Towson University
Department of Economics
Working Paper Series

Working Paper No. 2023-01

Means Testing and Social Security in the U.S.

By Shantanu Bagchi

March 2023

Means Testing and Social Security in the U.S.*

Shantanu Bagchi†

March 11, 2023

Abstract

This paper uses a heterogeneous-agent overlapping-generations model to examine the fiscal and distributional consequences of introducing a means test in U.S. Social Security. I find that a means test, i.e. conditioning benefit payments on a household's earnings and/or assets, leads to a higher implicit tax on old-age resources, but has desirable distributional effects. A 75% cut in the benefits to households with earnings more than 200% of the median leads to a 2.3% reduction in the overall size of Social Security, but has almost no effect on the average benefit level. A fiscally comparable payroll tax cut, on the other hand, leads to an across-the-board decline of 2% decline in the benefits. Finally, an asset-based means test causes a decline of 1% in average benefits, but has a large negative effect on the accidental bequests left behind by deceased households.

JEL Classifications: E21, E62, H55

Keywords: Social Security; means test; mortality risk; labor income risk; incomplete markets; accidental bequests

*I would like to thank Frank Caliendo, Matt Chambers, Jim Feigenbaum, Aspen Gorry, Juergen Jung, and Kai (Jackie) Zhao 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

With the Republican Party winning a narrow majority in the House of Representatives in the 2022 U.S. Midterm Elections, policy conversations about making Federal entitlement programs such as Social Security and Medicare more fiscally sustainable have regained media attention.[1] With respect to Social Security in particular, lawmakers on both sides of the aisle have recently offered proposals to either strengthen the revenue base of the program (such as Sen. Bernie Sanders (I-VT)), or to downsize its cost liability (such as Sen. Rick Scott (R-FL)).[2] One of the proposals in this mix calls for "means-testing" Social Security benefits, or conditioning benefit payments on a household's old-age resources, such as earnings and/or assets.

The first clearly articulated proposal for means-testing Social Security benefits likely came from former New Jersey Governor Chris Christie. During his 2015 presidential campaign, Christie proposed to reduce benefit entitlements for top earners by calling for reduced benefits for people earning in excess of $80,000 per year, and by eliminating benefit payments entirely for individuals making more than $200,000 per year. His campaign argued that this means-testing, coupled with an increase in the early as well as full retirement age, will reduce lifetime benefits payouts by about 10%. Even Donald Trump, during his first presidential campaign, indirectly suggested that he (and implying that other rich people like him) should forgo Social Security benefits in order to keep the program sustainable.[3]

The academic literature on Social Security reform has also examined a range of these proposals, starting from tax and retirement age increases (Kitao, 2014), to a partial or complete privatization of the program (Jeske, 2003). Surprisingly, proposals to explicitly means-test Social Security benefits based on households' old-age resources, such as earnings and/or assets, have not received much attention in academic research. This study aims to fill this gap in the literature.

In this paper, I examine if introducing a means test in U.S. Social Security can be an effective policy to make the program fiscally sustainable while retaining its core mission of preventing the likelihood of old-age poverty. To do this, I construct an overlapping-generations macroeconomic model with incomplete markets, an unfunded public pension system that mimics Social Security, and heterogeneous households that experience realistic mortality and labor income risks. Social Security provides partial insurance against these risks, because households do not have access to private insurance markets. Households in the model also face a progressive income tax schedule, factor markets are competitive, firms maximize profit, and the government provides public goods and Social Security. I calibrate this model to match key features of the U.S. economy, such as overall capital accumulation, pattern of labor supply over the life cycle, earnings and Social Security benefits distributions, and also the share of government expenditures in GDP.

Using this model, I compute the macroeconomic and distributional consequences of a means test, in which benefits to households with higher-than-average old-age resources are reduced. In the computations, I allow for all the household-level and macroeconomic adjustments to this policy change. I find that an earnings-based means test, i.e. conditioning benefit payments on a household's earnings, leads to a higher implicit tax on old-age resources, but has desirable distributional effects. A 75% cut in the benefits to households with earnings more than 200% of the median leads

[1] See, for example, recent articles from the NPR (https://www.npr.org/2023/02/11/1156132516/republicans-say-they-wont-cut-social-security-so-why-does-it-keep-coming-up) and ABC News (https://abcnews.go.com/Politics/us-holds-31-trillion-debt-shrink/story?id=97251512).

[2] For further details on their proposals, see https://www.washingtonpost.com/us-policy/2023/02/23/biden-white-house-social-security/ from the Washington Post and https://thehill.com/homenews/senate/3862794-under-fire-rick-scott-changes-plan-to-exempt-social-security-medicare-from-sunsetting/ from The Hill.

