Idiosyncratic Risk, Aggregate Risk, and the Welfare Effects of Social Security*

Daniel Harenberg[†] Alexander Ludwig[‡]

This Version: May 8, 2015

Abstract

We ask whether a pay-as-you-go-financed social security system is welfare improving in an economy with idiosyncratic productivity risk and aggregate business cycle risk. We show analytically that the whole welfare benefit from joint insurance against both risks is greater than the sum of benefits from insurance against the two isolated risk components. One reason is the convexity of the welfare gain in total risk. The other reason is a direct risk interaction which amplifies the utility losses from consumption risk. We proceed with a quantitative evaluation of social security's welfare effects. We find that introducing a small social security system leads to substantial welfare gains in expectation, even net of the welfare losses from crowding out. This stands in contrast to the welfare losses documented in previous studies which all consider only one risk in isolation. About 60% of the welfare gains would be missing when simply summing up the isolated benefits.

JEL classification: C68; E27; E62; G12; H55

Keywords: social security; idiosyncratic risk; aggregate risk; welfare

*An earlier version of this paper circulated under the title "Social Security and the Interactions Between Aggregate and Idiosyncratic Risk". We thank Klaus Adam, Alan Auerbach, Martin Barbie, Tino Berger, Antoine Bommier, Johannes Brumm, Georg Dürnecker, Ayse Imrohoroglu, David Jaeger, Philip Jung, Tom Krebs, Dirk Krueger, Felix Kubler, Michele Tèrtilt, Fabrizio Zilibotti, and various seminar participants at several place for helpful discussions. Daniel Harenberg gratefully acknowledges the financial support of the German National Research Foundation (SFB 884) and of Swiss Re. Alex Ludwig gratefully acknowledges financial support by the German National Research Foundation (SPP 1578) and by the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE.

†Chair of Integrative Risk Management and Economics, ETH Zürich; dharenberg@ethz.ch.

‡SAFE, Goethe University Frankfurt; MEA; ludwig@safe.uni-frankfurt.de.

1 Introduction

Many countries operate large social security systems. One reason is that social security can provide insurance against risks for which there are no private markets. However, these systems also impose costs by distorting prices and decisions. The question arises whether the benefits of social security outweigh the costs.

We address this question in a model economy featuring two types of risk, aggregate business cycle risk in form of aggregate wage and asset return risk on the one hand and idiosyncratic productivity risk on the other hand. We follow the literature and assume that insurance markets for both types of risk are incomplete. In such a setting, social security can increase economic efficiency by partially substituting for missing markets. However, it also distorts decisions leading to welfare losses from crowding out of capital formation. Our analysis differs from the previous literature in that prior studies characterized social security's welfare effects in models with only one type of risk. One strand of the literature examined social security when only aggregate risk is present, e.g., Krueger and Kubler (2006). In that setting, social security—by pooling aggregate wage and asset return risks across generations—can improve intergenerational risk sharing. The other strand only considered idiosyncratic risk, cf., e.g., İmrohoroğlu, İmrohoroğlu, and Joines (1995, 1998) and Conesa and Krueger (1999). There, social security is valuable because it redistributes ex-post from high to low productivity households which provides intragenerational insurance from the ex-ante perspective. Broadly speaking, both strands of this previous literature conclude that the costs of introducing social security outweigh the benefits.

Such a segregated view is incomplete because households face both types of risk over the life-cycle and because social security, when appropriately designed, can (partially) insure both types of risk. We also argue that simply combining the findings from previous studies leads to severe biases in the overall welfare assessment. Our theoretical contribution is to show analytically why the whole insurance benefit is substantially greater than the sum of the benefits from insurance against isolated risk components. Our quantitative contribution is to establish that joint insurance against both types of risk leads to large net welfare gains thereby turning previous findings in the literature upside down: social security is welfare improving from the ex-ante perspective.

We show that there are two biases when simply combining previous findings.

