Social Security Reform: Three Rawlsian Options*

Erin Cottle Hunt† Frank N. Caliendo‡

Lafayette College Utah State University

This Version: March 11, 2020. First Version: April 18, 2011.

Abstract

To maintain long run solvency in light of changing demographics, the US Social Security system needs to be reformed. We present three reform options that protect the retirement benefits of the economically vulnerable while also balancing the Social Security budget. We refer to these three options as Rawlsian reforms because, with each reform option, the Social Security benefits of those at the low end of the income distribution are left intact. Two of our reforms break the link between the benefit cap and the tax cap by lowering the benefit cap. We explore the effect of each reform on ex-ante expected utility, the distribution of private savings, and the distribution of lifetime income in a life-cycle model.

JEL codes: H3, E6.

Keywords: Social Security Reform, Inequality, Rawlsian Welfare.

*We thank two anonymous referees and the editor of ITAX for very helpful recommendations. An earlier version of this paper circulated under the title "Rawls, Pensions, and Demographic Shocks" and was presented at the Central European Program in Economic Theory at the Universite d'Udine, Italy June 2011.

†cottlee@lafayette.edu.

‡frank.caliendo@usu.edu. Please send editorial correspondence to Frank Caliendo.

1. Introduction

Social Security in the US (and in many OECD countries) is an unfunded, or pay-as-you-go system. The financing of benefits in an unfunded system depends critically on the ratio of workers to retirees, and this ratio is predicted to decline dramatically in many countries. For example, in the US during the period 1974-2008 the ratio of workers to retirees was stable between 3.4 and 3.2. The ratio has declined since 2008 and is predicted to fall to 2.0 over the next few decades. The Social Security Administration (SSA) estimates that they will only be able to pay full benefits until the year 2033. To maintain long run solvency, Social Security taxes will need to be raised, benefits will need to fall, or some combination of both.

In addition to transferring resources across generations, Social Security also redistributes resources *within* a generation. Social Security benefits are paid according to a progressive formula, thus redistributing resources from high-income retirees to low-income retirees.[1] Many retirees rely on Social Security benefits for a large fraction of their retirement consumption, particularly retirees who had low income during their working years. In this paper, we look for reforms that protect the benefits of low-income retirees. Our focus on retirees who earned low wages during their working life is shared by many policy makers.

The Senate's Special Committee on Aging (2010) and the Government Accountability Office (2010) both encourage Congress to take a Rawlsian perspective when evaluating Social Security reform measures. In their estimation, a desirable reform should not only balance the budget, but it should also protect benefits for the economically vulnerable.[2]

Many recent proposals from members of Congress have a similar theme and attempt to balance the

[1] Cremer, De Donder, Maldonado, and Pestieau (2008) address the issue of how redistributive a pension system should be in the first place, absent demographic shocks.

[2] According to the Committee, "Congress should enact modest changes to Social Security in the near future in order to bring its long-term financing into balance and improve benefits for those who need them most." The GAO uses similar language: "...the nation faces the challenge of improving long-term program solvency, while also ensuring benefit adequacy for economically vulnerable beneficiaries."

budget (or reduce the long-run funding shortfall) while protecting the benefits of the poor. For example, a recent report analyzes the long-run budgetary impacts of roughly 150 provisions that would alter Social Security taxes or benefits (SSA 2017). Several of the provisions in the report protect (or increase) the benefits of the economically vulnerable by reducing the benefits of higher earners. The Congressional Budget Office has also considered several similar policy changes in reports and in an interactive, on-line platform that allows users to choose options from a menu of possible reforms to see how the reforms impact the Social Security budget (CBO 2015, 2018).[3] Several of the reform options protect or expand the benefits of the lowest earners.[4]

Consistent with these efforts, in this paper we introduce three of our own potential reforms to Social Security (explained in detail in the next subsection) that are likewise aimed at protecting the benefits of the poor. We study a stylized economy with a calibrated wage density function that captures the degree of heterogeneity in the earnings of US workers. Including heterogeneity allows us to move beyond the limited set of responses to a demographic transition that would arise from a representative agent model toward a richer set of recommendations that focus on distributional concerns. We impose the discipline of a balanced-budget constraint throughout all of our analysis. This generates trade-offs between reducing benefits and raising taxes, and also between reducing the benefits of some while protecting the benefits of others. We embed our reforms into a life-cycle consumption-savings model to explore the effects of each reform on the distribution of saving and lifetime income as well as the effect on ex-ante expected utility.[5]

We consider two model economies to explore these household responses and welfare. In the first, accidental bequests are distributed evenly to all living agents in the model. This is the standard assumption

[3] The interactive platform is available at https://www.cbo.gov/publication/54868.

[4] The CBO explains, this type of reform "would better target Social Security benefits toward people who need them more—protecting or expanding benefits for people with low average earnings while reducing payments to people with higher average earnings. This option would help make the Social Security system more progressive at a time when growing disparities in life expectancy by income level are making the system less progressive" (CBO 2018). For additional examples of Rawlsian Social Security reforms, see Diamond's (2018) description of two proposed reforms from opposite sides of the political spectrum and see also the Bipartisan Policy Center Commission on Retirement Security and Personal Savings (2016).

[5] The effect of Social Security on private savings is one of the most-studied behavioral questions in the public pension literature (see Slavov, Gorry, Gorry, and Caliendo (2019) for a review).

in macroeconomic models. In the second, bequests depend directly on wage type. Workers of a given wage inherit bequests from those who had the same wage type. Intergenerational inequality is more persistent in this framework. In both types of economies that we consider, the size of the bequests depend on Social Security, as Social Security crowds out private saving, but bequests are a particularly large part of income for low-wage workers in Economy 1 (with uniform bequests). Thus Social Security's crowding out of private saving interferes with the redistributive transmission of wealth across generations, and such "family insurance" considerations have important welfare implications that affect the performance of reform options that are intended to help the poor. Our study of two model economies captures these potential effects.

The need for Social Security reform in light of demographic changes has been explored in many papers (see, e.g., the references in Bagchi (2015, 2016)). Reform options include changes to taxes and benefits (Diamond and Orszag (2005) and De Nardi, İmrohoroğlu, and Sargent (1999)), changes to retirement age (İmrohoroğlu and Kitao (2012)), and changes to both retirement age and taxes or benefits (Kitao (2014)). Many papers also consider moving towards a fully-funded or privatized Social Security system.[6] Given the variety of policy options available, considerable uncertainty exists about the timing and structure of reform. A growing section of the literature examines the welfare effects of Social Security policy uncertainty.[7] The emphasis of this paper is more closely aligned with analysis conducted by the Social Security Administration and the Congressional Budget Office.

[6]Papers exploring full or partial privatization of Social Security as a response to demographic changes include: Feldstein (1996), Samwick (1998), Conesa and Krueger (1999), Huggett and Ventura (1999), Kotlikoff, Smetters, and Walliser (1999, 2007), Nishiyama and Smetters (2007), Conesa and Garriga (2008), and İmrohoroğlu and Kitao (2009). Attanasio, Kitao, and Violante (2007) model changes that maintain the pay-as-you-go nature of Social Security and also a move to a fully-funded system in a two-region model of the global economy.

[7]Social Security policy uncertainty papers include Bütler (1999), Gomes, Kotlikoff, and Viceira (2007), Caliendo, Gorry, and Slavov, (2019), Nelson (2017), Kitao (2018), Luttmer and Samwick (2018), and Cottle Hunt (2019).

1.1. Current System and Three Potential Reforms

The retirement portion of the current Social Security tax rate is 10.6% of wage income. The tax applies on wage income up to a taxable maximum, or tax cap. In 2019, the tax cap is $132,900.[8] The current formula for calculating Social Security benefits (primary insurance amount, PIA) is a piecewise linear function that is increasing in the average of the individual's highest 35 years of wage earnings (average indexed monthly earnings, AIME). The Social Security system pays benefits equal to 90% of earnings up to the first "bend point," 32% of earnings between the first and second bend points, and 15% of earnings between the second and third bend points. No benefits are paid on earnings beyond the third bend point. Under the current system, the third bend point is the same amount as the taxable maximum. That is, the maximum benefit and the taxable maximum are linked. The bend points are adjusted each year to maintain their relative positions at approximately 0.2, 1.24, and 2.47 times the average wage (Ortiz (2014) and the references therein). The declining ratio of workers to retirees requires an increase in the tax beyond 10.6%, or a decrease in some or all of the slopes of the different segments of the benefit-earning rule to bring down average benefits paid. It is also possible to increase tax revenue by increasing the tax cap, or to reduce benefits by reducing the maximum benefit.

We focus on three concrete adjustments to the current benefit-earning rule, two of which break the link between the tax cap and the maximum benefit. These three options are designed to protect the benefits of the economically vulnerable, and each option keeps the budget balanced under future demographics. In all cases, we assume a large demographic change (from 3.3 to 2.0 workers per retiree), and we assume that "IOUs" from the US Treasury (trust fund assets) are already fully depleted. This is equivalent to assuming the reform happens after 2033. Thus, we take the perspective of planning for the long term.

[8] Earnings above the taxable maximum are subject to a payroll tax for the Medicare Hospital Insurance program, but those earnings are not subject to Social Security taxes. The taxable maximum changes each year according the changes in the national average wage index. The 10.6% Social Security tax rate applies to the Old Age Survivors Insurance portion of the program. A separate tax funds the Disability Insurance portion of the program. For more details, see https://www.ssa.gov/oact/cola/cbb.html.

Our three proposals are summarized below.

Policy Option 1: Reduce the maximum benefit without changing taxes. In this option, the policy maker leaves the current benefit-earning rule in place for as many poor earners as possible by reducing the maximum benefit, leaving taxes (both the tax rate and the tax cap) unchanged. In effect, the benefit rule is left intact for as many low-income earners as possible and then the benefit rule becomes flat thereafter. In our calibrated model, the benefit rule can be left intact for everyone below 34% of the mean wage, which corresponds to protecting benefits for the bottom 20% of the population. To then balance the budget without a tax increase, it is necessary to reduce the benefit cap, such that all workers earning more than 34% of the mean wage receive the same benefit. In other words, it is necessary to completely flatten out the benefit-earning rule (zero slope) after 34% of the mean wage.

