t$, normalized with respect to productivity at age 25, is plotted in Figure 2.

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

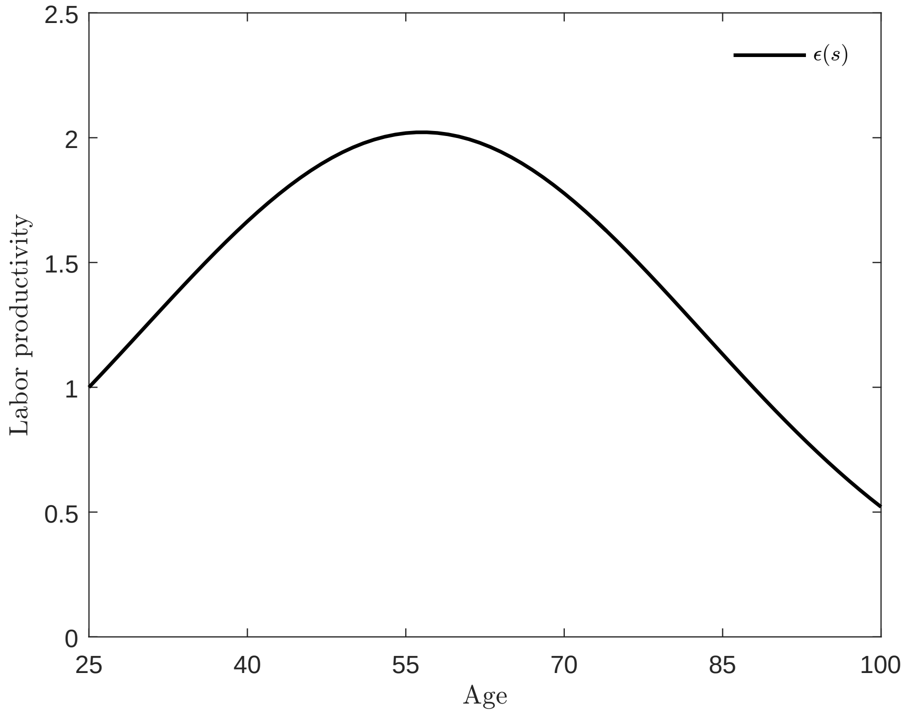


**Figure 2:** The age-dependent component of labor productivity from Kitao (2014).

### 4.4 Taxes

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

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

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (22) reduces to a proportional tax function with a marginal rate of $\tau_y$. With this income tax function, after-tax labor income is log-linear in before-tax labor income, and the parameter $\tau_1$ controls the progressivity of the tax code. Following Heathcote et al. (2014), I set the value of this parameter to $\tau_1 = 0.151$. Heathcote et al. (2014) estimate this value using data from the 2000, 2002, 2004, and 2006 waves of the PSID, and also NBER's TAXSIM program, accounting for federal and state income taxes plus public transfers.

### 4.5 Technology

The historically observed value of capital's share in total income in U.S. ranges between 30-40%, so I set $\alpha = 0.35$. Also, following Stokey and Rebelo (1995), I set the depreciation rate to $\delta = 0.08$.

### 4.6 Structural parameters

I jointly calibrate the remaining unobservable structural parameters of the model, i.e. the preference parameters $\sigma$, $\beta$, $\eta$, $\zeta$, the age-dependent labor force participation cost $\theta_P$, and also the labor income tax parameter $\tau_y$, to match key macroeconomic targets.

First, so that overall wealth accumulation in the model matches the U.S. economy, I fix the IES to $\sigma = 2.5$ and then calibrate the discount factor ($\beta$) to get an equilibrium capital-output ratio

| $\sigma$ | $\beta$ | $\eta$ | $\zeta$ | $\kappa_1$ | $\kappa_2$ | $\kappa_3$ | $\tau_y$ |
|---|---|---|---|---|---|---|---|
| 2.5 | 0.975 | 0.35 | $10^4$ | $2.31\times10^{-3}$ | 0.3059 | 3.75 | 0.8 |

**Table 1:** Structural parameter values under the baseline calibration.

| | Target | Model 1 |
|---|---|---|
| Capital-output ratio | 3.0 | 2.99 |
| Avg. hours of market work per week per worker (25-55) | 30-40 | 35.8 |
| Share of govt. expenditures in GDP | 0.2 | 0.197 |
| Annual Social Security benefit per retiree | – | \$36,907 |

**Table 2:** Model performance under the baseline calibration.

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] Third, the parameter $\zeta$ controls the strength of the "warm-glow" bequest motive, and I follow De Nardi (2004) and French (2005) to calibrate it to match average asset holdings between ages 70 and 90.

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

$$\theta_{P,t} = \kappa_1 + \kappa_2 t^{\kappa_3},$$

where $t$ is model age, and then parameterize $\kappa_1$, $\kappa_2$, and $\kappa_3$ to match the observed rapid decline in labor force participation after age 55. The consumption share parameter ($\eta$) controls the fraction of time a household spends on market work (conditional on participation), so I calibrate it to match the hours per week target.

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

# 5 Baseline steady-state equilibrium

The structural parameter values under which the baseline steady-state equilibrium reasonably matches the above targets are reported in Table 1. With the above values of $\kappa_1$, $\kappa_2$, and $\kappa_3$, the labor force participation cost is a positive and increasing function of age, which causes the labor force participation rate to be a decreasing function of age with a discrete drop at age 65 (the full retirement age). I report the equilibrium values of relevant macro aggregates in Table 2. The life cycle profiles of labor, consumption, and asset holdings are reported in Figures 3 and 4 respectively.