[3] For more details, see the Motley Fool article https://www.fool.com/retirement/2018/08/16/what-is-means-testing-and-how-could-it-affect-soci.aspx.

to a 2.3% reduction in the overall size of the program, but has almost no effect on the average benefit level. This is despite the fact that capital and labor decline slightly due this higher implicit tax. A fiscally comparable payroll tax cut, on the other hand, leads to an across-the-board decline of 2% decline in the benefits, despite an increase in capital and labor. Finally, a comparable means test that conditions benefit payments on asset holdings rather than earnings yields a 1% decline the average benefit level, but has a large negative effect on the accidental bequests left behind by deceased households. This is because the assets-based means test is an implicit tax on capital, which reduces old-age asset holdings, and therefore the accidental bequests.

To be clear, what I consider in the current paper is an "explicit" means test, i.e. an explicit conditioning of benefit payments on a household's earnings and/or assets. This is somewhat different from the progressivity implied by Social Security's concave benefit-earnings rule, which is arguably an "indirect" means test. Overall, my findings suggest that an "explicit" means test, administered through an earnings-based or an assets-based test, provides an effective strategy to put Social Security on a more sustainable fiscal path without sacrificing the program's core progressive mission of providing insurance against the likelihood of old-age poverty.

The remainder of the paper is organized as follows: Section 2 introduces the model, and Section 3 describes the baseline calibration and its results. I describe the main experiment of this paper and its results in Section 5, and then discuss the alternative experiments in Sections 6 and 7. Finally, I conclude in Section 8.

2 The model

The unit of the current 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, i.e. households do not have access to markets to privately insure against these risks.

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.

2.1 Preferences

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

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

where 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 unit of time), 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 newborn household is given by

$$U = E\left[\sum_{t=0}^{T} \beta^t Q_t \, u\left(c_t, 1 - h_t, P_t\right)\right], \tag{2}$$

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

2.2 Income

Conditional on labor force participation, a household earns before-tax wage income of $y_t = h_t w e_t(\varphi)$ at age t, where w is the wage rate, and $e_t(\varphi)$ is a labor productivity endowment that depends on age, as well as on a stochastic productivity shock φ. This wage 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 $T_{ss}(\cdot)$. After-tax wage income at age t is therefore given by

$$y_t^{at}(\varphi) = y_t(\varphi) - T_y(y_t(\varphi)) - T_{ss}(y_t(\varphi)). \tag{3}$$

Finally, a household's asset holdings at age t earn a risk-free gross interest rate of R.

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. 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%.[4]

2.3 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.[5]

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

[5] 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.

2.4 A household's optimization problem

Let use define the state vector of an age$-t$ household as $x = \{k, \varphi, AIME\}$, where k denotes the beginning-of-period assets, φ 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.

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

$$V_t(x) = \max_{c,k',h} \left\{ u(c, 1-h, P) + \beta \frac{Q_{t+1}}{Q_t} \, E\left[V_{t+1}(x')\right] \right\} \tag{4}$$

subject to

$$c + k' = Rk + y_t^{at}(\varphi) + \Theta(t - T_r)\, b(AIME) \tag{5}$$

$$y_t^{at}(\varphi) = h_t w e_t(\varphi) - T_y\left(h_t w e_t(\varphi)\right) - T_{ss}\left(h_t w e_t(\varphi)\right) \tag{6}$$

$$0 \leq h_t \leq 1 \tag{7}$$

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

$$AIME' = \begin{cases} \left[AIME \times (t-1) + \min\left\{h_t w e_t(\varphi), \bar{y}\right\}\right]/t & t < T_r \\ AIME & t \geq T_r \end{cases}, \tag{9}$$

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 and die with zero assets, i.e. $k_0 = k_{T+1} = 0$, and prior to age T_r, the earnings history measure AIME evolves based on the realized labor productivity shocks and the endogenous labor supply decisions.

2.5 Technology and factor prices

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

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

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

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

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

where δ is the depreciation rate of physical capital.

2.6 Aggregation

The population structure in the model is as follows: at each instant a new cohort is born and the oldest cohort dies, and cohort size grows at the rate of n over time. Let us denote the measure of households at age t with state x as $\mu_t(x)$. Then, the aggregate capital stock and labor supply are given by

$$K = \sum_{t=0}^{T} N_t Q_t \int k_{t+1}(x) d\mu_t(x) \tag{13}$$

$$L = \sum_{s=0}^{T} N_t Q_t \int h_t(x) e_t(x) d\mu_t(x), \tag{14}$$

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

$$\sum_{t=0}^{T} N_t Q_t \int T_{ss}\left(h_t(x) w e_t(x)\right) d\mu_t(x) = \sum_{t=0}^{T} N_t Q_t \int \Theta(t - T_r) b(x) d\mu_t(x). \tag{15}$$

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 (BEQ)

$$BEQ + \sum_{t=0}^{T} N_t Q_t \int T_y\left(y_t(x)\right) d\mu_t(x) = G. \tag{16}$$

2.7 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.

$$\begin{aligned} C + K' - (1-\delta)K + G &= C + (n+\delta)K + G \\ &= wL + (R + \delta - 1)K \\ &= Y. \end{aligned} \tag{17}$$

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 $N_0 = 1$.

