Towson University
Department of Economics
Working Paper Series

Working Paper No. 2024-11

Business Cycles and Public Pensions: Aggregate Risk and Social Security in the United States

By Shantanu Bagchi

September 2024

Business Cycles and Public Pensions: Aggregate Risk and Social Security in the United States*

Shantanu Bagchi†

August 29, 2024

Abstract

This paper uses a stylized overlapping-generations model to examine the effect of aggregate (or business cycle) risk on the macroeconomic and welfare implications of Social Security. In this model framework, unfunded public pensions provide partial insurance against inter- and intra-generational risks that are uninsured due to incomplete markets. I find that in this environment, Social Security's macroeconomic and welfare effects are considerably smaller than those in a framework without aggregate risk, and that the persistence of the aggregate shock process is an important determinant of this difference. I also find that aggregate risk changes how the redistribution implicit in Social Security's benefit-earnings rule interacts with its intergenerational risk sharing mechanism.

JEL Classifications: E21, E62, H55

Keywords: Social Security; aggregate risk; business cycles; incomplete markets; intergenerational risk

1 Introduction

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

---

[*] I would like to thank Frank Caliendo, Jim Feigenbaum, Juergen Jung, Kai (Jackie) Zhao, and participants at the CEF 2024 conference 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

Unfunded public pensions also facilitate the sharing of a third type of risk – intergenerational risk – because they establish a financial contract between current workers and *future* workers. In this sense, unfunded pensions are current workers' claims to *future* labor income, which is uncertain because of aggregate productivity (or business cycle) risk. While this role of public pensions is well known, there is a surprising paucity of research on how aggregate productivity risk interferes with a public pension system's ability to fulfill this role. Specifically, in the context of the U.S., almost the entire literature on public pensions is based on models without any aggregate risk to be shared between generations.[1] This is partly due to the fact that overlapping–generations models with aggregate risk are notoriously difficult to solve because of the (general) non–existence of steady–state wealth distributions. However, the absence of aggregate risk models from this discussion has likely led to an incomplete accounting of Social Security's insurance effects, particularly those related to current workers' claims to future labor income.[2]

The goal of this paper is to quantitatively examine the effect of aggregate risk on the macroeconomic and welfare implications of unfunded public pensions – U.S. Social Security in particular. To do this, I begin by designing a stylized overlapping-generations macroeconomic model with the following microeconomic building blocks: rational utility-maximizing households, profit-maximizing firms, and a government that provides public goods and Social Security. Households in the model survive for a maximum of three periods, and they vary in their earning ability depending on their human capital (education) level. Aggregate productivity is stochastic, and it affects overall factor returns. Most notably, this stochastic aggregate productivity leads to a general non-existence of the model's steady-state wealth distribution, so households have to forecast future macroeconomic aggregates based on their current values. I calibrate this model to match U.S. macroeconomic targets, and then compute the effect of aggregate risk on the macroeconomic and welfare implications of a 50% cut in Social Security's payroll tax rate.

In general, my findings suggest that introducing aggregate risk in this otherwise standard model framework has quantitatively important implications for Social Security's macroeconomic and welfare effects. First, downsizing Social Security in the presence of aggregate risk, on the average, causes a smaller increase in capital stock, labor supply, output, and aggregate consumption. Because of this reason, the overall welfare gains from the tax cut are smaller in the presence of aggregate risk, especially when the economy is in an unfavorable initial productivity state. I also find that the persistence of the aggregate shock process is a key determinant of the importance of current workers' claims to future labor income. With a shock process only half as persistent as the baseline, the macroeconomic and welfare effects of downsizing Social Security are much closer to those observed in the absence of aggregate risk, but in this case the smaller effects occur when the economy is initially in a favorable productivity state. Finally, in the presence of aggregate risk, the welfare gains from downsizing Social Security follow a hump-shaped pattern with respect to earnings. This hump shape is even more pronounced under a hypothetical linear benefit-earnings rule, which suggests that aggregate risk also changes how the implicit financial contract between current workers with varying labor-market outcomes interacts with its intergenerational risk sharing mechanism.

Starting with Abel (1985) and Hubbard and Judd (1987), a large literature has evolved to examine the importance of the different roles of unfunded public pensions in justifying the size of U.S. Social Security. Both Abel (1985) and Hubbard and Judd (1987) find a welfare-improving role for Social Security in a model with mortality risk and closed annuity markets, but Hubbard and Judd (1987) find that these welfare gains are significantly reduced or even eliminated when there

---

[1] Two notable exceptions are Krueger and Kubler (2006) and Harenberg and Ludwig (2015).

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

are borrowing constraints. In a related study, İmrohoroğlu et al. (1995) examine the optimality of Social Security in a life-cycle economy with mortality risk, missing annuity markets, idiosyncratic employment risk, and borrowing constraints. They find that the optimal social security arrangement features a replacement rate of 30% and a tax rate of 6.1%. While this literature does not arrive at a consensus regarding the optimal size of Social Security in the U.S., it generally concludes that the welfare-improving role of Social Security is much smaller once the consumption and labor supply distortions from Social Security are accounted for. However, almost all of these studies are based on models without any aggregate risk to be shared between generations, which potentially underestimates the importance of Social Security's implicit financial contract between current and *future* workers.

Two notable exceptions in this literature are Krueger and Kubler (2006) and Harenberg and Ludwig (2015). Krueger and Kubler (2006) were the first to seriously investigate the welfare consequences of introducing an unfunded public pension system in an overlapping-generations model with stochastic production and incomplete markets. They find that the introduction of such a program is Pareto-improving only in partial equilibrium, and that the severity of the capital crowding-out effect in general equilibrium reverses the welfare gains. In other words, their findings mirror those of studies conducted using models without aggregate risk. Similarly, Harenberg and Ludwig (2015) show analytically that under incomplete markets, a public pension system can partially insure against idiosyncratic and aggregate risks, but only in partial equilibrium.[3] However, neither of these studies weigh in on the *marginal* effect of aggregate risk on their welfare results, which is what this paper attempts to do.

Finally, it is worth mentioning Ríos-Rull (1996) as one of the earliest studies to consider the implications of aggregate risk in an overlapping-generations framework. Ríos-Rull (1996) considers if the quantitative implications of business cycles change when they are studied using models with sophisticated demographic structures and finitely-lived agents, rather than infinite-horizon representative-agent models. Ríos-Rull (1996) concludes that the implications are basically the same for the two kinds of models, but also that the overlapping-generations framework does a better job of explaining the relative volatility of hours across age groups, something that cannot be measured in a infinite-horizon representative-agent model. However, Ríos-Rull (1996) is silent on the implications of public insurance, because their model does not consider government expenditures and taxes.

The rest of the paper is organized as follows. I discuss the specific details of the stylized model in Section 2, explain the equilibrium computation method in Section 3, and I detail the calibration approach in Section 4. I report the primary quantitative results in Section 5, explain the underlying mechanisms behind the results in Section 6, and I conclude in Section 7.

2 The Model

In this section, I develop a stylized model framework that features both intra- and intergenerational risk in an incomplete markets environment. This model framework, although simplistic, captures the essential elements that allow us to study Social Security's implicit contracts: between current workers and current retirees, between current workers (or future retirees), and between current workers and *future* workers. I summarize these elements below.