It is clear from Figure 3 that the model does a reasonable job of matching observed labor supply behavior in the U.S. It replicates the rapid decline in participation after age 50 (top panel), and it also reasonably matches the general declining trend of weekly hours over the life cycle (bottom panel). However, the model fails to replicate the mild hump-shape in the hours profile, and it also

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

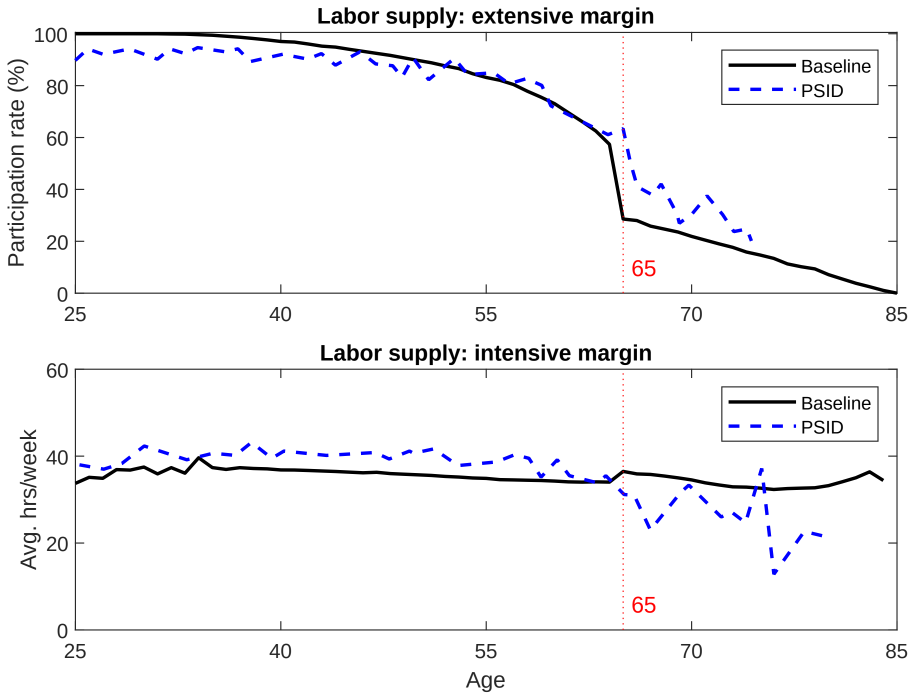


**Figure 3:** Baseline labor supply: participation and hours.

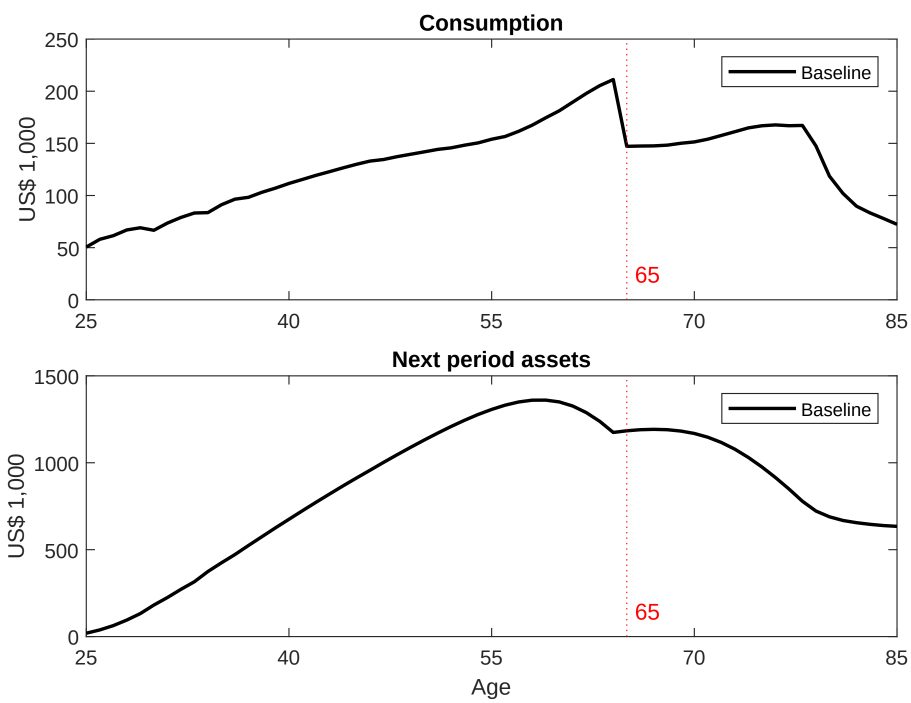


**Figure 4:** Baseline consumption and assets.

| | Target | Model 1 | Model 2 |
|---|---|---|---|
| Capital-output ratio | 3.0 | 2.99 | 2.47 |
| Avg. hours of market work per week per worker (25-55) | 30-40 | 35.8 | 45.3 |
| Share of govt. expenditures in GDP | 0.2 | 0.197 | 0.188 |
| Annual Social Security benefit per retiree | – | \$36,907 | \$34,264 |

**Table 3:** Model performance under the baseline calibration.

somewhat overestimates both participation and weekly hours at later ages.[7]. Also, Figure 4 shows that the model does a good job of matching the non-monotonicity of the life-cycle consumption profile (top panel), and also the old-age asset holdings in the data (bottom panel). This should not be surprising, because the current model accounts for a major driver of old-age savings: "warm glow" bequest motives (De Nardi et al., 2010).