3 Calibration

3.1 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.

3.2 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.[6] 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.[7] 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.

3.3 Labor productivity

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

$$\log e_t(\varphi) = \epsilon_t + \varphi_t, \tag{18}$$

[6] 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.

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

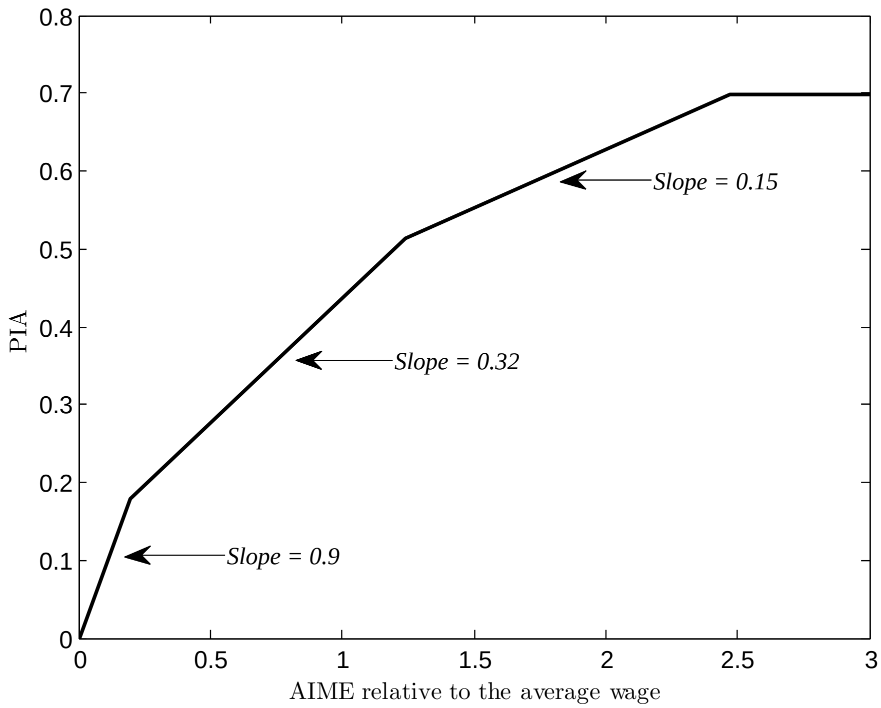

Figure 1: Benefit formula in the U.S.

where ϵ_t is a deterministic age-dependent component, and φ_t is a stochastic component, which is further decomposed as

$$\varphi_t = p + z_t + \nu_t. \tag{19}$$

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

$$z_t = \rho z_{t-1} + \upsilon_t \tag{20}$$

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

I parameterize ϵ_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 ϵ_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).

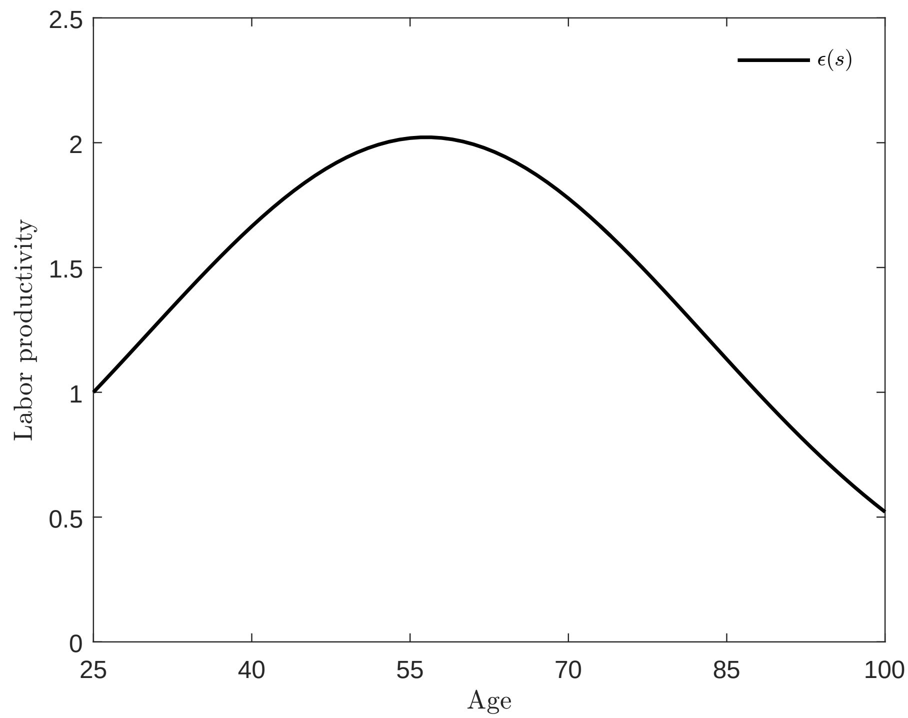

Figure 2: The age-dependent component of labor productivity from Kitao (2014).