The unit of the model is a life-cycle permanent-income household that survives for a maximum of

---

[3] It is worth noting that while the two-period model framework in Harenberg and Ludwig (2015) is useful for analytical tractability, it ignores an important channel of intergenerational risk sharing facilitated by Social Security: the implicit financial contract between current *young* and future *old* workers.

three periods, in which the first two are "work" periods where utility is derived from consumption and leisure, and the last period is "retirement" where only consumption provides utility.[4] This household experiences three types of risk over the course of the life cycle:

- labor income risk: the household draws a human-capital (education) fixed effect prior to entering the model,
- mortality risk: the household faces a non-zero risk of dying every period, and
- aggregate risk: aggregate labor productivity is uncertain, because of which returns to capital and labor each period are risky.

It is important to note that all the three types of risk are uninsurable, i.e. households do not have access to markets to privately insure against these risks, nor do they have access to riskless financial instruments (such as bonds). During the "retirement" period, a surviving household receives Social Security benefits, which depend on their "work" period contributions. Lifetime expected utility is given by

$$U = E\left[Q_0\ u\left(c_0, \ell_0\right) + \beta Q_0 Q_1\ u\left(c_1, \ell_1\right) + \beta^2 Q_0 Q_1 Q_2\ u\left(c_2, 1\right)\right], \tag{1}$$

where Q_j is the probability at age j of surviving to the next period, and β is the discount factor. The respective period budget constraints are

$$\begin{aligned}
\text{Period 0:}\quad c_0 + k_1 &= \left[(1-\ell_0)W_0 e_{0,\theta} - T_y\left\{(1-\ell_0)W_0 e_{0,\theta})\right\}\right] - T_{SS}\left\{(1-\ell_0)W_0 e_{0,\theta})\right\} && (2)\\
\text{Period 1:}\quad c_1 + k_2 &= \widetilde{R_1}k_1 + \Big[\left(\widetilde{R_1}-1\right)k_1 + (1-\ell_1)\widetilde{W_1}e_{1,\theta} \\
&\quad - T_y\left\{\left(\widetilde{R_1}-1\right)k_1 + (1-\ell_1)\widetilde{W_1}e_{1,\theta}\right\}\Big] - T_{SS}\left\{(1-\ell_1)\widetilde{W_1}e_{1,\theta})\right\} && (3)\\
\text{Period 2:}\quad c_2 &= \widetilde{R_2}k_2 + \left(\widetilde{R_2}-1\right)k_2 - T_y\left\{\left(\widetilde{R_2}-1\right)k_2\right\} + \widetilde{B}\left(\bar{y}\right), && (4)
\end{aligned}$$

where $\widetilde{W_j}$ and $\widetilde{R_j}$ are stochastic wages and the gross rates of return, $e_{j,\theta}$ are deterministic labor productivity endowments that vary with age j and a human capital fixed effect θ, and $\widetilde{B}(\cdot)$ is a stochastic Social Security benefit that depends on a measure of past income $\bar{y}$. The tax functions $T_y(\cdot)$ and $T_{SS}(\cdot)$ respectively denote general labor plus capital income taxes and the Social Security payroll tax. Note that from the perspective of an age-0 household, the period 0 wage and the rate of return is deterministic because the period-0 aggregate shock has already been realized. Finally, because of mortality risk, a finite number of age-0 and age-1 agents die with unused assets. I assume that the government imposes a confiscatory tax on these assets, which also become a part of the general tax revenues.

Firms operate competitively and produce output using capital, labor, and a constant returns to scale technology with uncertain aggregate labor productivity, given by

$$Y_t = F\left(K_t,\ A_t L_t\right) \tag{5}$$

$$A_t = A_0(1+g_A)^t \times A_t^\epsilon \tag{6}$$

$$\ln A_t^\epsilon = \rho \ln A_{t-1}^\epsilon + \sigma_A \epsilon_t \tag{7}$$

where K_t and L_t are aggregate capital and labor, g_A is the productivity growth rate, ρ and σ_A are respectively the persistence and variance of the aggregate productivity shock, and ϵ_t is a white

---

[4] It should be noted that a two-period model with only one "work" period, such as the one used by Harenberg and Ludwig (2015), fails to account for Social Security's implicit contract between current *young* and future *old* workers.

noise error term. Competitive factor markets imply that

$$\widetilde{W}_t = \widetilde{MP}_L \tag{8}$$

$$\widetilde{R}_t = \widetilde{MP}_K + 1 - \delta, \tag{9}$$

where α is capital's share in output, and δ is the depreciation rate.

The government provides public goods and Social Security; the public goods purchases are funded using the revenues from the general labor and capital income taxes, and Social Security is funded through the payroll tax on labor income on a Pay-As-You-Go (PAYG) or unfunded basis. As mentioned above, Social Security facilitates three different implicit financial contracts in this framework – between current workers and current retirees, between current workers (or future retirees), and between current workers and *future* workers. The general government and Social Security budget constraints at any given point of time are

$$\sum_{j=0}^{2} N_j \; T_y \left\{ \left(\widetilde{R}_j - 1 \right) k_j + (1 - \ell_j) \widetilde{W}_j e_j \right\} + \widetilde{BEQ}_t = G_t \tag{10}$$

$$\sum_{j=0}^{1} N_j \; T_{SS} \left\{ (1 - \ell_j) \widetilde{W}_j e_j \right\} = \widetilde{TB}_t, \tag{11}$$

where N_j is the mortality-adjusted size of the age-j cohort, which grows exogenously at the rate of population growth g_N. Total accidental bequests collected from the deceased age-0 and age-1 households are $\widetilde{BEQ}_t$, aggregate public good expenditures are G_t, and $\widetilde{TB}_t$ are total Social Security benefits at time t. Finally, capital and labor markets clear in the aggregate, which implies that

$$K_t = E \left[\sum_{j=0}^{2} N_j \; k_{j+1} \right] \tag{12}$$

$$L_t = E \left[\sum_{j=0}^{2} N_j \; (1 - \ell_j) e_j \right]. \tag{13}$$

With this specification of the model economy, I next define the competitive equilibrium in this framework, and also discuss the methodological challenges posed by the presence of aggregate productivity risk.

3 Equilibrium computation

Let us assume that this stylized model economy starts at some arbitrary point of time $t = t_0$ in a given aggregate productivity state A_0. Then, the steady-state competitive equilibrium at a point of time $t > t_0$ is defined as

- a sequence of household choices $\{c_j(A_t), l_j(A_t), k_{j+1}(A_t)\}_{j=0}^{2}$,
- firm choices $\{K_t(A_t), L_t(A_t)\}$,
- Social Security benefits $\{B_t\left(\bar{y}|A_t\right)\}$, and
- wages and gross rates of return $\{W_t(A_t), R_t(A_t)\}$

under which

- the households' expected utility and the firm profits are maximized,
- the government's budget constraints are satisfied, and
- the labor and capital markets clear.

I normalize the initial aggregate productivity state and cohort size to $A_0 = N_0 = 1$.

The presence of aggregate productivity risk in an otherwise standard overlapping-generations framework poses a number of methodological challenges. Without aggregate risk, steady-state wages and interest rates are functions of the steady-state wealth distribution, which in general exists and can be calculated. However, in the presence of aggregate risk, a steady-state wealth distribution does not exist in general. This implies that the wages and interest rates that households need in order to solve their dynamic optimization problems in the steady state are, in general, unknown.

There are two broad approaches to solving this problem in the literature. Krusell and Smith (1998) were the first to approximate the steady-state wealth distributions by forecasting a summary statistic, generally the mean, of the distribution. While their application was in an infinite horizon setting, Storesletten et al. (2007) and Gourinchas (2000) apply the same idea to an overlapping-generations model with aggregate risk. Krueger and Kubler (2004), on the other hand, develop a projection algorithm for approximating the steady-state wealth distribution in the presence of aggregate risk. They show that the two approaches yield similar results when the variations in saving propensity across model households is not large – a condition under which "quasi-aggregation" is obtained; i.e. aggregation of saving across households is "close enough" to aggregate capital in the steady state.