# 6 An endogenous no-borrowing constraint

Next, I proceed to computing a hypothetical version of the baseline model in which the hard no-borrowing constraint is relaxed to allow borrowing consistent with life-cycle permanent income behavior. In this model, the highest borrowing allowed at any age is equal in magnitude to the present value of the lowest future labor income stream at the maximum possible work hours. Notably, because households exit the labor market at late ages, the present value of the lowest future labor income stream in late life is equal to zero, but is strictly positive in early life.

I compare the macro aggregates from this model (Model 2) with those from the original exogenous no-borrowing model (Model 1) in Table 3, and the life cycle profiles in Figures 5 and 6. It is clear from the table that relaxing the exogenous hard no-borrowing constraint leads to a lower capital-output ratio, increased work hours, and a small decline in the average Social Security benefit level. The effect on labor supply is further evident in Figure 5, which shows that participation declines slightly (top panel), but the workers who continue to participate work more hours, especially earlier in life (bottom panel). On the other hand, both consumption and asset holdings are lower until retirement in Model 2, but are higher post-retirement (see Figure 6). In other words, borrowing consistent with life-cycle behavior causes early-life indebtedness (bottom panel), but allows a higher level of consumption late in the life cycle (top panel).

# 7 The computational experiment

The quantitative importance of the hard no-borrowing constraint can be determined by comparing the effects of an identical change in the size of Social Security across two alternative models: a canonical one with hard no-borrowing constraints (Model 1), and a hypothetical one in which the hard no-borrowing constraint is relaxed to allow borrowing consistent with life-cycle permanent income behavior (Model 2). I define this experiment to be a relatively straightforward one: a roughly 50% cut in the overall size of Social Security. However, two additional details about the computational experiments are relevant. First, because of the associated general equilibrium effects, the payroll tax rate cut needed to accomplish this downsizing varies across the two models. Second,

[7] One potential strategy to improve the model's fit along this dimension is to treat the age-dependent component of labor productivity $\epsilon_t$ as an unobservable structural parameter. Treating $\epsilon_t$ as an unobservable parameter would eliminate any selection bias arising from measuring it as residual wages (Bullard and Feigenbaum, 2007; Bagchi and Feigenbaum, 2014)