3.4 Taxes

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

$$T_y(y) = y - (1 - \tau_y)y^{1-\tau_1}, \tag{21}$$

where $\tau_y < 1$ and $\tau_1 > 0$. Note that with $\tau_1 = 0$, equation (21) reduces to a proportional tax function with a marginal rate of τ_y. With this income tax function, after-tax labor income is log-linear in before-tax labor income, and the parameter τ_1 controls the progressivity of the tax code. Following Heathcote et al. (2014), I set the value of this parameter to $\tau_1 = 0.151$. Heathcote et al. (2014) estimate this value using data from the 2000, 2002, 2004, and 2006 waves of the PSID, and also NBER's TAXSIM program, accounting for federal and state income taxes plus public transfers.

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 $\delta = 0.08$.

3.6 Structural parameters

I jointly calibrate the remaining unobservable structural parameters of the model, i.e. the preference parameters σ, β, and η, the age-dependent labor force participation cost θ_P, and also the labor income tax parameter τ_y, to match key macroeconomic targets.

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

σ	β	η	κ_1	κ_2	κ_3	τ_y
2.5	0.95	0.4	2.31×10^{-3}	0.3059	3.75	0.84

Table 1: Structural parameter values under the baseline calibration.

	Target	Model
Capital-output ratio	3.0	2.99
Avg. hours of market work per week per worker (25-55)	40	40.2
Share of govt. expenditures in GDP	0.2	0.193
Annual Social Security benefit per retiree	–	$31,490

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.[8]

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 κ_1, κ_2, and κ_3 to match the observed rapid decline in labor force participation after age 55. The consumption share parameter (η) controls the fraction of time a household spends on market work (conditional on participation), so I calibrate it to match the hours per week target.

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

4 Baseline results

The structural parameter values under which the baseline equilibrium reasonably matches the above targets are reported in Table 1. 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 3). The model-generated values for the targets under the baseline calibration are reported in Table 2, and the cross-sectional labor force participation (the extensive margin) and labor hours, conditional on participation (the intensive margin), are reported in Figure 4.

It is clear from Figure 4 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 quite well, and it also reasonably matches the general declining trend of weekly hours over the life cycle. However, the model fails to replicate the mild hump-shape in the hours profile, and it also somewhat overestimates both participation and weekly hours at later ages. One potential strategy to improve the model's fit along this dimension is to treat the age-dependent component of labor productivity ϵ_t as an unobservable structural parameter. Treating ϵ_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).

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

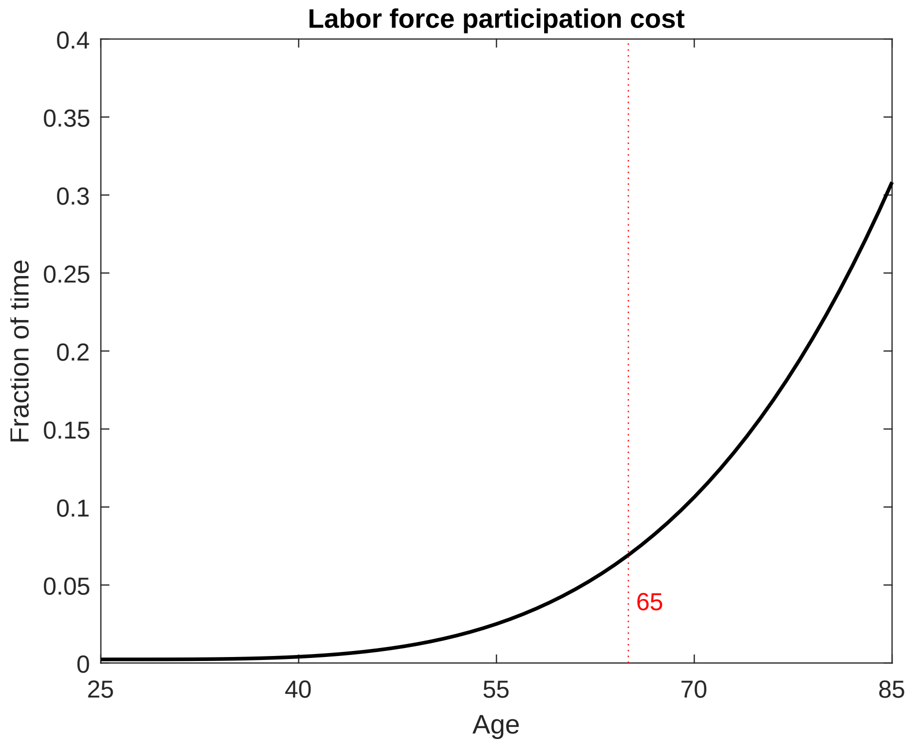

Figure 3: Age-dependent labor force participation cost.