Policy Option 2: Reduce the maximum benefit and increase the tax rate. Protecting benefits for more than the bottom 20% of the population while maintaining a weakly increasing benefit-earning rule as in Option 1 can be accomplished only with a tax increase. As an example, we create a post demographic change benefit-earning rule that maintains the current level of benefits for all individuals with below average wage earnings and is financed with an approximately seven percentage point increase in the tax rate, leaving the tax cap unchanged. The maximum benefit is reduced to balance the budget with the tax rate increase. Thus, in this option the new benefit-earning rule exactly traces the old rule up to the mean wage and then flattens out thereafter. This option protects the benefits of the bottom 58% of the population.[9]

[9] There is a very large literature prompted by Feldstein (1985) which studies the optimal size of the Social Security tax rate, independent of demographic shocks. A survey by Cremer and Pestieau (2011) details some of the advances in this literature. We intentionally remain agnostic about optimal program size since this issue depends critically on the assumptions the researcher makes about the degree and type of bounded rationality of households. We instead focus on distinguishing between adjustments to the benefit-earning rule that require additional taxes to finance and adjustments that do not require additional taxes, thereby providing policy makers with a menu of potential Rawlsian responses to the demographic shock.

Policy Option 3: Protect the benefits of the poor and also maintain a strictly increasing benefit-earning rule. As a third option, we show that the benefit-earning rule can be left exactly as it currently is up to the first bend point (90% slope up to 20% of the mean wage) if the slopes of the second and third segments of the benefit-earning rule drop from the current slopes of 32% and 15% to the new slopes of approximately 5% and 3%, respectively, without changing the tax cap or tax rate. The new rule balances the budget, protects the benefits of the very poorest segment of the population (those below the first bend point), maintains a strictly positive slope, and does not require additional taxes. Some variation of this third option may be the most attractive to policy makers since the Senate Aging Committee has argued that benefits received bear some relationship to the amount of taxes paid.

In all of our simulations, Policy Option 1 confers the highest expected utility. Policy Option 3 is almost as good in an expected utility sense, while Policy Option 2 confers much lower expected utility than the other two.

Our proposals are precise adjustments to the current benefit-earning rule that are Rawlsian in nature and focus on protecting the benefits of the most economically vulnerable in the face of a demographic transition. Two of our proposals break the link between the taxable maximum and the maximum benefit, by reducing the maximum benefit. By breaking the link between the taxable maximum and the maximum benefit (or benefit cap), the policy maker can protect the benefits of the poorest households without necessarily changing taxes.

While there are plenty of reform proposals already under consideration, we feel that our proposals may be useful because they highlight policy levers that are simple and have not received attention yet. We are aware of several proposals to increase the taxable maximum, but we are not aware of other proposals to reduce the benefit cap.[10]

[10] See, for example, Bagchi (2017) or the interactive CBO tool https://www.cbo.gov/publication/54868. Proposals to increase the taxable maximum often increase the benefit cap as well, while some proposals break the link. Because the benefit-earning rule is progressive, increasing both the taxable maximum and benefit cap increases net revenues. However,

2. Model

We augment the classic Yaari (1965) consumption-saving model with uncertain lifespan to include wage heterogeneity and a Social Security system that mimics key features of the US program. We use the model to study various reforms to Social Security that protect the benefits of the poor and ensure that the system remains solvent in the face of a large expansion in life expectancy.[11]

2.1. *Household Behavior*

Age is continuous and is indexed by t. At each moment in time an infinitely divisible cohort of unit mass is born. Individuals are born at $t = 0$ and die no later than $t = T$. The probability of surviving to age t from the perspective of age 0 is $S(t)$, where $S(0) = 1$ and $S(T) = 0$. Retirement from the workforce occurs at $t = t_R$.[12]

An individual's wage rate is drawn from the unit interval, $w \in [0, 1]$. Wages vary across the population according to the density function $g(w)$, and individuals have full information about their wage type from birth.

Individuals must pay Social Security taxes at rate τ on wage income up to a wage cap $w_c < 1$. That is, individuals below the wage cap ($w < w_c$) pay Social Security taxes τw for all $t \leq t_R$, and individuals

breaking the link can increase net revenues (or reduce funding short-falls) by a greater amount.

[11] We do not study disability risk or disability insurance within our model. Disability insurance is an important component of the US Social Security program: workers who suffer a disability shock may collect full benefits before reaching their normal retirement age without facing penalties for their condensed earnings history. While this feature provides important insurance against a specific risk, in this paper we study only longevity risk in particular, and we focus on reforms to the retirement portion of the Social Security program. Future work could consider the distributional effects of disability risk, especially if low wage earners are more likely to become disabled on the job (due to the nature of work conducted). Modeling disability risk and a disability insurance program requires a dynamic stochastic model of consumption and saving with career-length risk (as in Caliendo, Casanova, Gorry, and Slavov (2019)), and such a model could be used to understand how the disability insurance program would need to evolve in the future to remain solvent while providing coverage to disabled workers.

[12] A variety of factors are causing the population to age in the US. Increases in life expectancy, decreases in fertility, and the baby boom phenomenon all act to reduce the ratio of workers to retires. As a simplification, we focus on the increase in life expectancy in this paper and later when we calibrate the model, we assume that life expectancy is the sole source of the decline in the ratio of workers to retirees. We are comfortable with this simplification because, in our model, the projected shortfall in the aggregate Social Security budget depends on the ratio of workers to retirees but not specifically on how that ratio is generated. So loading the entirety of the demographic transition into changes in life expectancy as we do below will not have a quantitative effect on the reforms that we study.

at or above the wage cap ($w \geq w_c$) pay taxes τw_c for all $t \leq t_R$.

After retirement, individuals receive Social Security benefits $b(w)$ that depend on their wage earnings over the working period. In the calibration section below we will be more explicit about the precise functional form of $b(w)$, but for now it is sufficient to make a few points. First, in reality and in our model, benefits are weakly increasing in wages and there is a limit to the level of benefits that can be received. We denote the benefit cap as b_c. Second, and crucial to our analysis is this: under current law the wage and benefit caps are *synchronized* in the sense that individuals at or above the wage cap $w \geq w_c$ will collect benefits $b(w) = b_c$. However, in our model we allow for the possibility that as we move up the wage distribution, individuals hit the benefit cap before hitting the taxable wage cap. That is, at some wage $\hat{w} < w_c$, the benefit-earning rule hits the maximum benefit, $b(\hat{w}) = b_c$.

Throughout the paper we refer to $b(w)$ as the benefit-earning rule. In the US system, this function has a specific name, the "primary insurance amount" (PIA). For our purposes, these terms can be used interchangeably.

The individual's consumption is $c(t)$ and utility from consumption is $u(c(t))$, with $u_c > 0$ and $u_{cc} < 0$. Private annuity markets do not exist and all saving is done in a risk-free account that pays interest at rate r. The individual's savings balance at time t is denoted $k(t)$. The individual starts the life cycle with nothing $k(0) = 0$ and must pay off all debts by the maximum survival age, $k(T) = 0$. Future utility is discounted at rate ρ.

The assets of the deceased are bequeathed to survivors. We will say more below about the nature of bequest income, but for now it is sufficient to denote the flow of bequest income received by wage type w as $B(w)$. The individual does not have a bequest motive; instead, bequest income is "accidental".

The individual chooses $(c(t), k(t))_{t \in [0,T]}$ to solve

$$\max \int_0^T e^{-\rho t} S(t) u(c(t)) dt,$$

subject to

$$\dot{k}(t) = rk(t) + w - \tau \min(w, w_c) + B(w) - c(t), \text{ for } t \in [0, t_R],$$

$$\dot{k}(t) = rk(t) + b(w) + B(w) - c(t), \text{ for } t \in [t_R, T],$$

$$k(0) = k(T) = 0.$$

We consider two types of economies that differ according to the way in which wealth is transmitted across generations. First, we consider the standard macroeconomic framework in which accidental bequests are distributed equally to all of the living. All agents in the model receive the same bequest, regardless of wage. Thus, bequests provide partial insurance against drawing a low wage in this framework. We view this as the benchmark model. As an alternative and potentially more realistic specification, we also consider a world in which bequests depend directly on wage type. Workers of a given wage inherit bequests from those who had the same wage type. Income inequality is persistent across generations and is more severe than in first economy.

2.2. *Economy I: Mean-Reverting Wealth with Uniform Bequests*

In this economy wealth is transmitted across generations in the traditional manner in the macroeconomics literature. Bequest income is spread evenly across all surviving individuals (of all ages and wage types), that is $B(w) = B$ for all w and all t. Essentially, wage income is mean reverting in the sense that a given individual's wage type is a random draw from a distribution that bears no resemblance to the wage type of one's family predecessors.

Total bequest income received by survivors equals the total assets of the deceased according to the following aggregate resource constraint

$$\int_0^T S(t)Bdt = \int_0^1 \int_0^T g(w)\left(-\dot{S}(t)\right) k(t|w, B)dtdw.$$

And the Social Security budget must balance

$$\int_0^{w_c}\int_0^{t_R} g(w)S(t)\tau w dt dw + \left(\int_{w_c}^1 g(w)dw\right)\int_0^{t_R} S(t)\tau w_c dt = \int_0^1\int_{t_R}^T g(w)S(t)b(w)dtdw.$$

The left side of the equation is aggregate (average) taxes collected and the right side is benefits paid. Solving for the balanced budget Social Security tax we have

$$\tau = \frac{\int_0^1\int_{t_R}^T g(w)S(t)b(w)dtdw}{\int_0^{w_c}\int_0^{t_R} g(w)S(t)wdtdw + \left(\int_{w_c}^1 g(w)dw\right)\int_0^{t_R} S(t)w_c dt}.$$

A **Stationary Equilibrium** is comprised of household allocations $(c(t|w), k(t|w))_{t\in[0,T]}$ for all w that solve each individual's optimization problem, bequest income B that satisfies the aggregate resource constraint, and Social Security tax rate τ, taxable wage cap w_c, and benefit cap b_c that jointly balance the Social Security budget.