The first arises even when the two types of risk are statistically independent. This bias is a consequence of the *convexity of the welfare gain* (CWG) in total risk. The welfare gain is convex in the amount of total risk because the marginal utility of insurance increases disproportionately as risk increases. The crucial aspect to notice is that joint presence of idiosyncratic productivity and aggregate business cycle risk strongly fans out the earnings and consumption distributions thereby leading to a substantial risk exposure for households over their life-cycle. If social security is designed as a Beveridgean system with lump-sum pension benefits it provides partial insurance against this total life-cycle risk. Because of CWG, the whole benefit from insurance is therefore greater than the sum of benefits from insurance against the single risk components. We call this difference in welfare assessments the "CWG bias" and show that it increases in the total amount of risk. Since total life-cycle risk is large, we can expect this bias to be large.

The second bias stems from a direct interaction of risks in form of a *countercyclical cross-sectional variance* (CCV) of idiosyncratic productivity risk: the variance of idiosyncratic shocks is higher in a downturn than in a boom. The CCV has been documented in the data (Storesletten, Telmer, and Yaron 2004) and analyzed with respect to its asset pricing implications (Mankiw 1986; Constantinides and Duffie 1996; Storesletten, Telmer, and Yaron 2007).[1] It leads to a high variance of the idiosyncratic income component when the aggregate income component is low. Due to concavity of the utility function this amplifies the welfare gains from insurance against both risks.

To illustrate the main insurance channels of social security and the welfare consequences of CWG and CCV, we start our analysis by employing a simple two-period life-cycle model in which a household faces idiosyncratic and aggregate wage risk in the first period of life. In absence of social security, retirement consumption is financed by private savings with aggregate return risk. We study the welfare consequence of introducing a pay-as-you-go (PAYG) financed social security system. Social security payments are lump-sum transfers. By pooling idiosyncratic wage risks within and aggregate risks across generations, this Beveridgean system jointly provides partial insurance against idiosyncratic

[1] Guvenen, Ozkan, and Song (2014) find that the skewness of the earnings distribution, not its variance, is countercyclical. Attributing this observation to shocks, both specifications capture a direct interaction of risks. A countercyclical left-skewness would likely strengthen our results.

and aggregate risks. We measure welfare gains by a consumption equivalent variation. Initially abstracting from CCV, we derive a term capturing the welfare difference between the whole insurance benefit and the sum of the benefits from insurance against the isolated risk components. This difference reflects the bias that occurs from combining the findings of the previous literature, the CWG bias. We subsequently modify the two-period model to account for the CCV mechanism and show how an additional welfare difference emerges because of the increased variance of idiosyncratic risk in the low aggregate state of the economy.

Our arguments so far ignore behavioral reactions, i.e., the reduction of savings caused by social security. In general equilibrium, this savings reaction leads to crowding out of aggregate capital so that there may be similar biases with respect to the cost from crowding out. In our companion work (Harenberg and Ludwig 2015), we show that joint presence of idiosyncratic and aggregate risk may in- or decrease these welfare costs. Thus, while the CWG bias unambiguously increases welfare benefits, the bias in the welfare costs of crowding out is ambiguous. It therefore remains a quantitative question how large the effects are and whether the benefits outweigh the costs.

To address these quantitative questions we here build a large-scale overlapping generations model in the tradition of Auerbach and Kotlikoff (1987), extended by idiosyncratic productivity risk and aggregate wage and asset return risk. Households can save privately by investing in a risk-free bond and a risky stock. Including this portfolio choice is important. It allows us to appropriately calibrate the risk-return structure of the private savings technologies, which directly affects the value of social security. The possibility to save in two assets also implies that households have additional means of self-insurance. This reduces the welfare benefits of social security. We consider a pure Beveridgean pay-as-you-go (PAYG) social security system as in our simple two-period model. Our policy experiment consists of an introduction of such a system with a contribution rate of 2%. This is the size of the US system when first introduced in 1935. With this size and our design of the system we hence study the welfare implications of introducing a minimum pension.[2]

As our main quantitative contribution we find—by calibrating the model to

[2]Most real world pension systems feature distributional components, and almost all have a minimum pension. Our system features similarities to the Danish public pension system.

the U.S. economy—that such a marginal introduction of social security leads to a strong welfare gain of approximately 2% in terms of a consumption equivalent variation. This welfare improvement is obtained despite substantial welfare losses from crowding out of capital in general equilibrium. This finding stands in stark contrast to the previous literature. We also replicate this earlier literature by considering economies with only one type of risk. We indeed observe welfare losses in these economies.