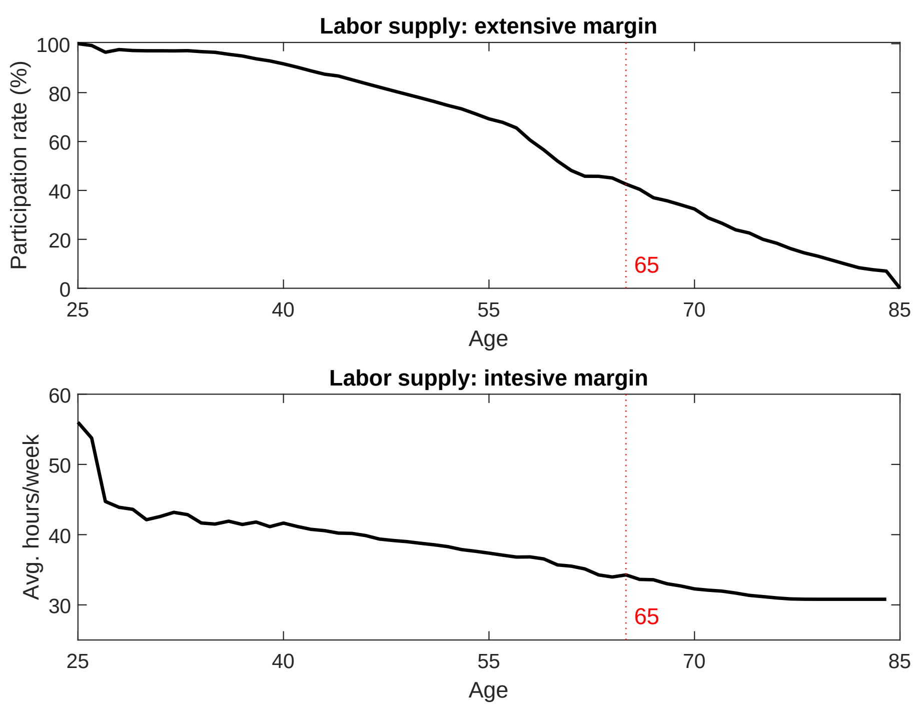

Figure 4: Cross-sectional labor supply under the baseline calibration.

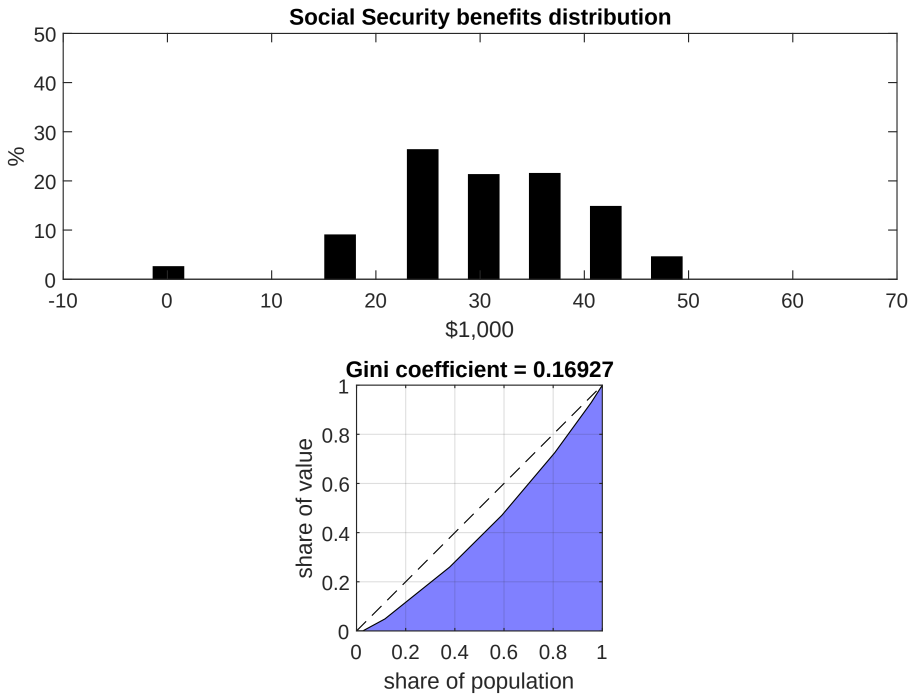

Figure 5: Baseline Social Security distribution and the Lorenz Curve.

	Low	Medium	High
Avg. Annual Social Security benefits	\$25,063	\$31,967	\$36,009

Table 3: Social Security benefits by productivity type.

It is worthwhile at this point to examine the distribution of Social Security benefits in the baseline calibration. In Figure 5, I report the cross-sectional Social Security benefit distribution (in \$1,000) in the top panel, and the associated Lorenz curve in the bottom panel. For further illustration, I report the average level of Social Security benefits by permanent productivity shock type in Table 3. As is clear from Figure 5, about 2.5% of the population receives almost no benefits. These are households who do not participate in the labor market for long enough to even qualify for Social Security. However, in the \$20,000 to \$50,000 range, the benefit distribution is positively skewed, as evidenced by its thicker right tail. Comparing the average level of Social Security benefits by productivity type, Table 3 shows that while benefits are positively related to earnings, average benefits of top earners are roughly 13% higher than medium earners, but average benefits of medium earners are more than 27% higher than low earners. Finally, the Gini coefficient of Social Security benefits in the baseline calibration is 0.1693, as seen in Figure 5.