2.3. *Economy II (Dynasty Economy): Maximum Inequality with Wealth Persistence Across Generations*

In this economy wealth is dynastic and is therefore transmitted within wage types. Individuals of a given wage type collect bequest income from deceased individuals of the same wage type. Note that in this economy, the intergenerational transmission of wealth is no longer redistributive. In Economy I, all wage types receive the same bequest income, so in this sense the economy has a natural, built-in insurance mechanism to it: low wage earners get big bequests and high earners get small bequests, relative to their earnings. In Economy II, this insurance feature is removed, income inequality is more pronounced, and Social Security's redistributive benefit-earning rule is the only insurance against low wage realizations.

Essentially, this economy is composed of a continuum of equilibria, with bequest income differing by wage type $B(w)$, which satisfies the following

$$\int_0^T S(t)B(w)dt = \int_0^T \left(-\dot{S}(t)\right) k(t|w, B(w))dt \text{ for all } w \in [0,1].$$

The Social Security budget must balance in Economy II as in Economy I

$$\tau = \frac{\int_0^1 \int_{t_R}^T g(w)S(t)b(w)dtdw}{\int_0^{w_c} \int_0^{t_R} g(w)S(t)wdtdw + \left(\int_{w_c}^1 g(w)dw\right) \int_0^{t_R} S(t)w_c dt}.$$

A **Stationary Equilibrium** is comprised of household allocations $(c(t|w), k(t|w))_{t\in[0,T]}$ for all w that solve each individual's optimization problem, a distribution of bequest income $B(w)$ that satisfies a continuum of aggregate resource constraints (one for each wage type), and Social Security tax rate τ, taxable wage cap w_c, and benefit cap b_c that jointly balance the Social Security budget.

3. Three Rawlsian Reform Options

A Social Security system is a collection of parameters and functions $(\tau, b(w), b_c, w_c)$—a tax rate on wage earnings, a benefit-earning rule, and benefit cap, and a taxable wage cap that balance the Social Security budget constraint as explained above. We denote the "current" Social Security system as $(\tau^0, b^0(w), b_c^0, w_c^0)$. The current system is associated with current demographics, $S^0(t)$. However, the current Social Security arrangement is not feasible in a budgetary sense under a future state of demographics, $S^1(t)$ where $S^1(t) > S^0(t)$ for all $t \in (0,1)$. As people live longer and the ratio of workers to retirees declines, either taxes must be raised or benefits must fall, or both. In this section we study three possible reforms options that balance the budget and are Rawlsian in nature.

3.1. ***Policy Option 1: Reduce Benefit Cap without Changing Taxes***

In this option, the policy maker keeps this existing benefit-earning rule exactly in place for as many poor earners as possible by reducing the maximum benefit, leaving taxes (both the tax rate and the tax cap) unchanged. In effect, the benefit rule is left intact for as many low-income earners as possible and then the benefit rule becomes flat at the point that tax revenues are exhausted.

We solve the balanced-budget constraint for the new benefit cap $b_c^1 = b^0(\hat{w})$ (and hence we solve for the threshold wage associated with this new benefit cap, $\hat{w}$) using future survival probabilities $S^1(t)$ but holding the tax rate and taxable wage cap fixed at their original levels (τ^0, w_c^0)

$$\begin{aligned} & \int_0^{w_c^0}\int_0^{t_R} g(w)S^1(t)\tau^0 w dt dw + \left(\int_{w_c^0}^1 g(w)dw\right)\int_0^{t_R} S^1(t)\tau^0 w_c^0 dt \\ = & \int_0^{\hat{w}}\int_{t_R}^{T} g(w)S^1(t)b^0(w)dt dw + \left(\int_{\hat{w}}^1 g(w)dw\right)\int_{t_R}^{T} S^1(t)b_c^1 dt. \end{aligned}$$

Upon solving this equation for the new, threshold wage $\hat{w}$, we identify the fraction of the wage distribution whose benefits are left intact, $\int_0^{\hat{w}} g(w)dw$. All individuals with wages $w < \hat{w}$ continue to collect benefits according to the current benefit-earning rule $b^0(w)$ even though life expectancies have increased according to $S^1(t)$. For everyone else with wages $w \geq \hat{w}$, representing population share $\int_{\hat{w}}^1 g(w)dw$, benefits are equal to the new and lower benefit cap $b_c^1 = b^0(\hat{w})$.

3.2. ***Policy Option 2: Reduce Benefit Cap and Increase Tax Rate***

Suppose policy makers wish to protect a larger share of the low end of the wage distribution than what is possible holding taxes fixed as in Option 1 above. That is, suppose policy makers wish to preserve benefits for all wage earners $w \leq w^*$ where $w^* > \hat{w}$. For instance, policy makers may wish to preserve the benefits of all individuals with below-average wage income.

Utilizing the balanced-budget condition, the new Social Security tax rate this is needed to protect a given share $\int_0^{w^*} g(w)dw$ of individuals at the bottom of the distribution, holding the taxable wage cap fixed at its original level, is

$$\tau^1(w^*) = \frac{\int_0^{w^*} \int_{t_R}^T g(w)S^1(t)b^0(w)dtdw + \left(\int_{w^*}^1 g(w)dw\right) \int_{t_R}^T S^1(t)b^0(w^*)dt}{\int_0^{w_c^0} \int_0^{t_R} g(w)S^1(t)wdtdw + \left(\int_{w_c^0}^1 g(w)dw\right) \int_0^{t_R} S^1(t)w_c^0 dt}.$$

With this reform option, upon selecting the share of the population at the lower end of the distribution whose benefits are protected, $\int_0^{w^*} g(w)dw$, policy makers then know the required tax rate $\tau^1(w^*)$ that is needed to ensure solvency of Social Security under future demographics $S^1(t)$. All individuals with wages $w < w^*$ continue to collect benefits according to the current benefit-earning rule $b^0(w)$ even though life expectancies have increased according to $S^1(t)$. For everyone else with wages $w \geq w^*$, representing population share $\int_{w^*}^1 g(w)dw$, benefits are equal to the new and lower benefit cap $b_c^1 = b^0(w^*)$.

3.3. *Policy Option 3: Preserve Strict Earnings-Based Feature without Changing Caps*

Confining ourselves to leaving the tax rate and taxable wage cap as they currently are, another possible Rawlsian reform option is to leave benefits intact for all individuals with wages less than an even lower threshold $w' < \hat{w}$, which frees up tax revenue to preserve the strictly increasing feature of the benefit-earning rule thereafter. In other words, all individuals with wages $w < w'$ are allowed to keep their original Social Security benefits $b^0(w)$, while all individuals with wages $w \geq w'$ collect Social Security benefits $b^1(w) = b^0(w') + \delta(b^0(w) - b^0(w'))$. This option maintains the link between the taxable wage cap and the benefit cap, although the maximum benefit is lower than under the current system. Here δ is a scalar less than 1. The Social Security budget will balance if

$$\begin{aligned} & \int_0^{w_c^0} \int_0^{t_R} g(w)S^1(t)\tau^0 wdtdw + \left(\int_{w_c^0}^1 g(w)dw\right) \int_0^{t_R} S^1(t)\tau^0 w_c^0 dt \\ = & \int_0^{w'} \int_{t_R}^T g(w)S^1(t)b^0(w)dtdw + \int_{w'}^1 \int_{t_R}^T g(w)S^1(t)[b^0(w') + \delta(b^0(w) - b^0(w'))]dtdw. \end{aligned}$$

Notice that we can solve this equation for the scalar δ

$$\delta = \frac{\int_0^{w_c^0}\int_0^{t_R} g(w)S^1(t)\tau^0 w dt dw + \left(\int_{w_c^0}^1 g(w)dw\right)\int_0^{t_R} S^1(t)\tau^0 w_c^0 dt}{\int_{w'}^1\int_{t_R}^T g(w)S^1(t)(b^0(w)-b^0(w'))dtdw} - \frac{\int_0^{w'}\int_{t_R}^T g(w)S^1(t)b^0(w)dtdw + \int_{w'}^1\int_{t_R}^T g(w)S^1(t)b^0(w')dtdw}{\int_{w'}^1\int_{t_R}^T g(w)S^1(t)(b^0(w)-b^0(w'))dtdw}.$$

Hence, for a given threshold wage w' that is less than the threshold wage associated with maximum protection to the poor $\hat{w}$, policy makers can leave benefits intact for all individuals with wages $w < w'$ and can reduce benefits by a constant scale factor for all individuals with wages $w \geq w'$.

4. Calibrating the Theory to the US

To explore the quantitative implications of our model, we must first select numerical values for a variety of parameters, and we begin with the wage density function. The US Social Security Administration reports that average individual wage earnings in 2018 were \$52,146. Likewise the maximum taxable wage income for 2018 was \$128,400, and over the last few decades the share of workers who earn more than the taxable maximum has held fairly steady at slightly above 5% (SSA 2013). These facts give us quantitative targets to calibrate the model density function $g(w)$.

We assume wages follow a truncated beta distribution with density

$$g(w) = \frac{w^{\gamma-1}(1-w)^{\beta-1}}{\int_0^1 w^{\gamma-1}(1-w)^{\beta-1}dw}.$$

We make the (harmless) assumption that the maximum wage in the economy is twice the taxable wage cap \$256,800, which is normalized to model wage $w = 1$. We need to pick a maximum wage in order to calibrate the distribution, but what is relevant to our analysis is the share of people above the tax cap rather than the maximum wage per se, since all earners above the tax cap share the same tax liability.

With the maximum wage normalized to 1, the taxable wage cap in the model is $w_c^0 = 0.5$.