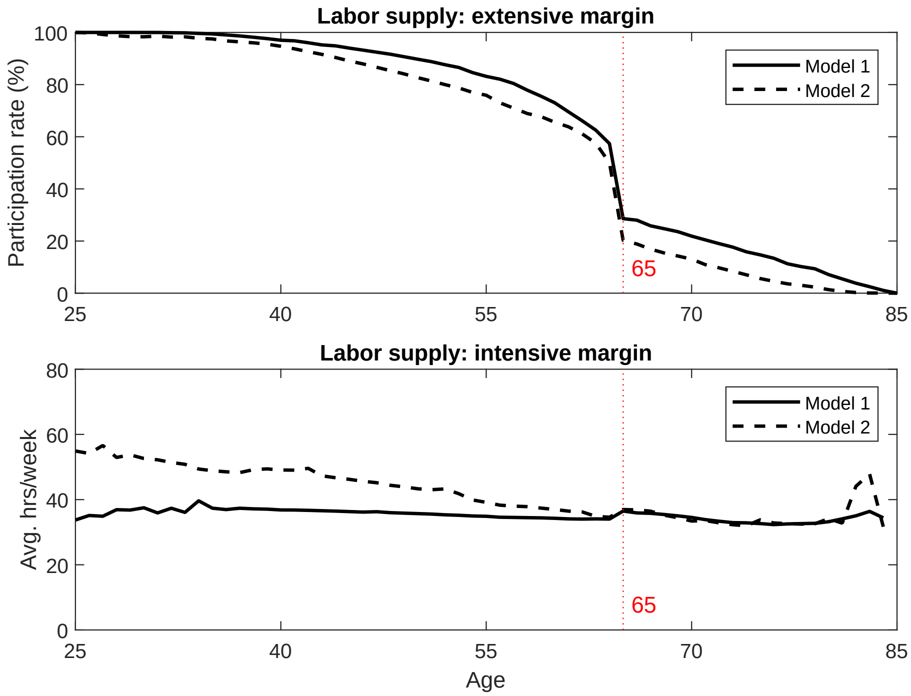


**Figure 5:** Labor supply in Models 1 and 2.

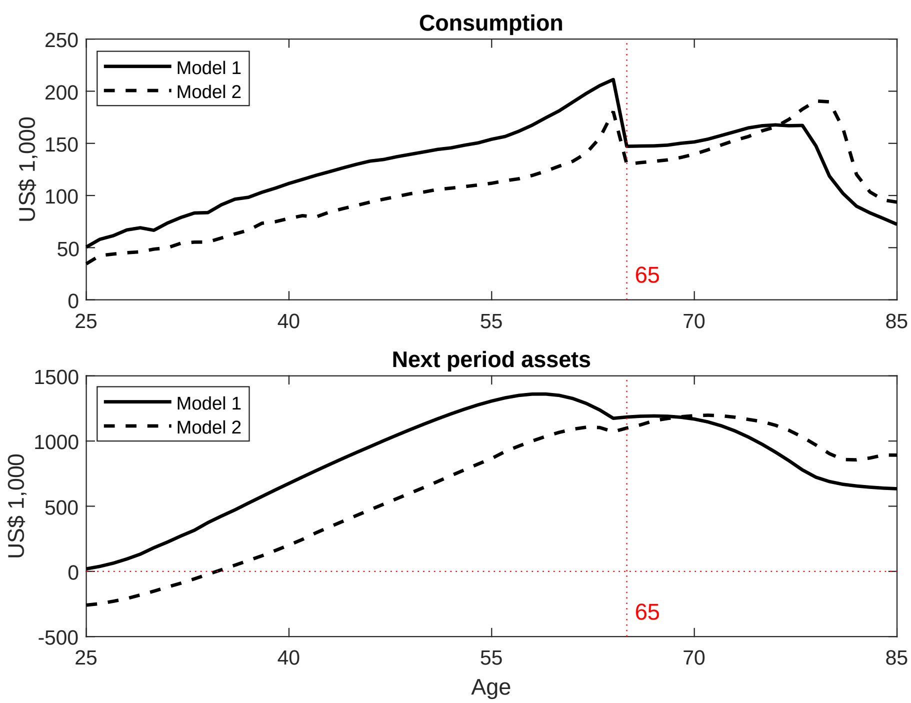


**Figure 6:** Consumption and assets in Models 1 and 2.

| | Baseline | Experiment | % Change |
|---|---|---|---|
| Capital ($K$) | 213.42 | 237.62 | 11.3 |
| Labor ($L$) | 39.50 | 40.63 | 2.9 |
| GDP | 71.22 | 75.43 | 5.9 |
| Consumption ($C$) | 33.18 | 35.30 | 6.4 |
| Government ($G$) | 17.9 | 17.9 | 0.0 |
| Social Security ($SS$) | 1.92 | 0.97 | −49.5 |
| Avg. Social Security benefit (US$1,000) | 36.91 | 15.49 | −58.0 |
| Avg. end-of-life assets (US$1,000) | 630.56 | 495.67 | −21.4 |

**Table 4:** Macroeconomic effects of the experiment in Model 1.

while downsizing Social Security, I adjust the income taxes to maintain revenue neutrality for the government's non-Social Security public goods expenditures.[8] For each model, I compute a new steady state under the experiment, and then compare the results across the models.

The steady-state values of key macroeconomic variables for Model 1 under this experiment are compared to the baseline in Table 4. It is clear from the table that downsizing Social Security in the presence of a hard no-borrowing constraint leads to an increase in both capital stock (11.3%) and labor supply (2.9%). Together, these lead to an increase of 5.9% and 6.4% in GDP and aggregate consumption respectively. Overall, these results are in line with the rest of the literature (see Kitao (2014), Bagchi (2016), and other papers cited therein), which finds generally positive macroeconomic effects of downsizing Social Security in a dynamically efficient economy with rational households.

These macroeconomic effects are driven by the underlying household-level responses to the tax cut. I compare the cross-sectional means of labor supply, consumption, and asset holdings under the experiment with those in the Model 1 baseline in Figures 7 and 8. The following facts are clear from the figures. First, in Model 1, a 50% cut in the size of Social Security in equilibrium increases labor force participation, especially at later ages. However, average labor hours among those who work remain roughly unchanged. In other words, the extensive margin appears to respond more to the tax cut. Second, average consumption and asset holdings both increase, with the largest increases in consumption accruing after age 65, i.e. the point at which households begin to receive Social Security benefits. The increased labor supply in response to the tax cut, therefore, is large enough to deliver higher old-age consumption despite a more than 50% cut in average Social Security benefits.

Next, I perform the same computational experiment on Model 2, i.e. when borrowing consistent with life-cycle behavior is allowed. I report the macro effects of this tax cut in Table 5, and I compare the cross-sectional means of labor supply, consumption, and asset holdings under the experiment with those in the Model 2 baseline in Figures 9 and 10. It is clear from Table 5 that similar to Model 1, the tax cut increases capital and labor in Model 2, but the effects are significantly smaller. The increase in capital stock is in the neighborhood of 6% in this scenario, compared to 11.3% in Model 1. Similarly, labor supply increases by 2% in Model 2, compared to by nearly 3% in Model 1. Together, these lead to a 3.5% increase in GDP in this model, compared to a nearly 6% increase in Model 1. In other words, relaxing the exogenous hard no-borrowing constraint significantly dampens the macroeconomic effects of a comparable downsizing in Social Security.

Figures 9 and 10 demonstrate the underlying household-level responses to the tax cut in Model 2. Similar to Model 1, labor supply responds more along the extensive margin (i.e. labor force participation), and the changes in consumption is considerably smaller. Most notably, the tax cut

[8] This is equivalent to assuming that the household's utility function is additively separable in public goods.