5 Introducing a Means Test

As discussed earlier, Social Security benefits in the U.S. are calculated as a function of average work-life income, and earnings only up to the cap are counted towards the benefits. Therefore, the maximum taxable earnings also sets a de-facto limit on the amount of benefit payments from Social Security, indicating an "indirect" means test. However, the payments themselves are not conditioned in any way on the other financial resources that older households have access to. One

	Baseline	With Means Test	% change
Output (Y)	65.33	64.93	−0.62
Capital (K)	195.92	194.46	−0.75
Consumption (C)	38.46	38.15	−0.81
Bequests (BEQ)	2.66	2.58	−2.76
Labor (L)	36.17	35.92	−0.69
K/Y	3.00	3.00	−0.12
T/Y	0.19	0.19	−0.13
Avg. hours/week workers	40.21	40.24	0.07
Social Security Spending	2.31	2.25	−2.29
Govt. Spending	15.26	15.09	−1.11
$k = K/(AL)$	5.42	5.42	0.17
Gross interest rate (R)	1.04	1.04	−0.01
Wages (W)	1.17	1.17	0.07
Median income (US$)	0.53	0.52	−2.31
Avg. Social Security benefit (US$1,000)	31.49	31.50	0.02

Table 4: Macroeconomic effects of the means test.

	Low	Medium	High
Baseline	$25,063	$31,967	$36,009
Means test	$26,205	$31,994	$34,797

Table 5: Annual Social Security benefits between the baseline and the means test.

proposal to improve Social Security's solvency while still retaining the program's core mission of preventing old-age poverty could be to introduce a means test that reduces benefit payments to households with above-average earnings and/or asset holdings.

To test such a policy, I examine the consequences of an experiment in which Social Security benefits are cut by 75% for households whose earnings (labor plus capital income) are more than 200% of median earnings. Specifically, I compute a new equilibrium of the model with all the institutional features of Social Security fixed at the baseline level, while introducing the means test. I report the macroeconomic effects of this experiment in Table 4.[9] The above results suggest that introducing the means test leads to a 2.3% reduction in the size of Social Security compared to the baseline, while keeping the average benefit per retiree roughly unchanged. More interestingly, this higher implicit tax on old-age resources has only a minor effect on the aggregate capital and labor, both of which decline by only less than one percentage point.

To illustrate the microeconomic (distributional) consequences of the means test, I report the distribution of Social Security benefits under this experiment, along with its corresponding Lorenz curve in Figure 6, and also the average levels of Social Security benefits by productivity type in Table 5. Comparing Figure 6 to Figure 5, it is clear that the means test leads to a slight rightward shift in the Social Security benefits distribution. Moreover, the Gini coefficient of Social Security benefits increases slightly from its baseline value of 0.1693 to 0.1696 under the means test. Finally, Table 5 shows that the means test causes an increase in the average benefit levels for low- and medium-earnings households, while reducing the average benefits for top earners. Together, these

[9] It is worth mentioning here that the current model, in general, is not a good framework to measure the welfare implications these policy changes. This is because Social Security has no welfare improving role in a rational-expectations general-equilibrium framework with endogenous labor (Bagchi, 2015). In a framework such as this, the welfare gains from increased capital accumulation and labor supply associated with a cut in Social Security almost always outweigh the welfare losses from the reduced public insurance. As a result, any downsizing of Social Security is always welfare improving from the perspective of a household that is born into the steady state.