In this paper, I follow Krusell and Smith (1998), and adopt the methodology in Storesletten et al. (2007) and Gourinchas (2000) of using a summary statistic of the current (stochastic) wealth distribution to forecast the future macroeconomic aggregates, which determines the wages and interest rates that households need to solve their dynamic optimization problems in the steady state. These forecasts are then used to make life-cycle consumption-saving and labor supply decisions under aggregate risk, which then determine the actual future wealth distributions under those forecasts. Then, the forecast functions are updated until the future macroeconomic aggregates cannot be forecast with any more precision, and "quasi-aggregation" is obtained under these "best" forecast functions.

There are three relevant macroeconomic aggregates that jointly govern the household- and firm-level decisions that satisfy the steady-state competitive equilibrium conditions above: the effective capital-labor ratio (K/AN), a benefit scaling factor for Social Security budget balance $(\bar{B})$, and an economy-wide "standard of living" measure, which in this case is mean earnings $(\bar{y})$. Specifically, I assume that

$$\log\left[X'\right] = \log\left[X\right]\Phi + \epsilon, \tag{14}$$

where $X = \left[(K/AN)\ \bar{B}\ \bar{y}\right]^T$ are the regressors, Φ is the matrix of regression coefficients with the first column being a unit vector, and ϵ is the error vector. Next, I discuss how I estimate this equation to approximate the steady-state competitive equilibrium of this overlapping-generations model with aggregate risk.

4 Calibration

Because the goal of the current paper is to examine the effect of intergenerational risk on the macroeconomic and welfare implications of unfunded public pensions – U.S. Social Security in particular, I begin by defining a "riskless" version of the above model as the benchmark, which I henceforth refer to as Model $R1$. In Model $R1$, aggregate labor productivity is deterministic, and is equal to the unconditional mean of stochastic realizations. Specifically,

$$Y_t = K_t^{\alpha} \left(A_t L_t\right)^{1-\alpha} \tag{15}$$

$$A_t = A_0(1+g_A)^t \times E\left[A_t^{\epsilon}\right] \tag{16}$$

$$E\left[A_t^{\epsilon}\right] = A_0 = 1 \tag{17}$$

Essentially, Model $R1$ is the canonical framework that has been widely used to examine the macroeconomic and welfare effects of Social Security in general. In this framework, Social Security facilitates implicit contracts between current workers and current retirees (through its unfunded structure), and also between current workers (or future retirees) through the progressive benefit-earnings rule. I calibrate the parameters of Model $R1$ using values that are commonly used in the literature, while also targeting relevant U.S. macroeconomic variables. Once Model $R1$ is calibrated, I then proceed to introduce aggregate labor productivity risk in the model environment (the "stochastic" version), while keeping the other model parameter values fixed.

4.1 Preferences and Production

To first calibrate Model $R1$, I begin by specifying the functional form for the utility and the production functions. I follow the macro public finance literature and assume a period utility function

$$u(c,\ell) = \frac{\left(c^{\eta}\ell^{1-\eta}\right)^{1-\gamma} - 1}{1-\gamma} \tag{18}$$

where γ is the inverse of the elasticity of intertemporal substitution, and η is the consumption share. I similarly specify the production function to the standard Cobb-Douglas form

$$Y = K^{\alpha} \left(AL\right)^{1-\alpha}, \tag{19}$$

where α is the capital's share and A is the aggregate labor productivity. This production function, along with equilibrium in spot markets for capital and labor, yield competitive factor prices that are equal to their net marginal products

$$W_t = (1-\alpha)A_t K_t^{\alpha} \left(A_t L_t\right)^{-\alpha} \tag{20}$$

$$R_t = \alpha K_t^{\alpha-1} \left(A_t L_t\right)^{1-\alpha} + 1 - \delta, \tag{21}$$

where I set the production-side parameters to $\alpha = 0.3$ and $\delta = 0.5$. Finally, I abstract from trend economic growth, i.e. I set the aggregate productivity growth rate to $g_A = 0.0$. Note that the factor prices are deterministic in Model $R1$, but are stochastic in the presence of aggregate risk.

4.2 Demographics and human capital

Next, I calibrate the demographic parameters using values typically used in the literature. Mapping a typical 75-period annual model framework into three periods implies that each period is approximately equal to 25 years. Using this scaling, I set the population growth rate to $g_N = 0.2824$, and then I use the age-25, 50, and 75 death rates from the U.S. Life Tables in Arias (2004) to calculate the conditional survival probabilities $\{Q_j\}_{j=0}^{2}$. To calibrate the household-level labor income process, I set the age-dependent productivity endowments $\{e_j\}_{j=0}^{2}$ based on estimates from Kitao (2013), who uses work hour and wage data from the PSID to derive this component as a residual of

Model age	0	1	2
Q_j	0.994	0.9443	10^{-4}
e_j	1.0	1.2	0.0

Table 1: Survival probabilities and the age-dependent productivity endowment.

γ	β	η	τ_1
2.0	0.96^{25}	0.33	0.72

Table 2: Unobservable parameters of Model $R1$.

	Target	Model $R1$
Capital-output ratio	3.0	2.92
Avg. time spent in market work	0.33	0.298
Share of govt. spending in GDP	0.2	0.19
Social Security as % of GDP	–	0.074

Table 3: Performance of Model $R1$.

wages, after accounting for hours worked and also the part-time wage penalty. These parameters are reported in Table 1. Finally, I calibrate the variance of the log of the human capital (education) fixed effect to $\sigma_\theta^2 = 0.012$, and I use Gaussian quadrature to approximate it using a three-point discrete distribution, which yields $\theta = \{0.8272, 1.0, 1.2089\}$.

4.3 Government

Next, I follow Karabarbounis (2012) and Heathcote et al. (2010) to calibrate the labor and capital income tax function

$$T_y(y) = y - \tau_1 y^{1-\tau_2}, \tag{22}$$

where $\tau_1 < 1$ and $\tau_2 > 0$. With this income tax function, after-tax labor income is log-linear in before-tax labor income, and the parameter τ_2 controls the progressivity of the tax code. Following Heathcote et al. (2010), I set the value of this parameter to $\tau_2 = 0.151$, and then calibrate the scale of taxes τ_1 so that model yields a realistic taxes-to-GDP ratio.

To calibrate Social Security in the model, I first set the payroll tax rate to $\tau_{SS} = 0.106$, and assume that this tax applies only up to the maximum taxable earnings, which is about 2.47 times the average earnings. Second, to compute the Social Security benefit amount (also known as the PIA), I incorporate the concave (piecewise linear) benefit-earnings rule used in the U.S. to calculate the replacement rate as a function of past earnings (see Figure 1).

4.4 Unobservable parameters

Finally, I calibrate the remaining unobservable parameters of Model $R1$ such that the model yields a steady-state equilibrium consistent with the U.S. macroeconomic data. In particular, the three preference parameters γ, β, and η, and the scale of taxes τ_y are chosen to yield a steady-state capital-output ratio of around 3.0, an average time spent in market work of 33%, and a share of government spending in GDP of around 20%. The value under which Model $R1$ matches these targets are reported in Table 2, and the model performance under these values is reported in Table 3. Note that even though I do not target it, Model $R1$ yields a realistic ratio of Social Security expenditures as a percentage of GDP of 7.4% in the steady state. The cross-sectional assets and labor supply profiles in this “riskless” steady state are reported in Figure 2.