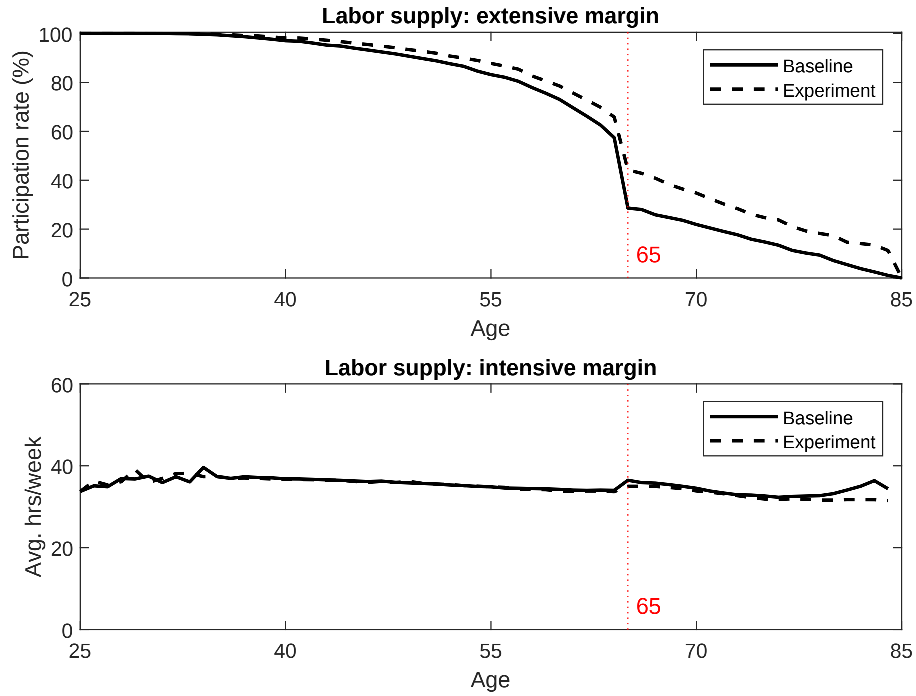


**Figure 7:** Labor supply in Model 1 in the baseline and after the tax cut.

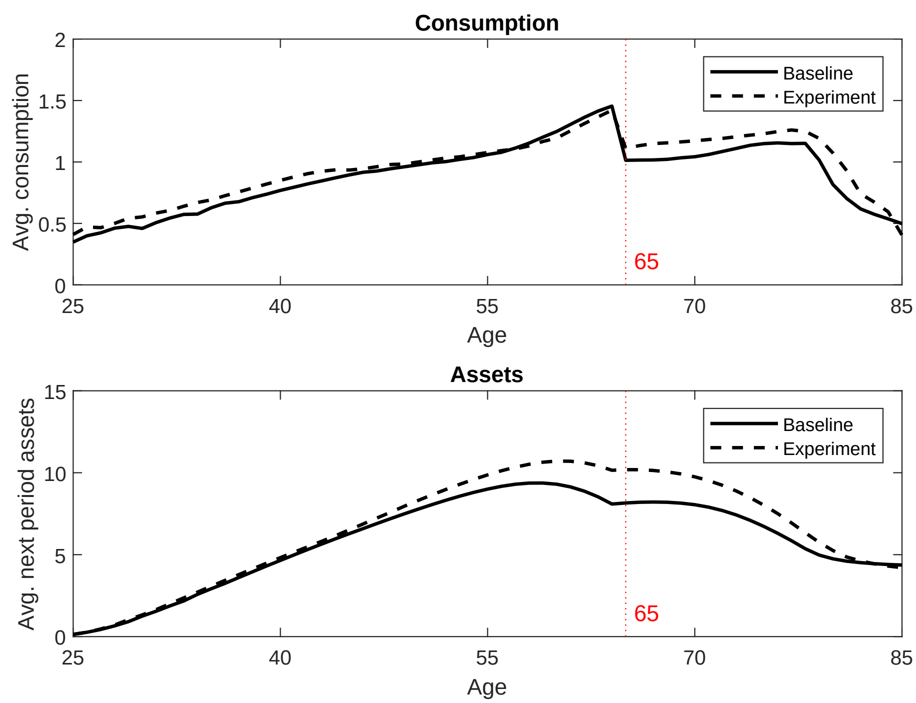


**Figure 8:** Consumption and assets in Model 1: baseline and after the tax cut.

| | Baseline | Experiment | % change |
|---|---|---|---|
| Capital ($K$) | 163.12 | 173.12 | 6.1 |
| Labor ($L$) | 40.69 | 41.52 | 2.0 |
| GDP | 66.16 | 68.45 | 3.46 |
| Consumption ($C$) | 31.66 | 33.00 | 4.3 |
| Government ($G$) | 17.6 | 17.6 | 0.0 |
| Social Security ($SS$) | 2.20 | 1.10 | −49.8 |
| Avg. Social Security benefit (US$1,000) | 34.26 | 15.81 | −53.9 |
| Avg. end-of-life assets (US$1,000) | 931.44 | 813.04 | −12.7 |