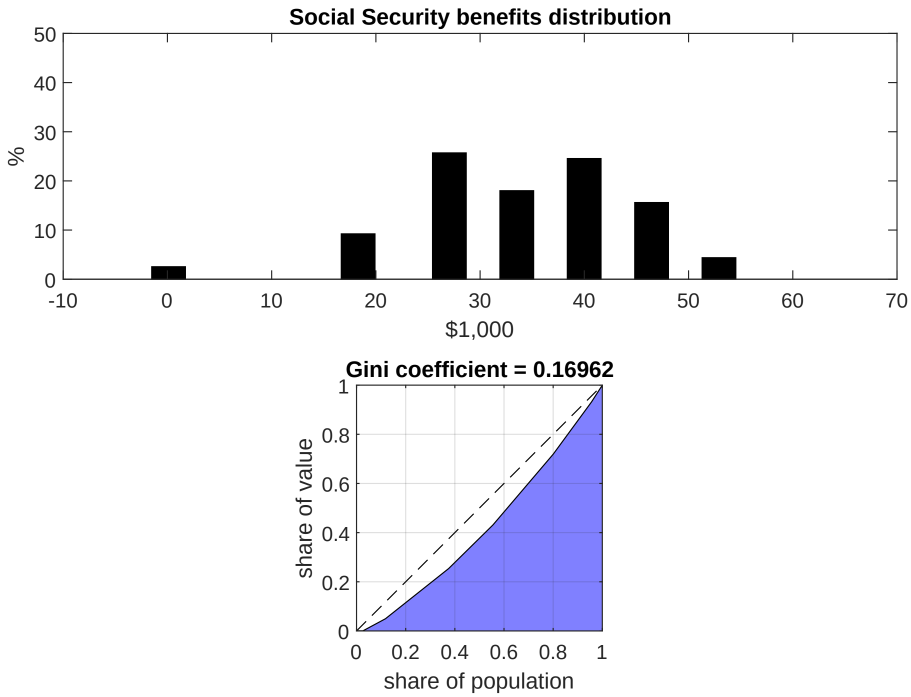

Figure 6: Social Security distribution and the Lorenz curve with the means test.

results suggest that the means test appears to facilitate a downsizing of Social Security without sacrificing the program's core progressive mission of providing insurance against the likelihood of old-age poverty.

6 A Comparable Tax Cut

Given the novel nature of the means test, especially in the context of Social Security reform, it is useful at this point to compare its macro- and microeconomic effects to those from a much more commonly discussed policy change: a fiscally comparable cut in the Social Security payroll tax rate. To do this, I compute a new equilibrium of the model with all the institutional features of Social Security fixed at the baseline level, while cutting the payroll tax rate to 10.17% from its baseline level of 10.6%, which accomplishes a roughly identical reduction in the size of the program. I compare the results of this experiment to those under the means test in Table 6, and I report the average Social Security benefits by productivity type under both experiments in Table 7.

It is clear from the table that the effects of the tax cut are quite different from those of the means test. First, the tax cut leads to a small increase in capital, labor, and output, which is the exact opposite of the effect of the means test. Second, average Social Security benefits decline by roughly 2% with the tax cut, whereas it remains roughly unchanged under the means test. Finally, in terms of the level of Social Security benefits, the low- and medium-income households are strictly better off under the means test compared to the tax cut. This is confirmed in Table 7, which shows that the tax cut leads to an across-the-board decline in Social Security benefits for all households, compared to the means test when benefits decline only for the top earners.

	Baseline	With Means Test	% change	With tax cut	% change
Output (Y)	65.33	64.93	-0.62	65.45	0.19
Capital (K)	195.92	194.46	-0.75	196.84	0.47
Consumption (C)	38.46	38.15	-0.81	38.54	0.21
Bequest (BEQ)	2.66	2.58	-2.76	2.68	0.99
Labor (L)	36.17	35.92	-0.69	36.21	0.12
K/Y	3.00	3.00	-0.12	3.01	0.28
T/Y	0.19	0.19	-0.13	0.19	0.01
Avg. hours/week workers	40.21	40.24	0.07	40.20	-0.03
Social Security Spending	2.31	2.25	-2.29	2.25	-2.24
Govt. Spending	15.26	15.09	-1.11	15.31	0.33
$k = K/(AL)$	5.42	5.42	0.17	5.43	0.19
Gross interest rate (R)	1.04	1.04	-0.01	1.04	-0.01
Wages (W)	1.17	1.17	0.07	1.17	0.06
Median income (US$)	0.53	0.52	-2.31	0.53	-0.20
Avg. Social Security benefit (US$1,000)	31.49	31.50	0.02	30.85	-2.05

Table 6: The means test and the tax cut compared.

	Low	Medium	High
Baseline	$25,063	$31,967	$36,009
Means test	$26,205	$31,994	$34,797
Tax cut	$24,487	$31,321	$35,304

Table 7: Average Social Security benefits compared.

7 Earnings Vs Asset Tests

While labor and capital income are flows that accrue to households over the life cycle, a second category of financial resources that older households have access to is their stock of asset holdings. So far, I have interpreted means testing as conditioning Social Security benefits on a household's earnings. In this section, I examine how conditioning the means test on asset holdings, rather than earnings, changes the macroeconomic and welfare consequences of means testing in general.

To do this, I now compute an experiment in which, as before, Social Security benefits are cut for households whose asset holdings are more than 200% of the economy's median asset holdings. However, in order to keep the asset-based means test fiscally comparable to the earnings-based means test, I adjust the benefit reduction rate to 76%. This rate yields a 2.2% decline in the size of Social Security in equilibrium, which is roughly identical to those under the earlier experiments. I report the macro- and the microeconomic effects of this experiment in Table 8, and the average level of Social Security benefits under the two different means tests in Table 9.

It clear from the tables that the macroeconomic effects of an asset-based means test are quite different from those of an earnings-based means test. First, the assets-based means test, being a direct tax on capital, has a large negative effect on aggregate capital stock. Together with a small decline in labor, this causes the largest decline in output across the three experiments (-2.4%). However, aggregate consumption remains mostly unchanged despite this fact, driven by a large decline in the government expenditures (-4.3%), caused by an almost 14% decline in accidental bequests. Average Social Security benefits decline by about 1%, mostly due to the fact that the benefits for top earners decline by the largest percentage under this experiment.