We select (γ, β) to approximate two targets. First, we calibrate the mean wage to $\$52,146/\$256,800 = 0.2031$. The mean wage from the beta distribution is $\mathbb{E}(w) = \gamma/(\gamma+\beta)$, and we use this fact to ensure that, for any choice of β, the parameter γ is chosen such that $\mathbb{E}(w) = 0.2031$. That is, $\gamma = 0.2031\beta/(1-0.2031)$. Second, we can then choose β to match the share of individuals above the taxable wage cap (5%): $\int_{w_c^0}^1 g(w)dw = 0.05$. Doing this yields $\beta = 4.8810$ and $\gamma = 1.2437$. The calibrated wage density appears in Figure 1. This density has 5% of individuals above the taxable wage cap, and the mean and taxable wage cap are in the correct positions relative to one another.[13]

We set the Social Security tax rate to the full employer and employee tax of 10.6% (the Old-Age and Survivors Insurance (OASI) tax rate since 1990), hence $\tau^0 = 0.106$. Current Social Security benefits (PIA) are a piecewise linear function of earnings,

$$b^0(w) = \begin{cases} wm_1, \text{ for } w \leq p_1, \\ p_1m_1 + (w - p_1)m_2, \text{ for } p_1 \leq w \leq p_2, \\ p_1m_1 + (p_2 - p_1)m_2 + (w - p_2)m_3, \text{ for } p_2 \leq w \leq p_3 = w_c^0, \\ p_1m_1 + (p_2 - p_1)m_2 + (p_3 - p_2)m_3 = b_c^0, \text{ for } w \geq w_c^0, \end{cases}$$

where p_1, p_2, and p_3 are the "bend points" and m_1, m_2, and m_3 are the slopes (marginal replacement rates) of the three distinct segments. Beyond the third bend point, the function is flat. We use a conventional estimate of the bend points relative to average wages, $p_1 = 0.2\mathbb{E}(w)$, $p_2 = 1.24\mathbb{E}(w)$, and $p_3 = 2.47\mathbb{E}(w)$ (as in, e.g., Ortiz (2014)).[14] The slopes are $m_1 = 90\%$, $m_2 = 32\%$, and $m_3 = 15\%$.

[13] Although our wage distribution was calibrated to match the mean wage and the share of the population above the cap, it provides a reasonably close fit to the lower tail of the actual wage distribution as well. For example, the 2018 Annual Statistical Supplement produced by the Social Security Administration provides data on the distribution of wage income, and in Series 4.B7 of this publication the distribution of wages bears positive skewness: the mean wage among earners below the tax cap is more than three times larger than the mode. Our calibrated wage distributions bears very similar skewness with the mean exceeding the mode by approximately a factor of 3 as well. Moreover, the variance of the wage distribution among those below the cap is very close in the data to that our calibrated wage distribution.

[14] When we calculate the 2018 relative bend points using the national average wage index, as reported by the Social Security Administration, we obtain bend points that are slightly different than those used by Ortiz and others. Nevertheless,

We assume individuals begin work at age 18 (model age 0), retire at age 67 (model age $t_R = 49$) and pass away no later than age 100 (model age $T = 82$). The baseline survival function $S^0(t)$ is calibrated to ensure that the ratio of workers to retirees is 3.3, which is approximately the average value in the US during the period 2000-2010, before Social Security tax revenues began to fall short of benefits paid. Although the ratio of workers to retirees has been falling since the inception of the Social Security program, we think of the period 2000-2010 as the "initial" period because budget deficits did not materialized until after the close of this decade. The survival function $S^0(t) = 1 - (t/T)^{2.58}$ produces of ratio of workers to retirees $\int_0^{t_R} S^0(t)dt / \int_{t_R}^{T} S^0(t)dt = 3.3$. The ratio of workers to retirees is projected to fall to 2.0 by the year 2075 according to the intermediate projections of the 2017 Social Security Trustee's Report. The worst case projection of the Trustee's report puts the ratio of workers to retirees at 2.0 in the year 2035. In any case, we think of 2.0 as the future, equilibrium ratio of workers to retirees, and the survival function $S^1(t) = 1 - (t/T)^{8.47}$ produces $\int_0^{t_R} S^1(t)dt / \int_{t_R}^{T} S^1(t)dt = 2.0$, holding fixed the retirement age t_R and maximum lifespan T. Figure 2 shows these survival functions.

Finally, we assume CRRA preferences, $u(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma = 3$. We set the rate of time preference $\rho = 0$ and the risk free interest rate to $r = 2.9\%$ according to Trustee projections.[15] We summarize our baseline calibration in Table 1.

---

the magnitude of the discrepancy is not large enough to affect our quantitative results in a material way.

[15] Our selection of the discount rate ρ does not end up being important for our main results. The policy reforms themselves do not in any way depend on this parameter, and the things that do depend on this parameter (wealth accumulation and welfare) are affected in predictable ways. For instance, increasing ρ has the expected effect of decreasing mean private saving at retirement and mean lifetime income (due to reductions in bequest income). Increasing ρ also increases wealth inequality at retirement. But all of our welfare rankings of the different reforms, across both Economies I and II, are invariant to the particular values of ρ that we consider.

Table 1. Summary of Baseline Calibration of Parameters

Wages:	
$g(w) = \frac{w^{\gamma-1}(1-w)^{\beta-1}}{\int_0^1 w^{\gamma-1}(1-w)^{\beta-1}}$, for $w \in [0,1]$	truncated beta pdf over wage type
$w_c^0 = 0.5$	taxable wage cap ($128,400 in 2018)
$\mathbb{E}(w) = 0.2031$	mean individual wage ($52,146 in 2018)
$\int_{w_c^0}^1 g(w)dw = 0.05$	share of individuals above taxable wage cap
$\gamma = \mathbb{E}(w)\beta/(1-\mathbb{E}(w)) = 1.2437$	calibrated to match mean wage
$\beta = 4.8810$	calibrated to match share above tax cap
Social Security:	
$\tau^0 = 10.6\%$	current law SS tax rate (excluding DI)
$b^0(w)$ piecewise continuous	current law SS benefit-earning formula
Demographics and misc:	
$T = 82$	normalized maximum lifespan (age 18 to age 100)
$t_R = 49$	length of career (age 18 to 67)
$S^0(t) = 1-(t/T)^{2.58}$	survival function yielding 3.3 workers:retirees
$S^1(t) = 1-(t/T)^{8.47}$	survival function yielding 2.0 workers:retirees
$\rho = 0$	discount rate (free parameter)
$\sigma = 3$	midpoint CRRA value from literature
$r = 0.029$	real interest rate from Trustee's Report

5. Quantitative Analysis

In this section we begin by quantifying the three Rawlsian policy reforms introduced above. Then we examine Economies I and II and quantify the effects of these three reforms on consumption, saving, and the distribution of wealth.

Briefly, the three policy options can be summarized as follows. Option 1 lowers the maximum benefit until the budget is balanced, with no changes in the tax rate or tax cap. This results in the bottom 20% of earners maintaining current-law benefit. Option 2 increases the payroll tax rate and then the maximum benefit is reduced until the budget is balanced. Policy makers can choose the fraction of low-wage earners to protect by selecting the appropriate tax increase. Option 3 changes the top two (marginal) replacement rates within the PIA formula without changing the tax rate or the tax cap. This approach changes the maximum benefit amount implicitly.

5.1. *Policy Option 1: Reduce Benefit Cap without Changing Taxes*

Policy Option 1 is designed to protect the benefit of as many poor individuals as possible without increasing the Social Security tax rate or the taxable maximum wage. Figure 3 shows the future benefit-earning rule, $b^1(w)$, which traces the current benefit rule, $b^0(w)$, up to the threshold value $\hat{w}$ and then continues as a flat line thereafter.

The future benefit-earning rule balances the budget without reducing benefits for anyone with wage earnings up to $\hat{w} = 0.069$, which is 34% of the mean economy wide wage. This policy amounts to protecting the benefits of the bottom 20% of the population. Thus the policy maker is able to preserve benefits for the poorest 20% of the population and also balance the budget under future demographics without an increase in the tax rate or the tax cap.

5.2. ***Policy Option 2: Reduce Benefit Cap and Increase Tax Rate***

Relative to Policy Option 1, the government can protect the benefits of an even larger share of low wage earners through a tax increase. Our model allows policy makers to quantify how many additional poor individuals can be protected for a given tax increase. For example, holding the taxable wage cap fixed at its current level, the benefits of all individuals with below average wage earnings can be preserved if the Social Security tax rate is set to 17.05%, that is $\tau^1(\mathbb{E}(w)) = 0.1705$. This would protect the benefits of the bottom 58% of the population. Likewise, the benefits of all individuals with earnings less than half the mean wage can be preserved if the Social Security tax rate is set to 12.55%, that is $\tau^1\left(\frac{\mathbb{E}(w)}{2}\right) = 0.1255$.

Figure 4 depicts the reforms described in the preceding paragraph, together with Policy Option 1 for comparison. Likewise, Table 2 summaries these potential reforms.

Table 2. Required Tax Rate to Protect a Given Fraction of Low Wage Earners

Share protected:	Highest earner protected:	Required tax rate:
bottom 20%	$0.34\mathbb{E}(w)$	10.6%
bottom 31%	$0.5\mathbb{E}(w)$	12.55%
bottom 58%	$\mathbb{E}(w)$	17.05%

5.3. ***Policy Option 3: Preserve Strict Earnings-Based Feature without Changing Caps***

We explore a particular version of Option 3 that is easy to communicate—we consider the example of protecting benefits up to the first bend point in the current US rule (which is 20% of the mean wage), and then we reduce the slopes of the other segments of the benefit-earning rule as needed to balance the budget without an increase in the tax rate or the tax cap. We find that the required adjustment is

$\delta = 0.17$. Hence, under this potential reform option, policy makers leave all bend points in place, they leave the tax rate and tax cap in place, and they leave the slope of the first segment at 90%. Then, to balance the budget they must make the following adjustments: reduce the slope of second segment from 32% down to $32\% \times 0.17 \approx 5\%$, reduce the slope of third segment from 15% down to $15\% \times 0.17 \approx 3\%$, and leave the slope of the benefit-earning rule at 0 beyond the tax cap. Figure 5 shows the benefit-earning rule under this reform option. Notice that benefits following the current rule up to the first bend point, and then they represent a scaled down version of the current rule thereafter, maintaining a strictly positive relationship between benefits and earning (up to the tax cap).