**Table 5:** Macroeconomic effects of the experiment in Model 2.

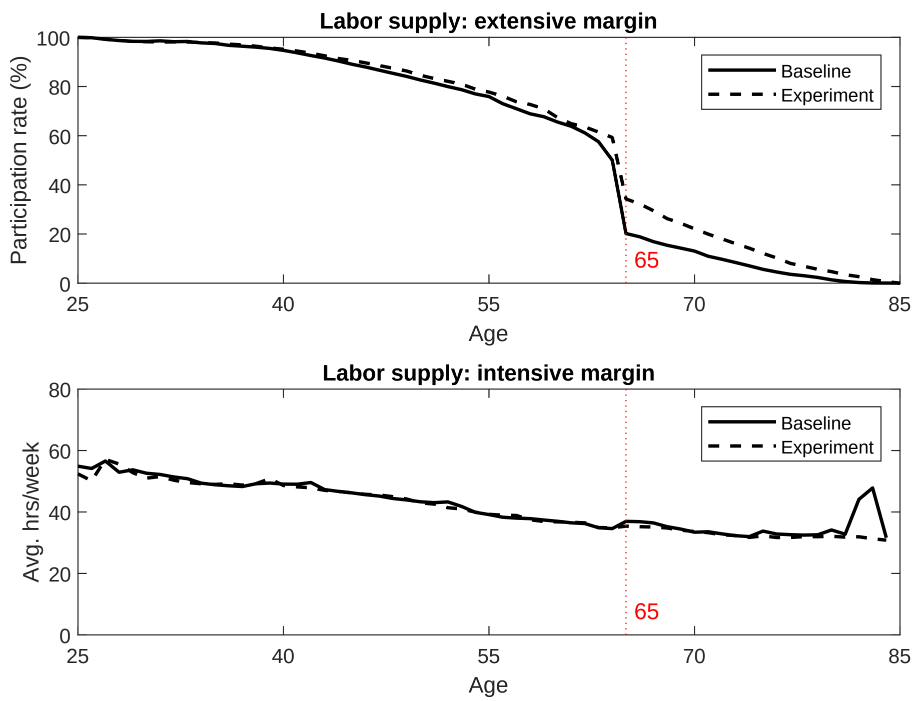


**Figure 9:** Labor supply in Model 1 in the baseline and after the tax cut.

leads to higher borrowing in early life compared to the baseline – a possibility only in Model 2. This increased borrowing plays a key role in weakening the effect of the Social Security downsizing on aggregate capital stock in Model 2.

Given the absence of the exogenous hard no-borrowing constraint in Model 2, the highest borrowing allowed at any age is equal in magnitude to the present value of the lowest future labor income stream at the maximum possible work hours. With the hump-shaped pattern of labor productivity over the life-cycle, this endogenous borrowing limit increases in early to mid-life. This allows households in Model 2 to use future labor supply as a stronger precautionary tool against earnings uncertainty. In Model 1, where borrowing is prohibited, this mechanism is absent, so labor supply is a less potent tool for households looking to self-insure against idiosyncratic risk. This is the key mechanism that explains the divergence of results across Models 1 and 2.

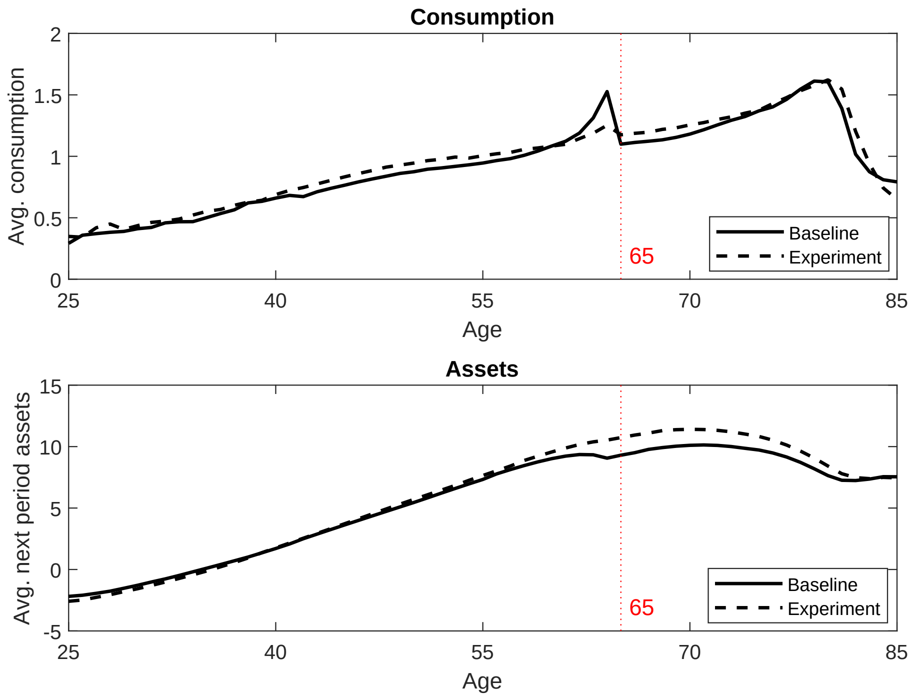


**Figure 10:** Consumption and assets in Model 2: baseline and after the tax cut.

| | Baseline | Experiment | % change |
|---|---|---|---|
| Capital ($K$) | 137.04 | 164.88 | 20.3 |
| Labor ($L$) | 30.45 | 30.45 | 0.00 |
| GDP | 51.48 | 55.07 | 7.0 |
| Consumption ($C$) | 27.56 | 28.48 | 3.3 |
| Government ($G$) | 10.6 | 10.6 | 0.0 |
| Social Security ($SS$) | 2.28 | 1.16 | −49.3 |
| Avg. Social Security benefit (US$1,000) | 35.52 | 16.59 | −53.3 |
| Avg. end-of-life assets (US$1,000) | 400.77 | 330.39 | −17.6 |

**Table 6:** Macroeconomic effects of the experiment in Model 3.

## 8 The importance of labor supply

A precise way to examine the importance of the labor supply mechanism in the context of borrowing constraints is to consider alternative versions of Models 1 and 2 in which labor supply is inelastic. Specifically, I define two alternative models – Models 3 and 4 – as follows. In Model 3, labor supply is exogenously fixed at the average equilibrium age-specific levels in Model 1, and borrowing is prohibited via the exogenous hard no-borrowing constraint. On the other hand, in Model 4, labor supply is exogenously fixed at the average equilibrium age-specific levels in Model 2, but the hard no-borrowing constraint is relaxed to allow borrowing consistent with life cycle behavior. Essentially, these two models are best understood as "inelastic labor" counterparts of Models 1 and 2 respectively. Then, I compute the effect of a comparable downsizing of Social Security across Models 3 and 4.