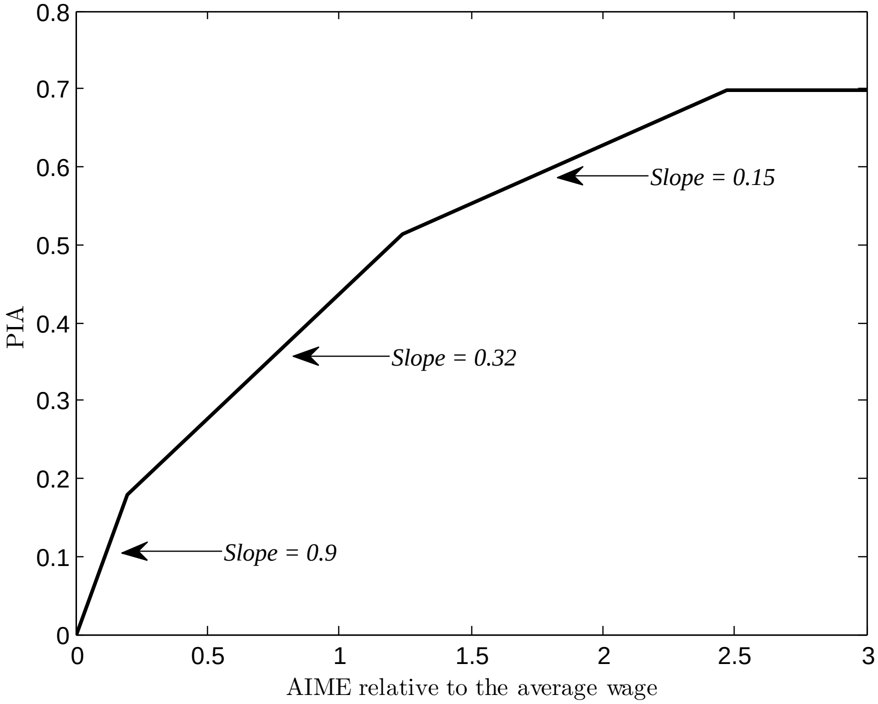

Figure 1: Benefit formula in the U.S.

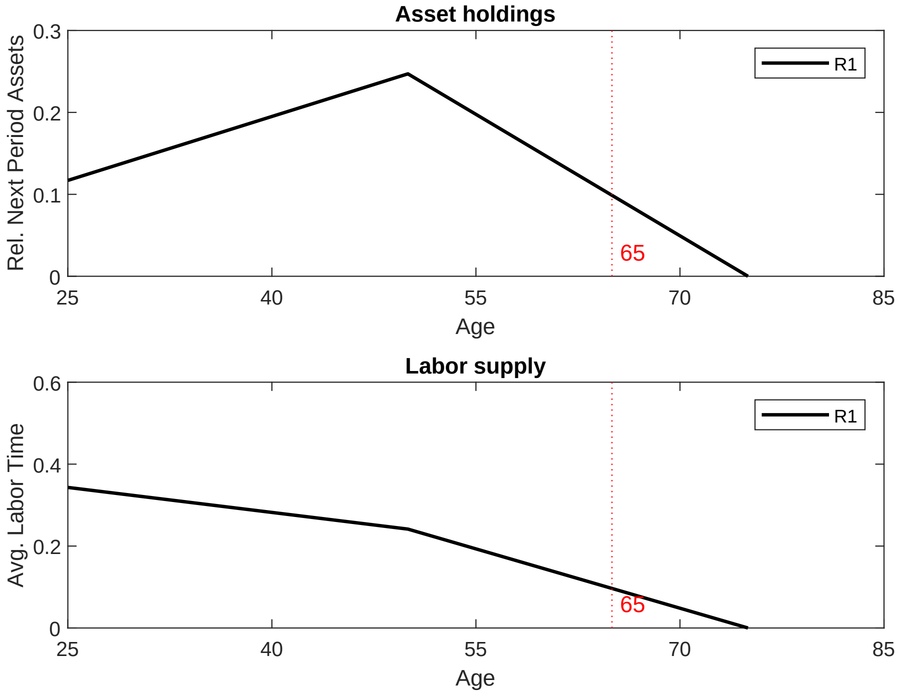

Figure 2: Equilibrium assets and labor supply in Model $R1$.

	Target	Model $R1$	Model $S1$ $[-,+]$	Mean
Capital-output ratio	3.0	2.92	[2.9, 3.0]	2.96
Avg. time spent in market work	0.33	0.298	[0.3116, 0.2944]	0.303
Share of govt. spending in GDP	0.2	0.21	[0.017, 0.158]	0.084
Social Security as % of GDP	–	0.074	[0.075, 0.074]	0.0748

Table 4: Performance of Model $S1$ Vs Model $R1$.

	Intercept	$\ln(K/AN)$	$\ln(\bar{B})$	$\ln(\bar{y})$	Forecast/Actual
$\ln(K'/A'N')$	−3.2985	−0.2007	0.0175	0.1336	0.98
$\ln(\bar{B}')$	−8.8485	−1.4487	0.3173	−1.6152	1.02
$\ln(\bar{y}')$	−2.5248	0.0916	0.0187	−0.1458	0.99

Table 5: Regression coefficients of the "best" forecast functions for Model $S1$.

Next, I compute the "stochastic" version of the model, in which the log of aggregate labor productivity (A) follows the $AR(1)$ process

$$\ln A_t = \rho \ln A_{t-1} + \sigma_A \epsilon_t \tag{23}$$

$$E\left[A_t^{\epsilon}\right] = A_0 = 1, \tag{24}$$

where ρ is the persistence parameter, σ_A is the variance, ϵ_t is a white-noise disturbance term, and $A_0 = 1$. I approximate this stochastic path with a two-state shock distribution

$$\ln A_t \in \left[\frac{-\sigma_A}{\sqrt{1-\rho^2}}, \frac{+\sigma_A}{\sqrt{1-\rho^2}}\right], \tag{25}$$

which implies the transition matrix

$$Tr\left(A_t\right) = \begin{bmatrix} \frac{1+\rho}{2} & \frac{1-\rho}{2} \\ \frac{1-\rho}{2} & \frac{1+\rho}{2} \end{bmatrix}.$$

I henceforth refer to this "stochastic" model as Model $S1$. I set the variance of the log of the aggregate shock process to $\sigma_A = 0.01$, and its persistence to $\rho = 0.97$. As we will see, the persistence of this $AR(1)$ process, governed by the value of the parameter ρ, will play an important role in our computational experiments. With this specification for the aggregate shock, the approximate steady state of Model $S1$ is reported in Table 4. I also report the "best" forecast functions that households use in order to solve their dynamic programming problems in this approximate steady state in Table 5.

At this point, it is worth considering the effect of aggregate risk on the performance of the calibrated overlapping-generations model. First, as Table 4 shows, introducing aggregate risk causes the mean capital-output ratio and the average labor time to slightly increase, and the mean government's share in GDP to decrease. In other words, households appear to respond by changing their saving and labor supply behavior in the presence of aggregate risk, which suggests that both margins are important consumption-stabilizing tools. I report the cross-sectional assets and labor time profiles under this "stochastic" approximate steady-state in Figure 3, which illustrates that both saving and labor supply respond more strongly to a favorable aggregate productivity shock.