5.4. *Comparison of Replacement Rates*

Policy makers and analysts commonly speak of Social Security programs in terms of replacement rates, which is the fraction of taxable wage income replaced by Social Security benefits. Table 3 reports replacement rates for three different workers across the three policy options that we consider. The wage rate of the first worker is very small, $w = \epsilon$. The wage rate of the second worker is average, $w = \int_0^1 g(w) w dw$. And the wage rate of the third worker equals the taxable maximum, $w = w_c$.

All three policy options share the same replacement rate for the poorest segment of the population by definition. With each of these policies, the new benefit-earning rule traces the current benefit-earning rule up to some point of departure. The policies differ in the location of the point of departure and in the shape of the benefit-earning rule after departure. Options 1 and 3 clearly have similar effects on replacement rates across the wage distribution. Option 2 of course provides more generous replacement rates to average workers and workers at the cap (but not to the poorest workers). Understanding exactly how the taxes and replacement rates associated with these policy options affect capital accumulation, inequality, and welfare is an important topic that we discuss in the next section.

Table 3. Replacement Rates across Policy Options

	min wage, $w = \epsilon$	avg wage, $w = \int_0^1 g(w)wdw$	tax cap, $w = w_c$
Option 1	90%	22.5%	9.1%
Option 2	90%	43.6%	17.7%
Option 3	90%	22.4%	10.9%

5.5. *Comparing the Effects of Options 1-3 on Economies I and II*

The notion that mandatory public saving through Social Security tends to crowd out private saving, at least to some degree, is both intuitive and commonly studied in the theoretical and empirical Social Security literatures. It is commonly understood that individuals may reduce their own private saving in response to mandatory public saving, thereby somewhat reducing the effectiveness of mandatory saving. Moreover, at the macro level such crowding out of private saving reduces the aggregate capital stock and hence reduces the equilibrium level of bequest income that is passed on from generation to generation. Caliendo, Guo, and Hosseini (2014) show that this second, macroeconomic effect is large enough under certain conditions to completely eliminate the welfare gains from Social Security's longevity insurance role. Consistent with that paper and with common understanding about the implications of crowding out, the ultimate welfare effects of the various policy reform options that we introduce will also depend on the degree to which individuals substitute public saving for private saving and the degree to which Social Security affects equilibrium bequest income. In this paper we go beyond Caliendo, Guo, and Hosseini in an attempt to introduce a realistic degree of wage heterogeneity and a realistic, progressive Social Security program into the model. We seek to understand the distributional welfare effects of the different reform options in the context of a model that accounts for both the micro and macro crowding out effects mentioned above.

Specifically, we are interested in computing the following equilibrium objects across all reform options: first, mean private savings at retirement

$$\mathbb{E}(k) = \int_0^1 g(w)k(t_R|w)dw,$$

where $k(t_R|w)$ is equilibrium savings at retirement of a type w individual; second, wealth inequality at retirement

$$Gini_k = \frac{1}{2\mathbb{E}(k)} \int_0^1 \int_0^1 g(x)g(y)|k(t_R|x) - k(t_R|y)|dxdy;$$

third, mean lifetime income[16]

$$\mathbb{E}(Y) = \int_0^1 g(w)Y(w)dw,$$

where

$$Y(w) = \left(\int_0^{t_R} e^{-rt}(w - \tau \min(w, w_c) + B)dt + \int_{t_R}^{T} e^{-rt}(b(w) + B)dt \right);$$

fourth, lifetime income inequality

$$Gini_Y = \frac{1}{2\mathbb{E}(Y)} \int_0^1 \int_0^1 g(x)g(y)|Y(x) - Y(y)|dxdy;$$

and fifth, ex-ante expected utility (before individuals realize their wage type)

$$\mathbb{E}(U) = \int_0^1 \int_0^T g(w)e^{-\rho t}S(t)u(c(t|w))dtdw.$$

To facilitate utility comparisons across reforms, we also report the consumption equivalent (CEV) which

[16] We exclude interest income from the calculation of lifetime income. We are trying to understand the degree to which lifetime budget constraints differ across individuals in the model, and our calculation of income is the standard calculation of a lifetime budget constraint. Also, by discounting future cash flows at the market rate of interest in our calculation of lifetime income, we are already accounting for intertemporal saving opportunities.

is the fraction of lifetime consumption an individual is willing to give up (ex ante) to live in a world with Policy Option 1 instead of Options 2 or 3,

$$CEV = 1 - \left(\frac{\mathbb{E}(U_i)}{\mathbb{E}(U_1)} \right)^{\frac{1}{1-\sigma}} \text{ for } i = \{2, 3\}.$$

Table 4 reports the results of our calculations. Panel A of Table 4 corresponds to Economy I and Panel B corresponds to Economy II. We will begin by focusing our discussion on Economy I.

Recall that the taxable wage cap w_c^0 is not altered in any of our reform options. In the case of Policy Option 1, the tax rate stays at the current rate of 10.6% and the new benefit-earning rule follows the current benefit-earning rule up to $0.34\mathbb{E}(w)$, which covers the bottom 20% of wage earners, and then the benefit-earning rule becomes flat thereafter. For Policy Option 2, we consider the case of a new tax rate of 17.05%, which is enough to keep the new benefit-earning rule equal to the current benefit-earning rule up to the average wage $\mathbb{E}(w)$, which covers the bottom 58% of wage earners, and then the benefit-earning rule becomes flat thereafter. And for Policy Option 3, the tax rate stays at 10.6% and the new benefit-earning rule follows the current benefit-earning rule up to the first bend point, and then the slope of the second segment drops from 32% to 5% and the slope of the third segment drops from 15% to 3%. The location of the bend points themselves ($0.2\mathbb{E}(w)$, $1.24\mathbb{E}(w)$, and $2.47\mathbb{E}(w)$) stay the same.

Policy Options 1 and 3 have virtually the same effect on capital accumulation, and have similar effects on inequality and welfare. This is not too surprising since these two options are mechanically similar. They both leave the benefit-earning rule completely intact for a portion of the left tail of the wage distribution and then either flatten out from there (Option 1) or retain a slight upward slope (Option 3); and neither option comes with a tax increase. Option 2 involves a significant tax increase which has a large crowding out effect on capital accumulation and hence a large reduction in bequest income received. Because bequest income is redistributive in the sense that all individuals, regardless of wage type, receive the same bequest (by assumption), the crowding out of bequest income occurring through this reform option is particularly painful as it weakens a natural insurance mechanism already at play in the economy. This in turn significantly reduces expected utility in Option 2 relative to the other options. In other words, in effort to protect a larger share of low-wage earners than can be protected in Options 1 and 3, the extra tax revenues required by Option 2 have the unintended consequence of unwinding private saving and hence unwinding the insurance that occurs naturally through the transmission of wealth from one generation to another. Hence, according to the model, policy makers may be better off pursuing Rawlsian reforms that can be financed without additional taxation (such as Options 1 and 3) rather than Rawlsian reforms that are financed through large increases in taxation (such as Option 2).

We view this set of results as containing potentially important lessons for policy making. Notice that the tax increases needed to finance Option 2 are undesirable even without considering the potential negative labor supply effects associated with higher taxation. The potential distortionary effect of Social Security taxation on labor supply is well known (among many examples, see Bagchi (2015) for recent computational experiments), and the debate about Social Security's effect on private saving is decades old (see Slavov, Gorry, Gorry, and Caliendo (2019) for a review). What is new here is the idea that Social Security's crowding out of private saving can interfere with the redistributive transmission of wealth across generations, because this mechanism not only redistributes wealth from the deceased to the living but also from the rich to the poor (since poor workers get a bequest that is large relative to their wage type). Such "family insurance" considerations have important welfare implications that could affect the performance of reform options that are intended to help the poor. Indeed, in terms of consumption equivalent units, the ex-ante welfare gains associated with Option 1 over Option 2 are very large (14.1%). This is because the crowding out of bequest income associated with higher taxes on wages is especially painful to those at the low end of the wage distribution who are almost totally reliant on such intergenerational transfers.[17]

[17] It is worth mentioning that the income inequality that we measure in our model is not intended to capture all of the inequality in income appearing in the real world. Our model captures just the inequality that is associated directly with

Turning to Economy II in Panel B, the first moments of capital accumulation $\mathbb{E}(k)$ and lifetime income $\mathbb{E}(Y)$ are the same in the two economies. This is intuitive since the distribution of wage income is fixed across the two economies and hence aggregate income or resources are the same as well. Economy II has more inequality though, and this is expected as well since we have introduced additional inequality by assuming that bequest income is transmitted in a dynastic manner (from rich to rich and from poor to poor). The main lesson that we wish to highlight from Panel B is that the welfare ranking of our three Rawlsian Social Security reforms remains the same in Economies I and II. Finally, we learn from our comparison of certainty equivalents across Economies I and II that the relative welfare cost of increasing the Social Security tax rate is exaggerated in Economy I. In both economies, an increase in the Social Security tax rate crowds out equilibrium bequest income, but in Economy II bequest income to the poor is already very low because it flows from ancestors who also were poor, and so there is less to crowd out.[18]

Of course, our model is a simplified version of reality and therefore is an imperfect instrument to fully assess the welfare effects of Social Security reform. Individuals in our model make rational, forward-looking decisions to maximize their utility, and they do so in the face of full information about the risks that they face. While this follows in the neoclassical tradition, it might be interesting in future work to extend our analysis to include the case of boundedly rational agents who make decisions without access to capital markets or who lack information about the distribution of risks that they face. However, while we have explored only one set of such assumptions, we do expect that the *distributional mechanisms* at play in the present paper will continue to be vitally important to a full understanding of the welfare effects of Social Security reform.