The microeconomic effects of the asset-based means test are also quite different. As noted above,

	Baseline	Earnings-based	% change	Assets-based	% change
Output (Y)	65.33	64.93	−0.62	63.76	−2.41
Capital (K)	195.92	194.46	−0.75	186.88	−4.62
Consumption (C)	38.46	38.15	−0.81	38.46	−0.02
Bequest (BEQ)	2.66	2.58	−2.76	2.29	−13.65
Labor (L)	36.17	35.92	−0.69	35.72	−1.26
K/Y	3.00	3.00	−0.12	2.93	−2.26
T/Y	0.19	0.19	−0.13	0.19	0.12
Avg. hours/week workers	40.21	40.24	0.07	40.04	−0.43
Social Security Spending	2.31	2.25	−2.29	2.26	−2.19
Govt. Spending	15.26	15.09	−1.11	14.61	−4.27
$k = K/(AL)$	5.42	5.42	0.17	5.24	−3.30
Gross interest rate (R)	1.04	1.04	−0.01	1.04	0.25
Wages (W)	1.17	1.17	0.07	1.16	−1.16
Median income (US$)	0.53	0.52	−2.31	0.52	−1.29
Avg. Social Security benefit (US$1,000)	31.49	31.50	0.02	31.20	−0.92

Table 8: The earnings- and asset-based means tests compared.

	Low	Medium	High
Baseline	$25,063	$31,967	$36,009
Earnings-based	$26,205	$31,994	$34,797
Assets-based	$28,995	$32,531	$28,089

Table 9: Average Social Security benefits compared.

the decline in average Social Security benefits in this case is slightly larger than the earnings-based means test. But this decline masks large increases in the benefits for low- and medium-earnings households, larger than those under the earnings-based means test (see Table 9). Overall, the above results suggest that both the earnings- and asset-based means tests yield a downsizing of Social Security, but without sacrificing the program's core progressive mission of providing insurance against the likelihood of old-age poverty.

8 Conclusions

In this paper, I examine if introducing a means test in U.S. Social Security can be an effective policy to improve the program's fiscal sustainability while retaining its core mission of preventing the likelihood of old-age poverty. To do this, I construct an overlapping-generations macroeconomic model with incomplete markets, an unfunded public pension system that mimics Social Security, and households that experience realistic mortality and labor income risks. Social Security provides partial insurance against these risks, because households do not have access to private insurance markets. Households in the model also face a progressive income tax schedule, factor markets are competitive, firms maximize profit, and the government provides public goods and Social Security. I calibrate this model to match key features of the U.S. economy, such as overall capital accumulation, pattern of labor supply over the life cycle, the earnings and Social Security benefits distributions, and the share of government expenditures in GDP.

Using this model, I first compute the macroeconomic and distributional consequences of a means test, in which benefits to households with earnings more than 200% of median income are cut by 75%. In the computation, I allow for all the household-level and macroeconomic adjustments to the policy change. I find that a means test, i.e. conditioning benefit payments on a household's earnings, leads to a higher implicit tax on old-age resources, but has desirable distributional effects. A 75% cut in the benefits to households with earnings more than 200% of the median leads to a 2.3% reduction in the overall size of the program, but has almost no effect on the average benefit level. This is despite the fact that capital and labor decline slightly due this higher implicit tax. A fiscally comparable payroll tax cut, on the other hand, leads to an across-the-board decline of 2% decline in the benefits, despite an increase in capital and labor. Finally, a comparable means test that conditions benefit payments on asset holdings rather than earnings yields a 1% decline the average benefit level, but has a large negative effect on the accidental bequests left behind by deceased households. This is because the assets-based means test is an implicit tax on capital, which reduces old-age asset holdings, and therefore the accidental bequests.

There is a large literature that has considered modifications to the various institutional features of Social Security to improve the program's fiscal sustainability in the long run. These modifications have ranged from changing the payroll tax rate and the eligibility age, to a complete privatization of the existing Pay-As-You-Go structure. The current paper contributes to this literature by evaluating whether or not an earnings- or assets-based means test on Social Security benefits can play an important role in this debate. The results in this paper suggest that this particular institutional feature may help in partially solving Social Security's long-run budgetary problems.

References

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

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

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

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

Caliendo, F. N., Guo, N. L., Hosseini, R., 2014. Social security is NOT a substitute for annuity markets. Review of Economic Dynamics 17 (4), 739–755.

Conesa, J. C., Garriga, C., 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.

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

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

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

İmrohoroğlu, A., İmrohoroğlu, S., Joines, D. H., 2003. Time-Inconsistent Preferences and Social Security. The Quarterly Journal of Economics 118 (2), 745–784.

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.

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

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

Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59 (2), 371–396.

Zhao, K., 2014. Social security and the rise in health spending. Journal of Monetary Economics 64 (0), 21–37.