Second, because future factor prices are uncertain in the presence of aggregate risk, households forecast future values of key macroeconomic variables as a function of their current aggregates. Table 5 shows that households make relatively small errors in making these forecasts. The forecast

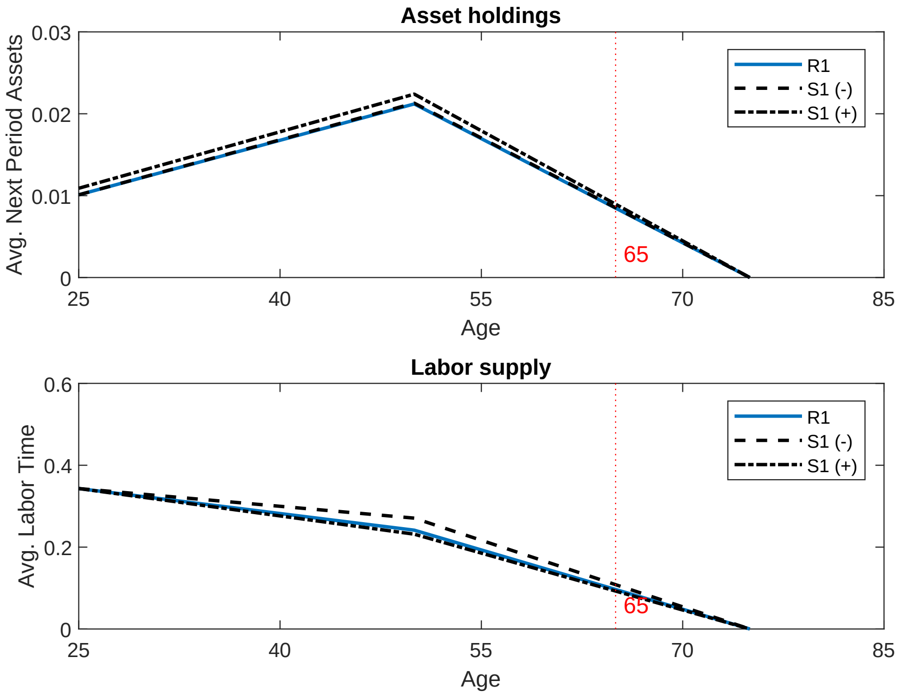

Figure 3: Equilibrium assets and labor supply in Model $S1$.

future values of the effective capital-labor ratio ($K'/A'N'$), the benefit scaling factor for Social Security budget balance ($\bar{B}'$), and the mean earnings ($\bar{y}'$) are reasonably close to their actual mean values in the approximate steady state. Overall, I find that household behavior does respond to the introduction of aggregate risk in a quantitatively meaningful way.

5 The Importance of Aggregate Risk

The goal of this paper is to quantitatively examine how the introduction of aggregate risk in an otherwise standard model framework affects the macroeconomic and welfare implications of unfunded public pensions – Social Security in particular. To accurately assess this, two sets of computational experiments are needed: one with the baseline model (or Model $R1$), and another with the baseline model augmented with aggregate risk, or Model $S1$. For both models, I choose the computational experiment to be a relatively straightforward one: a 50% cut in Social Security's payroll tax rate, i.e. a reduction in the payroll tax from its baseline value of 10.6% to 5.3%. For each model, I compute a new steady state under this lower payroll tax rate of 5.3%, and then compare the results across the models.

5.1 Social Security Tax Cut: Macroeconomic Effects

The steady-state values of key macroeconomic variables for Model $R1$ under this lower payroll tax rate are compared to the baseline in Table 6. It is clear from the table that downsizing Social Security in Model $R1$, as expected, leads to a large increase in both capital stock and labor supply. Capital increases by more than 23% and labor increases by 6.4%, which together leads to a 10.5% increase in GDP and a 9.6% increase in aggregate consumption. On the other hand, Social Security benefits per retiree declines by about 44%, which is slightly smaller than the size of the payroll tax

	Model $R1$ 10.6%	Model $R1$ 5.3%	Change (%)
Capital stock	0.0265	0.0327	23.4
Labor	0.5713	0.6078	6.4
Output	0.20	0.22	10.5
Consumption	0.2347	0.2572	9.6
Benefit per retiree	0.0294	0.0166	-43.5

Table 6: Model $R1$ under the experiment.

	Model $S1$ 10.6%	Model $S1$ 5.3%	Mean % Ch. given initial state
	$[-,+]$	$[-,+]$	$[-,+]$
Capital stock	[0.017, 0.0182]	[0.0191, 0.0215]	[12.6, 17.9]
Labor	[0.3827, 0.3593]	[0.3827, 0.3827]	[0.0, 6.5]
Output	[0.15, 0.151]	[0.1514 0.1661]	[3.8, 9.9]
Consumption	[0.1516, 0.1541]	[0.1535, 0.1689]	[1.4, 9.5]
Benefit per retiree	[0.0295, 0.0295]	[0.0167, 0.0167]	$[-43.4, -43.4]$

Table 7: Model $S1$ under the experiment.

	Intercept	$\ln(K/AN)$	$\ln(\bar{B})$	$\ln(\bar{y})$	Forecast/Actual
$\ln(K'/A'N')$	-2.2423	0.4002	-0.1134	-0.1238	0.99
$\ln(\bar{B}')$	-2.1718	-0.0475	0.0646	-0.3527	1.0
$\ln(\bar{y}')$	-2.7147	-0.0675	0.0069	-0.0577	1.02

Table 8: Regression coefficients of the "best" forecast functions in Model $S1$ under the experiment.

cut. This should not be surprising, because an important effect of the tax cut in general equlibrium is to increase Social Security's tax base (Bagchi, 2016). All of these results are in line with the rest of the literature (see Bagchi (2016) and other papers cited therein).

On the other hand, the approximate steady-state values of the macroeconomic variables for Model $S1$ under the lower payroll tax rate are reported in Table 7, and the corresponding "best" forecast functions are reported in Table 8. It is clear from Table 7 that the effects of an identical downsizing of Social Security are significantly different in Model $S1$. Both aggregate capital stock and labor supply increase, but the size of the increase is state-dependent. Specifically, when the economy is initially in an unfavorable productivity state ($-$), the mean percentage increase in capital stock from the tax cut is 12.6%, and when it is initially in a favorable productivity state (+), the mean percentage increase is nearly 18%. It is worth noting that both of these are smaller than the 23% increase in the absence of aggregate risk. Similarly, conditional on an unfavorable initial state ($-$), the tax cut leaves labor supply unchanged on the average, whereas given a favorable initial state (+), the average increase is 6.5%. Collectively, these changes correspond to average increases in output of 3.8% and 9.9% respectively, which are both smaller than what is observed without aggregate risk. Finally, the tax cut increases aggregate consumption by 1.4% when the economy is initially in an unfavorable productivity state ($-$), and by 9.5% starting from a favorable productivity state (+), compared to a larger increase of 9.6% in the absence of aggregate risk. To summarize, I find that the effects of a 50% cut in Social Security's payroll tax are somewhat muted under aggregate risk, particularly when the economy is initially in an unfavorable productivity state.