---

wage inequality. For example, in reality people experience different rates of return on their investments and this could be systematically correlated with wage earnings, and differences in utilization of capital markets altogether would also affect income inequality. Our model also abstracts from many other issues that affect the level of income inequality such as inter vivos transfers within families, utilization of insurance markets, and multiple-earner households among other things.

[18]To disaggregate our results, we examine the instantaneous utility of consumption for individuals with different wage types across the entire life cycle for each reform in both economies. The analysis confirms our main result: individuals in either economy prefer Policy Options 1 and 3 to Policy Option 2 (which raises taxes). This is true for a worker in the bottom, second, third, or top income quartile, for the young, and for the old. Across the life cycle, individuals of any wage type experience the lowest instantaneous utility with Policy Option 2.

Table 4. The Effects of Policy Reform Options on Savings, Inequality, and Welfare

Panel A: Economy I

Opt:	new tax:	new benefit-earning rule:	B	$\mathbb{E}(k)$	$Gini_k$	$\mathbb{E}(Y)$	$Gini_Y$	$\mathbb{E}(U)$	CEV
1	10.6%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq 0.34\mathbb{E}(w) \\ b^0(0.34\mathbb{E}(w)) \text{ for } w > 0.34\mathbb{E}(w) \end{cases}$	0.014	3.57	0.47	5.41	0.363	-6843	—
2	17%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq \mathbb{E}(w) \\ b^0(\mathbb{E}(w)) \text{ for } w > \mathbb{E}(w) \end{cases}$	0.011	2.83	0.49	5.10	0.369	-9268	14.1%
3	10.6%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq 0.2\mathbb{E}(w) \\ 0.17b^0(w) + 0.83b^0(0.2\mathbb{E}(w)) \text{ for } w > 0.2\mathbb{E}(w) \end{cases}$	0.014	3.57	0.46	5.41	0.365	-6875	0.2%

Panel B: Economy II

Opt:	new tax:	new benefit-earning rule:	$\mathbb{E}(B)$	$\mathbb{E}(k)$	$Gini_k$	$\mathbb{E}(Y)$	$Gini_Y$	$\mathbb{E}(U)$	CEV
1	10.6%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq 0.34\mathbb{E}(w) \\ b^0(0.34\mathbb{E}(w)) \text{ for } w > 0.34\mathbb{E}(w) \end{cases}$	0.014	3.57	0.48	5.41	0.402	-11423	—
2	17%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq \mathbb{E}(w) \\ b^0(\mathbb{E}(w)) \text{ for } w > \mathbb{E}(w) \end{cases}$	0.010	2.81	0.50	5.09	0.403	-13070	6.5%
3	10.6%	$b^1(w) = \begin{cases} b^0(w) \text{ for } w \leq 0.2\mathbb{E}(w) \\ 0.17b^0(w) + 0.83b^0(0.2\mathbb{E}(w)) \text{ for } w > 0.2\mathbb{E}(w) \end{cases}$	0.014	3.57	0.47	5.41	0.403	-11452	0.1%

6. Conclusion

The Senate Aging Committee and the Government Accountability Office (GAO) both encourage Congress to take a Rawlsian perspective when evaluating Social Security reform measures. They believe a desirable reform should not only balance the budget, but it should also protect benefits for the economically vulnerable. Yet the GAO acknowledges that "time constraints did not allow [them] to undertake the complex analysis necessary to develop quantitative estimates."[19] We have attempted to provide policy makers with a menu of quantitative adjustments to the benefit-earning rule that are consistent with their stated preferences for taking care of the poor.

We have built a theoretical model with a calibrated wage density function that captures the degree of heterogeneity in the earnings of US workers. Including heterogeneity allows us to move beyond the limited set of responses to a demographic shock that would arise from a representative agent model, toward a richer set of recommendations that focus on distributional concerns. Throughout our analysis we imposed the discipline of a balanced-budget constraint to distinguish between Rawlsian policy responses that cost extra taxes to achieve and Rawlsian responses that can be financed without tax increases. We have focused on three Rawlsian adjustments to the current benefit-earning rule that are naturally easy to understand and technically easy to implement.

Two of our proposals break the link between the taxable maximum and the benefit cap by reducing the benefit cap. Social Security payroll taxes apply only to earnings below a threshold amount, often referred to as the taxable maximum. Similarly, benefits are based on earnings up to a threshold. Currently, the threshold for taxes is the same as the threshold for benefits. We show that the policy maker can protect the benefits of the poorest households by breaking that link and reducing the maximum benefit. This is equivalent to reducing the slope of the benefit-earning rule to zero above the benefit threshold. Our proposals are similar in spirit to many current proposals that either reduce the slope of the benefit-earning

[19] Quote from Government Accountability Office document: GAO-10-1010R.

rule or increase the tax cap. Our proposals are mechanically distinct by breaking the connection between the benefit cap and tax cap. For comparison, we also provide a proposal that maintains the link between the tax cap and benefit cap. We show that the policy maker can protect benefits for a larger share of the population by breaking the link between the tax cap and the maximum benefit than by maintaining a benefit-earning rule with a strictly positive slope.

By using a static model in our analysis we have abstracted from two potential economic responses to the demographic crisis. First, increased longevity tends to increase the optimal age of retirement in life-cycle models. Second, increased longevity tends to increase capital accumulation (due to the need to finance consumption during a longer retirement period), which in turn puts upward pressure on the marginal product of labor and hence wages. Both effects may lead to an expansion of the tax base, which partly offsets the negative budgetary effects of increased longevity.[20] If these economic feedback effects are relevant, then we have overstated the magnitude of the crisis and policy makers could achieve Rawlsian goals with less severe adjustments than we have discussed. However, it is potentially risky for policy makers to count on increased tax revenues through longer careers and higher wages. While these effects may indeed result from an increase in longevity in closed economy with rational expectation and rational behavior, reality may be more complex. And if neither of these things end up resulting from increased longevity and there is therefore no boost to the tax base, then policy makers would be well served by making plans according to a (less-optimistic) model such as ours that abstracts from these potential benefits.

Finally, we have made a simple comparison of policy options in a future steady state. This is a non-trivial simplification that does not provide any instruction on how policy makers might tackle key issues such as grandfathering along the transition path. In other words, if policy makers wish to pursue a Rawlsian reform like those discussed above, they would still need to work out difficult issues concerning

[20] See Bagchi (2016).

the timing of when and for whom such policy changes would be implemented. A recent set of papers suggests that whatever option policy makers pursue and whenever they wish to implement those changes, simply sharing that plan with the public will reduce costly distortions to consumption and saving that arise from uncertainty about the timing and structure of reform (Bütler (1999), Kitao (2018), Caliendo, Gorry, and Slavov (2019), Cottle Hunt (2019), and Nelson (2017)).

7. References

1. Attanasio, O., Kitao, S. & Violante, G. (2007). Global demographic trends and social security reform. *Journal of Monetary Economics*, 54(1), 144-198.

2. Bagchi, S. (2015). Labor supply and the optimality of social security. *Journal of Economic Dynamics and Control*, 58(3), 167-185.

3. Bagchi, S. (2016). Is the social security crisis really as bad as we think? *Macroeconomic Dynamics*, 20(3), 737-76.

4. Bagchi, S. (2017). Can removing the tax cap save social security? *B.E. Journal of Macroeconomics*, 17(2), 1-28.

5. Bipartisan Policy Center Commission on Retirement Security and Personal Savings (2016). Securing our financial future: Report of the commission on retirement security and personal savings.

6. Bütler, M. (1999). Anticipation effects of looming public-pension reforms. *Carnegie- Rochester Conference Series on Public Policy*, 50, 119-159.

7. Caliendo, F., Casanova, M., Gorry, A., & Slavov, S. (2019). Retirement timing uncertainty: Empirical evidence and quantitative evaluation. Utah State University working paper.

8. Caliendo, F., Gorry, A., & Slavov, S. (2019). The cost of uncertainty about the timing of social security reform. *European Economic Review*, 118, 101-125.

9. Caliendo, F., Guo, N., & Hosseini, R. (2014). Social security is NOT a substitute for annuity markets. *Review of Economic Dynamics*, 17, 739-755.

10. Cottle Hunt, E. (2019). Life-cycle horizon learning, social security reform, and policy uncertainty. Lafayette College working paper.

11. Conesa, J.C., & Garriga, C. (2008). Optimal fiscal policy in the design of social security reforms. *International Economic Review* , 49(1), 291-318.

12. Conesa, J. C. & Krueger, D. (1999). Social security reform with heterogeneous agents. *Review of Economic Dynamics*, 2(4), 757-795.

13. Congressional Budget Office (2015). Social security policy options. Report 51011.

14. Congressional Budget Office (2018). Options for reducing the deficit: 2019 to 2028. Report 54667.

15. Cremer, H. & Pestieau, P. (2011). Myopia, redistribution, and pensions. *European Economic Review*, 55(2), 165-175.

16. Cremer, H., De Donder, P., Maldonado, D., & Pestieau, P. (2008). Designing a linear pension scheme with forced savings and wage heterogeneity. *International Tax and Public Finance*, 15(5), 547-562.

17. Diamond, P. (2018). The future of social security. *Economic Inquiry*, 56(2), 661-681.

18. Diamond, P. & Orszag, P. (2005). Saving social security. *Journal of Economic Perspectives*, 19(2), 11-32.

19. De Nardi, M., İmrohoroğlu, S., & Sargent, T. (1999). Projected US demographics and social security. *Review of Economic Dynamics*, 2(3), 575-615.

20. Feldstein, M. (1985). The optimal level of social security benefits. *Quarterly Journal of Economics*, 100(2), 303-320.

21. Feldstein, M. (1996). The missing piece in policy analysis: Social security reform. *American Economic Review* 86(2), 1-14.