First, I report the macroeconomic effects of a 50% cut in the size of Social Security in Model 3 in Table 6. It is clear from the table that with exogenous labor, the tax cut leads to a considerably larger increase in capital stock: 20.3% in Model 3, compared to 11.3% in Model 1. This, in turn,

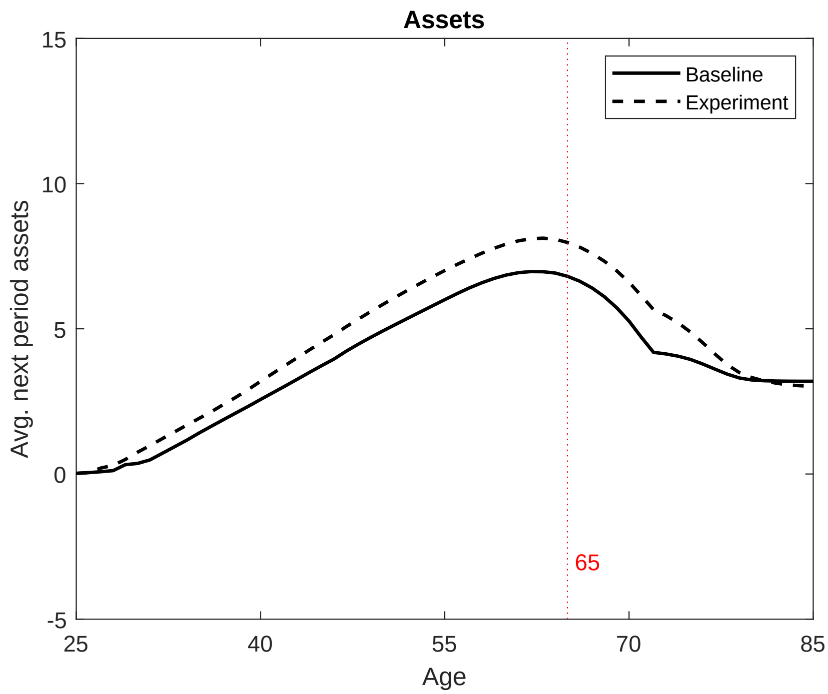


**Figure 11:** Assets in Model 3: baseline and after the tax cut.

| | Baseline | Experiment | % change |
|---|---|---|---|
| Capital ($K$) | 151.12 | 183.81 | 21.6 |
| Labor ($L$) | 37.03 | 37.03 | 0.00 |
| GDP | 60.59 | 65.04 | 7.3 |
| Consumption ($C$) | 32.14 | 32.86 | 2.3 |
| Government ($G$) | 14.25 | 14.2 | 0.0 |
| Social Security ($SS$) | 2.61 | 1.32 | $-49.37$ |
| Avg. Social Security benefit (US$1,000) | 34.24 | 16.03 | $-53.17$ |
| Avg. end-of-life assets (US$1,000) | 613.02 | 411.90 | $-32.81$ |

**Table 7:** Macroeconomic effects of the experiment in Model 4.

leads to a 7% increase in GDP in Model 3, compared to a 6% increase in Model 1. This should not be surprising, because with the exogenous hard no-borrowing constraint in Model 3, households can self-insure against idiosyncratic risk only via saving more. This increased saving (see Figure 11) has a stronger macroeconomic effect in Model 3.

On the other hand, in Model 4, where the exogenous hard no-borrowing constraint is relaxed to allow borrowing consistent with life-cycle behavior, a 50% cut in Social Security yields the macroeconomic effects reported in Table 7. Clearly, the tax cut in this case has an effect on capital stock that is very similar to that observed in Model 3, i.e. when the exogenous hard no-borrowing constraint is present. In Model 4, capital stock increases by 21.6%, which is only slightly larger than the 20.3% increase in Model 3. This yields only a slightly larger increase in GDP in this case: 7.3% in Model 4, compared to 7% in Model 3. At the household level, the tax cut in Model 4 leads to slightly larger borrowing in early life and lower late life savings, which are collectively responsible for this (see Figure 12).

To summarize, eliminating labor supply as a precautionary tool in Models 3 and 4 leads to