To illustrate how the household-level responses to a downsizing of Social Security drive these macroeconomic outcomes, I report the cross-sectional assets and labor time profiles under both models $R1$ and $S1$ in the baseline scenario as well as with the 50% payroll tax cut in Figures

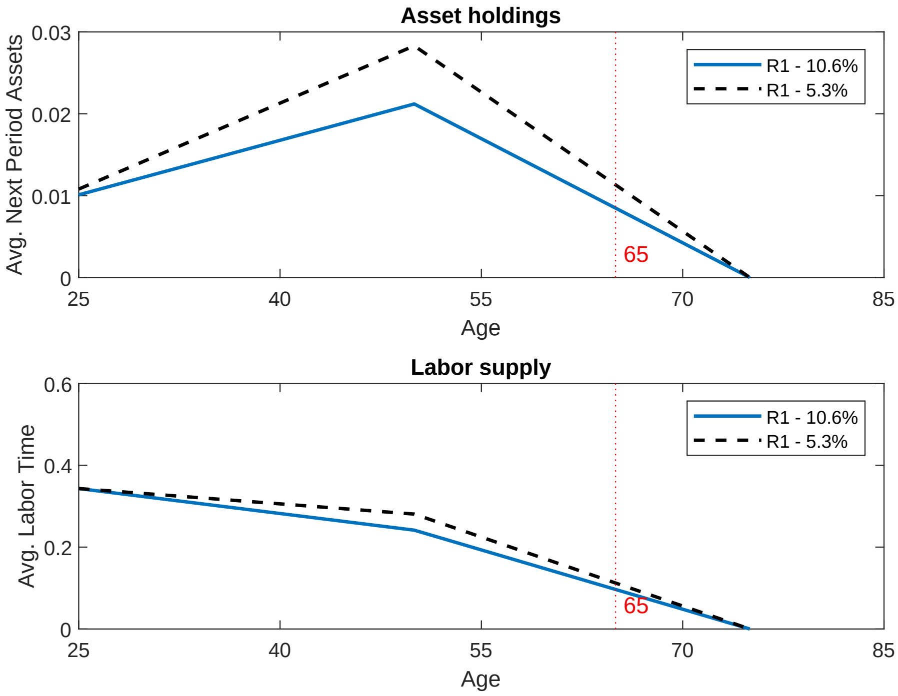
Asset holdings

Labor supply

Figure 4: Effect of the tax cut on assets and labor supply in Model $R1$.

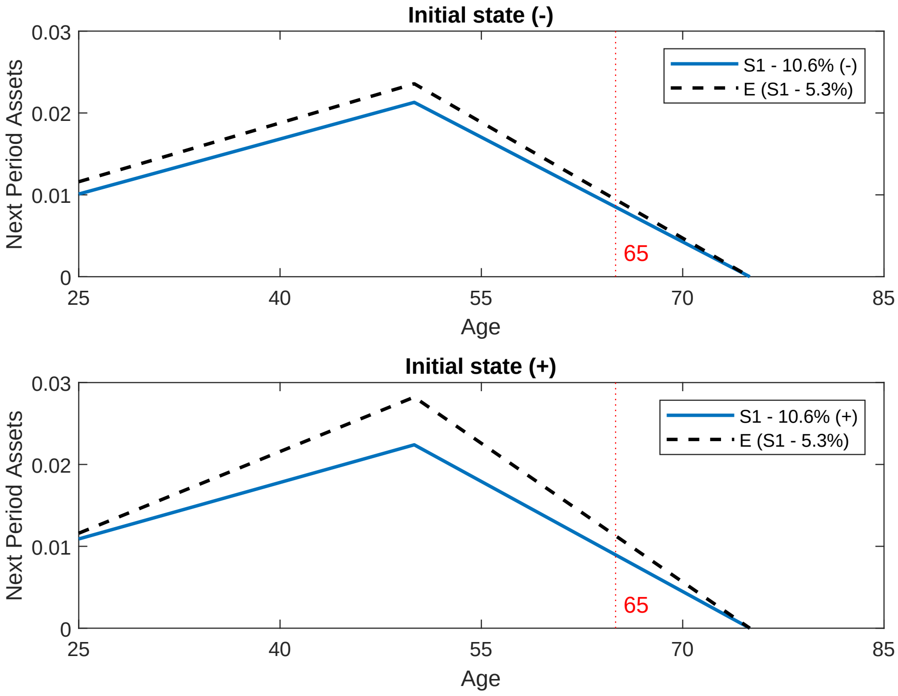
Initial state (-)

Initial state (+)

Figure 5: Effect of the tax cut on assets in Model $S1$.

4, 5, and 6. Figure 4 shows that as expected, the tax cut leads to an increase in savings and labor supply in the riskless model. However, the effects in the model with aggregate risk are more complex, as is clear from Figures 5 and 6. For example, if the economy is initially in an unfavorable

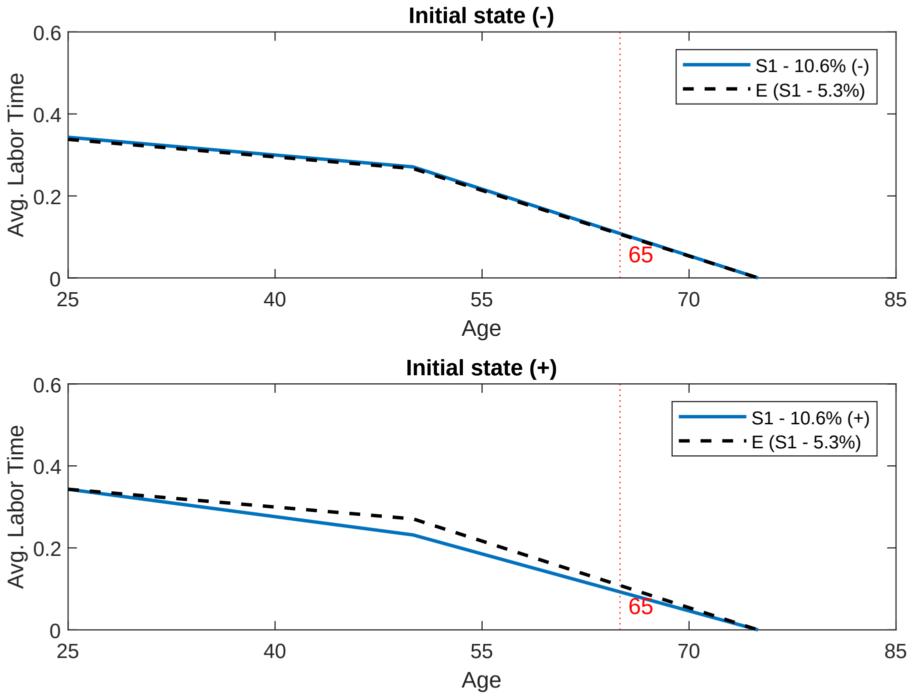

Figure 6: Effect of the tax cut on labor supply in Model $S1$.

productivity state $(-)$, the tax cut, on the average, has only a small effect on both household saving and life-cycle labor supply. On the other hand, if the economy is initially in a favorable productivity state $(+)$, the effects on both are large and positive. Together, these results suggest that the starting productivity state of the economy is an important determinant of the tax cut's macroeconomic effects.

As the reader will recall, the persistence of the aggregate shock process, denoted by the value of the parameter ρ, determines the relative likelihood of state transitions in the current model. For example, with a persistence of $\rho = 0.97$, the likelihood of an own-state transition in a two-state approximation of the $AR(1)$ shock process is 98.5%. Consequently, households do not expect the aggregate productivity state to change significantly in the near future, and this expectation dampens their saving and labor supply responses, particularly when the current state is unfavorable. I will further investigate the quantitative importance of this mechanism in Section 6, but before that let us consider the effect of aggregate risk on the welfare consequences of downsizing Social Security.