22. Gomes, F., Kotlikoff, L., & Viceira, L. (2012). The excess burden of government indecision. *Tax Policy and the Economy*, 26(1), 125-164.

23. Government Accountability Office (2010). Social security: Options to protect benefits for the vulnerable groups when addressing program solvency. Report GAO-10-101R.

24. Huggett, M. & Ventura, G. (1999). On the distributional effects of social security reform. *Review of Economic Dynamics*, 2(3), 498-531.

25. İmrohoroğlu, S., & Kitao, S. (2009). Labor supply elasticity and social security reform. *Journal of Public Economics*, 93, 867-878.

26. İmrohoroğlu, S., & Kitao, S. (2012). Social security reforms: Benefit claiming, labor force participation and long-run sustainability. *American Economic Journal: Macroeconomics*, 4 (3), 96-127.

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

28. Kitao, S. (2018). Policy uncertainty and cost of delaying reform: The case of aging Japan. *Review of Economic Dynamics*, 27, 81-100.

29. Kotlikoff, L., Smetters, K., & Walliser, J. (1999). Privatizing social security in the United States – comparing the options. *Review of Economic Dynamics*, 2(3), 532-574.

30. Kotlikoff, L., Smetters, K., & Walliser, J. (2007). Mitigating America's demographic dilemma by pre-funding social security. *Journal of Monetary Economics*, 54(2), 247-266.

31. Luttmer, E., & Samwick, A. (2018). The welfare cost of perceived policy uncertainty: Evidence from social security. *American Economic Review*, 108(2), 275-307.

32. Nelson, J. (2017). Welfare implications of uncertain social security reform. CBO working paper.

33. Nishiyama, S. & Smetters, K. (2007). Does social security privatization produce efficiency gains? *The Quarterly Journal of Economics*, 122(4), 1677-1719.

34. Ortiz, J. (2014). Social security and retirement across OECD countries. *Journal of Economic Dynamics and Control*, 47, 300-316.

35. Samwick, A.(1998). Discount rate heterogeneity and social security reform. *Journal of Development Economics*, 57(1), 117-146.

36. Slavov, S., Gorry, D., Gorry, A. & Caliendo, F. (2019). Social Security and saving: an update. *Public Finance Review*, 47(2), 312-348.

37. Social Security Administration (2013). Annual Statistical Supplement, 2013. Report 13-11700, Table 4.B1.

38. Social Security Administration (2017). Summary of provisions that change the social security program. *Office of the Chief Actuary of the Social Security Administration.*

39. Special Committee on Aging (2010). Social security modernization: Options to address solvency and benefits adequacy. *United State Senate. 111th Congress, 2nd session*, Report 111-187.

40. Yaari, M. (1965). Uncertain lifetime, life insurance and the theory of the consumer. *Review of Economic Studies*, 32(2), 137-150.

Figure 1. Wage Density Function

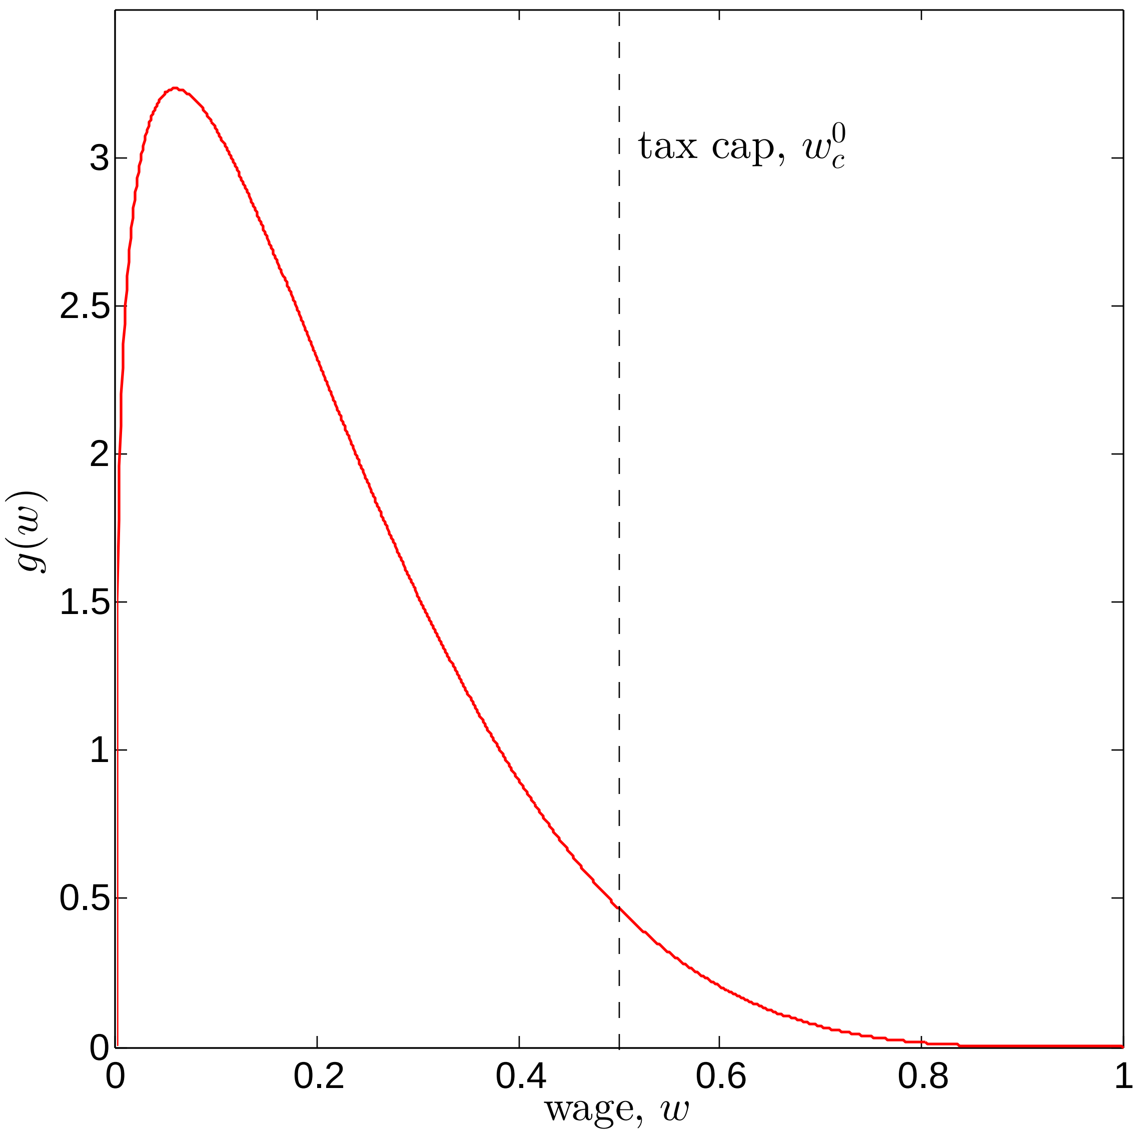

Figure 2. Survival Functions (Current and Projected)