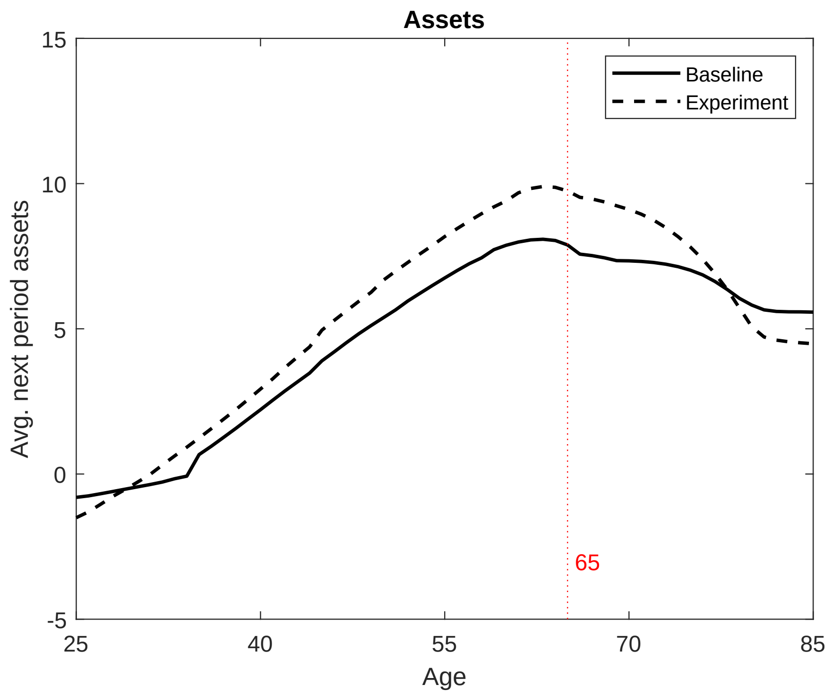


**Figure 12:** Assets in Model 4: baseline and after the tax cut.

both models responding very similarly to an identical downsizing in Social Security, despite their differences in underlying household-level borrowing behavior.

# 9 Conclusions

In this project, I examine the quantitative importance of hard no-borrowing constraints in a macroeconomic and welfare analysis of unfunded public pensions – U.S. Social Security in particular. Specifically, my objective is to compare the effects of an identical change in the size of Social Security across two alternative models: a canonical one with hard no-borrowing constraints, and a hypothetical one in which the hard no-borrowing constraint is relaxed to allow borrowing consistent with life-cycle permanent income behavior. A comparison of the effects from the two models will help illustrate the importance of the assumption.

I begin by constructing a stylized heterogeneous–agent overlapping–generations model that consists of rational utility-maximizing households, profit-maximizing firms, incomplete markets, and a government that provides public goods and Social Security. Households in the model experience permanent and idiosyncratic earnings risk, and they also face a hard exogenous no-borrowing constraint that affects their ability to self insure. I calibrate this model to match U.S. macroeconomic targets, and then compute the micro- and macroeconomic effects of a 50% downsizing in the size of Social Security. Next, I modify the stylized model by allowing borrowing consistent with life-cycle behavior, and then perform an identical computational experiment on this modified model. Finally, I compare the results across the two models to tease out the effect of the borrowing constraint.

I find that when borrowing is allowed in this model framework, the micro- and macroeconomic effects of a downsizing in Social Security are considerably smaller than when borrowing is prohibited. I also find that the key mechanism behind this result is labor supply: with endogenous borrowing, households are able to exploit increasing labor productivity in early life to better self-insure against income risk. I demonstrate this by performing the same computational experiment on two alternative versions of the original model with fixed labor supply: one with a hard exogenous no-borrowing constraint, and another with borrowing consistent with life-cycle behavior. A comparison of the results shows that the micro- and macroeconomic effects of the Social Security cut are quite similar across the two alternative models, confirming the importance of the household-level labor supply decision.

## References


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

Bagchi, Shantanu (2016), "Is the social security crisis really as bad as we think?" *Macroeconomic Dynamics*, 20, 737–776.

Bagchi, Shantanu and James Feigenbaum (2014), "Is Smoking a Fiscal Good?" *Review of Economic Dynamics*, 17, 170–190.

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

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

Conesa, Juan C. and D. Krueger (1999), "Social Security Reform with Heterogeneous Agents." *Review of Economic Dynamics*, 2, 757–795.

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

De Nardi, Mariacristina (2004), "Wealth inequality and intergenerational links." *The Review of Economic Studies*, 71, 743–768, URL `https://doi.org/10.1111/j.1467-937X.2004.00302.x`.

De Nardi, Mariacristina, Eric French, and John B. Jones (2010), "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy*, 118, 39–75.

French, Eric (2005), "The effects of health, wealth, and wages on labour supply and retirement behaviour." *The Review of Economic Studies*, 72, 395–427, URL `https://doi.org/10.1111/j.1467-937X.2005.00337.x`.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2010), "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy*, 118, 681–722.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2014), "Optimal tax progressivity: An analytical framework." Working Paper 19899, National Bureau of Economic Research.

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

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

İmrohoroğlu, Ayse, Selahattin İmrohoroğlu, and Douglas H. Joines (2003), "Time-Inconsistent Preferences and Social Security." *The Quarterly Journal of Economics*, 118, 745–784.

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

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

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

Nishiyama, Shinichi and Kent Smetters (2007), "Does Social Security Privatization Produce Efficiency Gains?" *The Quarterly Journal of Economics*, 122, 1677–1719.

Nishiyama, Shinichi and Kent Smetters (2008), "The Optimal Design of Social Security Benefits." Working Papers wp197, University of Michigan, Michigan Retirement Research Center.

Rojas, Juan A. and Carlos Urrutia (2008), "Social security reform with uninsurable income risk and endogenous borrowing constraints." *Review of Economic Dynamics*, 11, 83–103, URL `https://www.sciencedirect.com/science/article/pii/S1094202507000257`.

Smetters, Kent (2003), "Is the social security trust fund worth anything?" Working Paper 9845, National Bureau of Economic Research.

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

Tauchen, George and Robert Hussey (1991), "Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models." *Econometrica*, 59, 371–396.

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