5.2 Social Security Tax Cut: Welfare Effects

Downsizing Social Security weakens the implicit financial contracts between current workers and current retirees, between current workers (or future retirees), and also between current workers and *future* workers. Therefore, evaluating the welfare implications of this downsizing in an environment with aggregate risk requires one to consider the effect of the tax cuts on these implicit contracts. To do this, I first define the welfare measure as a consumption equivalence (CEV) needed for a newborn household to be indifferent (in terms of ex-ante expected utility) between the steady-state

Human capital	Low	Medium	High	Overall
Model $R1$	8.3	8.4	7.9	8.3
Model $S1$ $(-)$	4.4	6.5	5.1	5.9
Model $S1$ $(+)$	6.1	6.3	6.2	6.3

Table 9: The $CEV(\%)$ values under the experiment for models $R1$ and $S1$.

economies in the baseline and under the experiment. Specifically,

$$\begin{aligned} &E\left[Q_0\ u\left(c_0^B(1+CEV),\ell_0^B\right)+\beta Q_0 Q_1\ u\left(c_1^B(1+CEV),\ell_1^B\right)+\beta^2 Q_0 Q_1 Q_2\ u\left(c_2^B(1+CEV),1\right)\right] \\ &\qquad = E\left[Q_0\ u\left(c_0^H,\ell_0^H\right)+\beta Q_0 Q_1\ u\left(c_1^H,\ell_1^H\right)+\beta^2 Q_0 Q_1 Q_2\ u\left(c_2^H,1\right)\right], \end{aligned} \tag{26}$$

where B denotes baseline (both in models $R1$ and $S1$), and H denotes the hypothetical case with a 50% cut in Social Security's payroll tax rate. I calculate this CEV measure both at the aggregate level, and also by the human capital (education) fixed effects. In Table 9, I report both the aggregate and disaggregated CEVs under the tax cut in the models $R1$ and $S1$. Note that $(-)$ and $(+)$ respectively denote the initial productivity state of the Model $S1$ economy prior to the tax cut. The following facts are clear from the table. First, downsizing Social Security has a positive effect on welfare in both Models $R1$ and $S1$. This is a known result in general equilibrium (cite papers), and it is also consistent with Krueger and Kubler (2004). Second, the overall welfare gains from the tax cut in Model $S1$ are smaller than those in Model $R1$, regardless of the initial productivity state of the economy. This suggests that claims to future labor income are somewhat more valuable in an environment with aggregate productivity risk. Downsizing Social Security weakens its implicit intergenerational risk sharing mechanism, and the above results indicate that this mechanism is more important in the presence of aggregate risk.

Table 9 also shows that in Model $R1$, the welfare gains are fairly uniform across the human capital fixed-effect categories, but not in Model $S1$. If the economy is initially in an unfavorable productivity state $(-)$, mean welfare gains from the tax cut are a hump-shaped function of human capital. On the other hand, with a favorable initial productivity state $(+)$, mean welfare gains from the tax cut are similar across all household types. This suggests that the redistribution implicit in Social Security's benefit-earnings rule (i.e. the implicit financial contract between current workers with varying labor-market outcomes) interacts with its intergenerational risk sharing mechanism, especially when the economy is initially in an unfavorable productivity state.

To summarize, the above results indicate that aggregate risk has an important effect on the macroeconomic and the welfare consequences of downsizing Social Security. Both are markedly muted in the presence of aggregate risk, particularly when the economy is initially in an unfavorable productivity state. Moreover, the inter- and intra-generational insurance effects of Social Security appear to interact in the presence of aggregate risk. In the next section, I define and compute two additional experiments that help tease out the mechanisms underlying these effects.

6 The Underlying Mechanisms

As explained earlier, the intergenerational financial contract implicit in Social Security is effectively a claim on future labor income, which is uncertain in the presence of aggregate risk. The results above suggest that the quantitative importance of this implicit contract depends on the aggregate risk environment. Because the aggregate shock process is highly persistent in the baseline calibration, future productivity states are strongly correlated with the current state. Social Security's intergenerational risk sharing effect depends on the strength of this correlation.

	Target	$S1$ $[-,+]$	Mean ($S1$)	$S2$ $[-,+]$	Mean ($S2$)
Capital-output ratio	3.0	[3.0, 3.1]	3.05	[2.9, 2.8]	2.8
Avg. time spent in market work	0.33	[0.3073, 0.2902]	0.2988	[0.2923 0.2883]	0.2904
Share of govt. spending in GDP	0.2	[0.014, 0.154]	0.084	[0.02, 0.066]	0.043
Social Security as % of GDP	–	[0.02, 0.072]	0.0726	[0.074, 0.074]	0.074

Table 10: The "stochastic" Models $S1$ and $S2$.

	Intercept	$\ln(K/AN)$	$\ln(\bar{B})$	$\ln(\bar{y})$	Forecast/Actual
$\ln(K'/A'N')$	−2.1854	0.1797	0.055	0.1477	0.98
$\ln(\bar{B}')$	2.1178	−0.2683	0.0382	1.3256	0.99
$\ln(\bar{y}')$	−2.1695	−0.0268	−0.0111	0.1422	0.99

Table 11: Regression coefficients of the "best" forecast functions of Model $S2$.

	Model $S2$ 10.6%	Model $S2$ 5.3%	Mean % Ch. given initial state
	$[-,+]$	$[-,+]$	$[-,+]$
Capital stock	[0.0149, 0.0159]	[0.0183, 0.019]	[24.0, 18.4]
Labor	[0.3564, 0.3509]	[0.3814, 0.3827]	[7.1, 8.9]
Output	[0.1336, 0.143]	[0.1489, 0.1601]	[13.6, 9.9]
Consumption	[0.1322, 0.1398]	[0.1431, 0.1528]	[10.1, 7.5]
Benefit per retiree	[0.0428, 0.0428]	[0.0204, 0.0224]	[−51.1, −48.9]

Table 12: Model $S2$ under the experiment.

To further investigate this mechanism, I define an alternative calibration of Model $S1$ with an aggregate shock process that is only 50% as persistent as the baseline, while also adjusting its variance so that the spread of productivity realizations is unchanged at its baseline level. Specifically, I set the persistence parameter to $\rho_A = 0.97/2 = 0.485$ and the variance to $\sigma_A = 0.037$, which leads to a roughly identical spread of productivity realizations around its unconditional mean of $E\left[A_t^{\epsilon}\right] = A_0 = 1.0$. I refer to this version of the "stochastic" model as Model $S2$.

In Table 10, I report the baseline characteristics of Model $S2$ along with those of Model $S1$. I report the estimated coefficients of the "best" forecast functions corresponding to Model $S2$ in Table 11. It is clear from Table 10 that a less persistent aggregate shock process leads to a decrease in the mean capital–output ratio and the mean share of government spending in GDP, but leaves average labor time roughly unchanged. Next, I implement the Social Security payroll tax cut experiment on Model $S2$, and I report the corresponding macroeconomic and welfare consequences in Tables 12 and 13 respectively. As Table 12 shows, the macroeconomic effects of downsizing Social Security under a less persistent aggregate shock process are quite different. First, the mean increases in capital stock, labor, output, and consumption are all larger compared to Model $S1$ (see Table 7. Second, under an unfavorable initial productivity state $(-)$, the mean percentage increase in capital stock from the tax cut under Model $S2$ is 24%, but under a favorable initial productivity state $(+)$, it is around 18%. This pattern is quite different from that observed in Model $S1$, where the larger increase is associated with a favorable initial productivity state $(+)$. This is also true for the other macroeconomic variables. Third, larger increases in capital and labor imply larger increases in output and consumption from the tax cut under Model $S2$, which are now much closer to the increases observed in the absence of aggregate risk (Model $R1$). In other words, a less persistent aggregate shock process appears to undercut the quantitative importance of aggregate risk in the context of the tax cut experiment.

The welfare effects of the tax cut under this model (Table 13) further highlight this mechanism.

	Human capital	Low	Medium	High	Overall
Model $R1$		8.3	8.4	7.9	8.3
Model $S1$ $(-)$		4.4	6.5	5.1	5.9
Model $S1$ $(+)$		6.1	6.3	6.2	6.3
Model $S2$ $(-)$		10.1	10.1	8.6	9.8
Model $S2$ $(+)$		6.3	5.4	6.7	5.8

Table 13: The $CEV(\%)$ values under the experiment for models $R1$, $S1$, and $S2$.