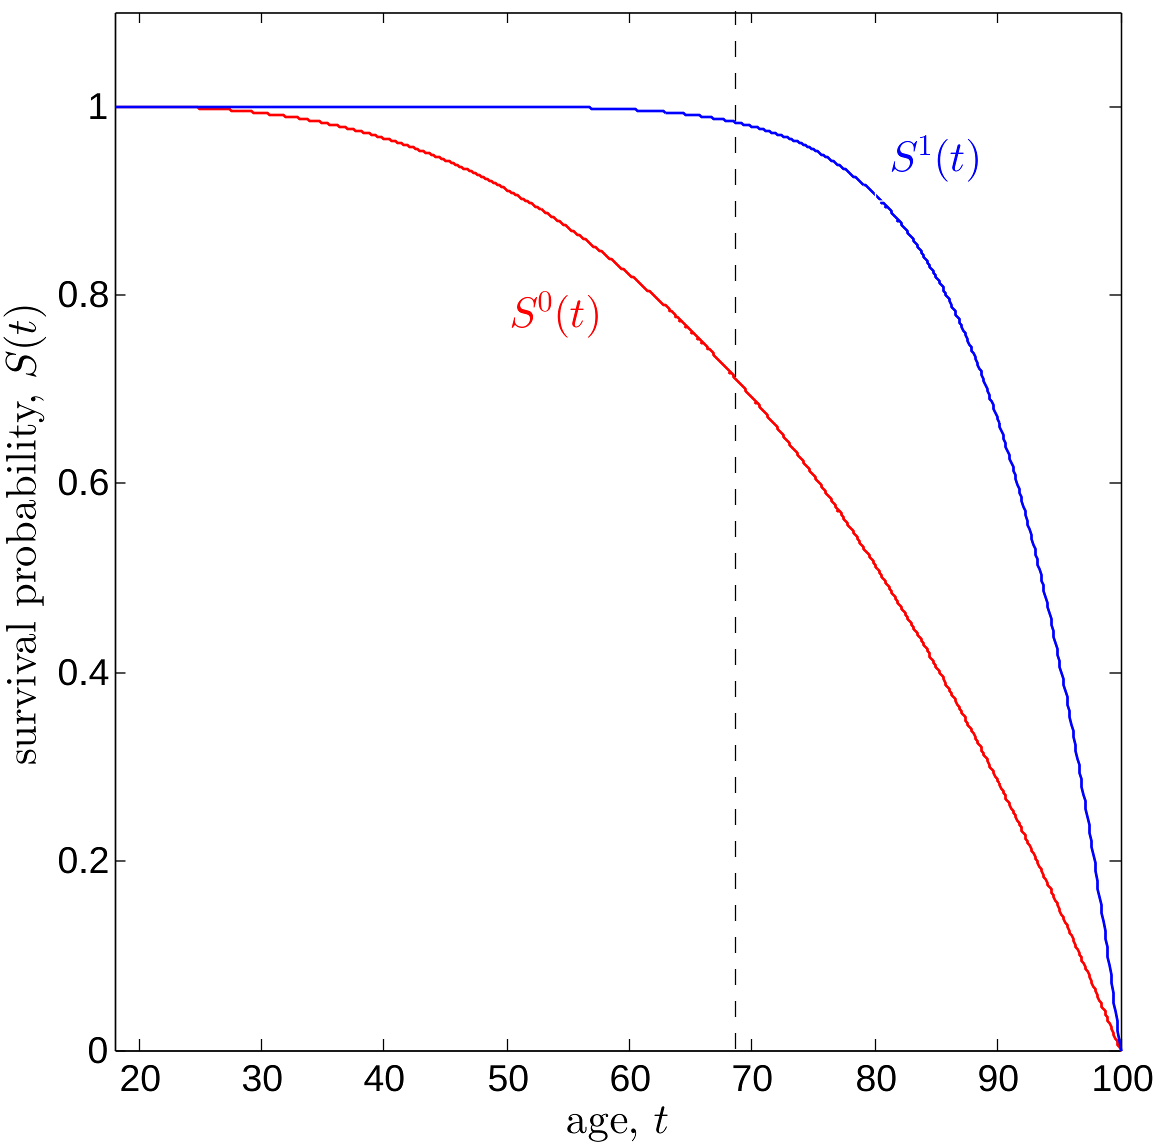

Figure 3. Policy Option 1: Reduce Benefit Cap without Changing Taxes

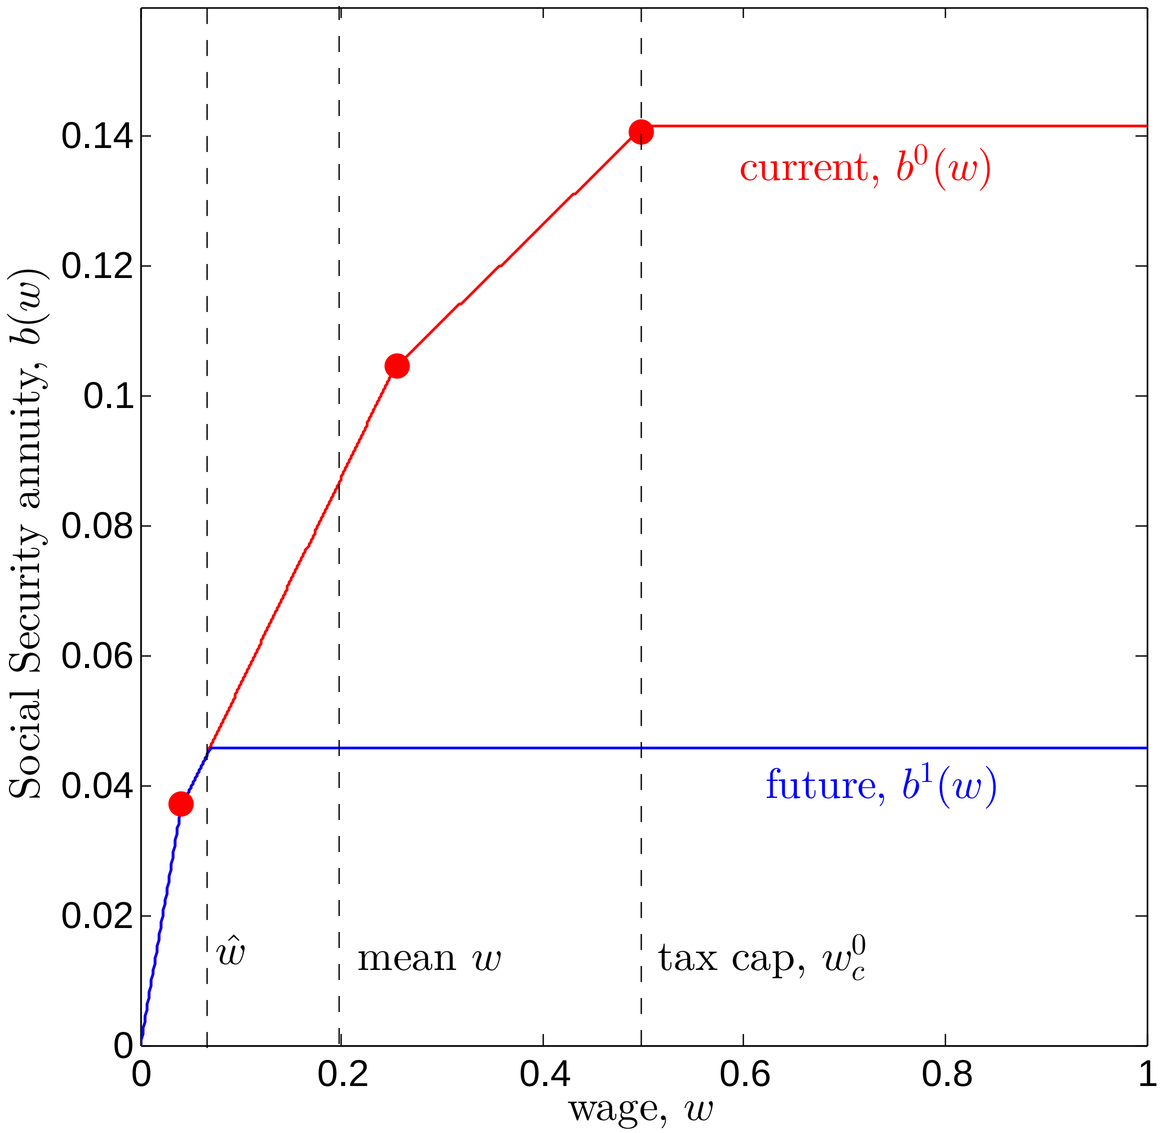

Figure 4. Policy Option 2: Reduce Benefit Cap and Increase Tax Rate

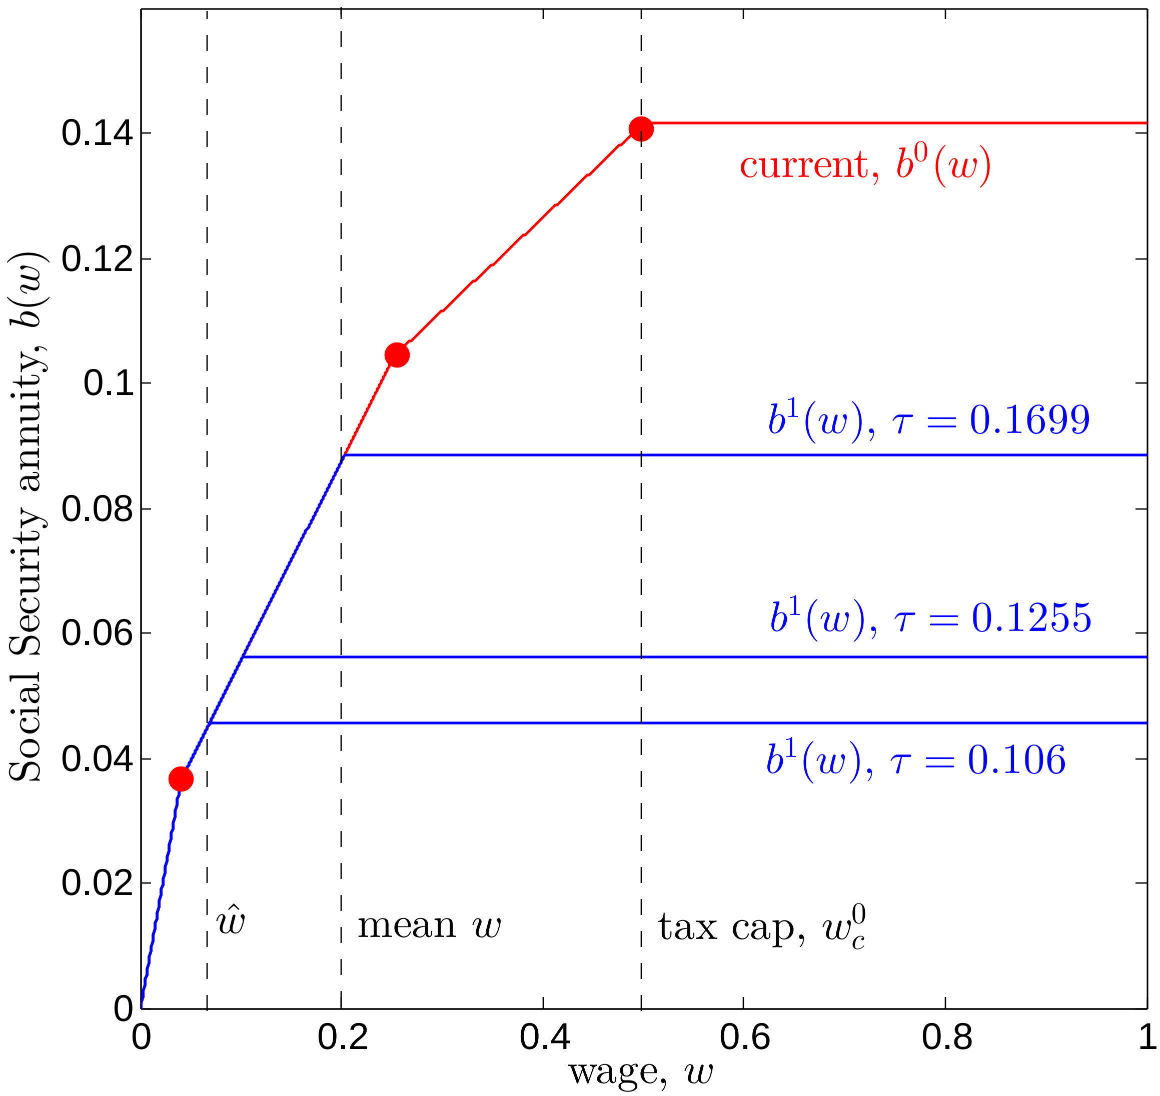

Figure 5. Policy Option 3: Maintain Strict Benefit-Earning Link