	Human capital	Low	Medium	High	Overall
Model $R1$		8.3	8.4	7.9	8.3
Model $S1$ $(-)$		4.4	6.5	5.1	5.9
Model $S1$ $(+)$		6.1	6.3	6.2	6.3
Model $R2$		8.4	8.4	7.5	8.2
Model $S3$ $(-)$		3.2	9.3	4.3	7.4
Model $S3$ $(+)$		7.8	12.3	8.0	10.7

Table 14: The $CEV(\%)$ values under the experiment for $R1$, $S1$, $R2$ and $S3$.

From an unfavorable initial productivity state $(-)$, the overall welfare gain from downsizing Social Security in Model $S2$ are higher than Model $S1$ (and also Model $R1$). This suggests that when aggregate risk is less persistent over time (i.e. when future states different from the current one are more likely), households do not value Social Security's implicit claim to future labor income as much, especially when the downsizing happens when the economy is in an unfavorable productivity state.

It is worth noting at this point that while the persistence of the aggregate shock process helps explain the relative importance of Social Security's implicit claim to future labor income, it does not provide any insights on the distribution of the welfare gains from the tax cut. Table 9 shows that in general, the welfare gains from downsizing Social Security are hump-shaped in the human-capital fixed effect. This pattern appears to be more pronounced in the presence of aggregate risk, i.e. in Model $S1$, particularly when the initial productivity state is unfavorable $(-)$. The human-capital distribution of the welfare effects, in turn, depends on the redistribution implicit in Social Security's benefit-earnings rule. Next, I investigate the interaction between aggregate risk and this intra-generational contract between current workers (or future retirees) implied by Social Security.

To examine this interaction, I define a second alternative calibration for both the initial "riskless" Model $R1$ as well as the "stochastic" Model $S1$, in which I replace Social Security's calibrated benefit-earnings rule with a hypothetical rule that removes the progressive relationship between benefits and past earnings. Specifically, I set Social Security's replacement rate (the ratio of the benefit annuity to average work-life earnings) equal to about 64%, regardless of a household's earnings history. With this specification, benefits are a strictly linear function of past earnings, which removes any insurance implicit in the contract between current workers (or future retirees) with varying earnings outcomes. I refer to these alternative "riskless" and "stochastic" versions of the original models as Models $R2$ and $S3$ respectively, and I report the welfare effects of a 50% cut in Social Security's payroll tax rate in Table 14.

Two facts are clear from Table 14. First, in the absence of aggregate risk, downsizing Social Security has similar overall welfare effects both with and without the progressive benefit-earnings rule. However, in the presence of aggregate risk, that is no longer the case: overall welfare gains from downsizing Social Security are generally larger with the hypothetical linear benefit-earnings rule. This should not be surprising, as the redistribution implicit in Social Security's current benefit-earnings rule strengthens its insurance effects, which play a larger role in the presence of aggregate risk. Second, the hump-shaped relationship between human capital and welfare gains in the presence of aggregate risk is even more pronounced under the linear benefit-earnings rule. In particular, under this hypothetical rule, welfare gains for low- and high-human capital households are *lower* when the downsizing happens in an unfavorable productivity state. This suggests that for these households, the absence of Social Security's within-cohort risk sharing mechanism makes the implicit claim to future labor income more valuable, particularly when the economy is initially in an unfavorable state.

To summarize, my computations from the alternative calibrations suggest that the persistence of the aggregate shock process is an important determinant of Social Security's intergenerational risk-sharing ability. When the persistence of the calibrated shock process is cut by half, downsizing Social Security yields smaller welfare gains, especially when the economy is initially in an unfavorable aggregate productivity state. This is because in this case, future states different from the current one are more likely, so households do not value Social Security's implicit claims to future labor income as much. I also find that aggregate risk changes how the redistribution implicit in Social Security's benefit-earnings rule interacts with its intergenerational risk sharing mechanism. When this implicit within-cohort redistribution is eliminated, low- and high-income households find claims to future labor income more valuable, particularly when Social Security is downsized under an unfavorable productivity state.

7 Conclusions

This paper examines the effect of aggregate risk on the macroeconomic and welfare implications of unfunded public pensions – U.S. Social Security in particular. To do this, I design a stylized overlapping–generations macroeconomic model with rational utility-maximizing households, profit-maximizing firms, and a government that provides public goods and Social Security. Households in the model survive for a maximum of three periods, and they vary in their earning ability depending on their human capital (or education) level. Aggregate productivity is stochastic, and it affects overall factor returns. Most notably, this stochastic aggregate productivity leads to a general non-existence of the model's steady-state wealth distribution, so households have to forecast future macroeconomic aggregates based on their current values. I calibrate this model to match U.S. macroeconomic targets, and then compute the effect of aggregate risk on the macroeconomic and welfare implications of a 50% cut in Social Security's payroll tax rate.

In general, my findings suggest that introducing aggregate risk in this otherwise standard model framework has quantitatively important implications for Social Security's macroeconomic and welfare effects. First, downsizing Social Security in the presence of aggregate risk, on the average, leads to smaller increases in capital stock, labor supply, output, and consumption. Moreover, these increases are highly state-dependent, with the smaller effects occurring when the economy is initially in an unfavorable productivity state. These smaller increases generate smaller welfare gains from the tax cut in the presence of aggregate risk. Second, the persistence of the aggregate shock process is a key determinant of the importance of current workers' claims to future labor income. I find that with a shock process only half as persistent as the baseline, the macroeconomic and welfare effects of downsizing Social Security are much closer to those observed in the absence of aggregate risk, but in this case the smaller effects occur when the economy is initially in a favorable productivity state. Finally, in the presence of aggregate risk, the welfare gains from downsizing Social Security follow a hump-shaped pattern with respect to earnings. This hump shape is even more pronounced under a hypothetical linear benefit-earnings rule, which suggests that aggregate risk also changes how the implicit financial contract between current workers with varying labor-market outcomes interacts with its intergenerational risk sharing mechanism.

References

Abel, Andrew B. (1985), “Precautionary Saving and Accidental Bequests.” *The American Economic Review*, 75, 777–791.

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

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

Gourinchas, Pierre-Olivier (2000), “Precautionary Savings, LifeCycle and Macroeconomics.” Technical Report 0793, Econometric Society.

Harenberg, Daniel and Alexander Ludwig (2015), “Social security in an analytically tractable overlapping generations model with aggregate and idiosyncratic risks.” *International Tax and Public Finance*, 22, 579–603.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010), “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006.” *Review of Economic Dynamics*, 13, 15–51.

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

İmrohoroğlu, Ayse, Selahattin İmrohoroğlu, and Douglas H Joines (1995), “A Life Cycle Analysis of Social Security.” *Economic Theory*, 6, 83–114.

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

Kitao, Sagiri (2013), “Sustainable social security: Four options.” *Review of Economic Dynamics*, –.

Krueger, Dirk and Felix Kubler (2004), “Computing equilibrium in olg models with stochastic production.” *Journal of Economic Dynamics and Control*, 28, 1411–1436.

Krueger, Dirk and Felix Kubler (2006), “Pareto-improving social security reform when financial markets are incomplete!?” *The American Economic Review*, 96, 737–755.

Krusell, Per and Anthony A. Smith, Jr. (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106, 867–896.

Ríos-Rull, José-Víctor (1996), “Life-cycle economies and aggregate fluctuations.” *The Review of Economic Studies*, 63, 465–489.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron (2007), “Asset pricing with idiosyncratic risk and overlapping generations.” *Review of Economic Dynamics*, 10, 519–548.