To uncover the sources of the welfare gain in our model with both risks we decompose the welfare gain into its various components and oppose it with the welfare losses from crowding out in general equilibrium. The sum of the welfare components attributable to insurance against idiosyncratic and the aggregate risk is substantially smaller than the total welfare gain. Across various scenarios, the difference varies between 50% to 70% of the total gains from insurance. Roughly half of this difference is attributable to the CCV, the other half to the CWG. Hence, the biases are indeed sizeable. Correctly modeling and quantifying the total life-cycle risk exposure of households is therefore essential for an accurate welfare analysis of social security.

The notion that social security can insure against aggregate risks dates back to Diamond (1977) and Merton (1983). They demonstrate how it can partially complete financial markets thereby increasing economic efficiency. Building on these insights, Shiller (1999) and Bohn (2001, 2009) show that social security can reduce consumption risk of all generations by pooling labor income and capital income risks across generations. Gordon and Varian (1988), Matsen and Thogersen (2004), Krueger and Kubler (2006), and Ball and Mankiw (2007) use a two-period partial equilibrium model in which households only consume in the second period of life, i.e., during retirement. For our analytical results, we extend this model by adding idiosyncratic risk. Among the few quantitative papers with aggregate risk and social security, Krueger and Kubler (2006) is the most similar to our work.[3] They examine the same introduction of a small PAYG system and conclude that it does generally not constitute a Pareto-improvement.[4] The concept of a Pareto-improvement requires that they take an ex-interim welfare perspective, whereas we calculate welfare from an ex-ante perspective. Our

[3]Ludwig and Reiter (2010) assess how pension systems should optimally adjust to demographic shocks. Olovsson (2010) contends that pension payments should be highly risky because this increases precautionary savings and thereby capital formation.

[4]The recent work by Hasanhodzic and Kotlikoff (2013) mirrors these findings.

analysis differs substantially because we also include idiosyncratic risk and analyze interactions between the risks.

Many quantitative papers consider idiosyncratic risk and social security, e.g., Conesa and Krueger (1999), İmrohoroğlu, İmrohoroğlu, and Joines (1995, 1998), Huggett and Ventura (1999), and Storesletten, Telmer, and Yaron (1999). One general conclusion from this literature is that welfare in a stationary economy without social security is higher than in one with a PAYG system. More recent work such as Nishiyama and Smetters (2007) and Fehr and Habermann (2008) focuses on modeling the institutional features of existing social security systems in detail which we abstract from. Our results demonstrate the benefits of a flat minimum pension.[5]

We derive our analytical results in Section 2. Section 3 describes the quantitative model, Section 4 presents the calibration and Section 5 the main results of our quantitative analysis. We conclude in Section 6. Proofs as well as computational and calibration details are relegated to separate appendices.

2 A Two-Generations Model

We adopt the partial equilibrium framework of Gordon and Varian (1988), Matsen and Thogersen (2004), Krueger and Kubler (2006), and Ball and Mankiw (2007), among others, who assume that members of each generation consume only in the second period of life. This literature considers only aggregate risk. We extend it by adding idiosyncratic risk. We use the model to derive simple and insightful expressions for the welfare gains from insurance against individual risk components and for the difference between the whole insurance benefits and the sum of its parts, the CWG bias. We subsequently modify the setup to account for the CCV effect.

2.1 Model

In each period t, a continuum of households is born who live for two periods only. A household has preferences over consumption in the second period. In the first period of life, a household experiences an idiosyncratic productivity

[5] Finally, Gomes, Michaelides, and Polkovnichenko (2012) use a very similar model to study how changes in fiscal policy and government debt affect asset prices and capital accumulation.

shock, denoted by η. This shock induces ex-post heterogeneity, so that we denote ex-post different households by i. Age is indexed by j with $j = 1$ being working age and $j = 2$ being retirement. Denoting by $c_{i,2,t+1}$ consumption in retirement, the expected utility of a household born in period t is given by $E_t\left[u(c_{i,2,t+1})\right]$. We assume a CRRA per period utility function with coefficient of relative risk aversion θ, $u(c_{i,2,t+1}) = \frac{c_{i,2,t+1}^{1-\theta}-1}{1-\theta}$.

Gross wage income is given by $\eta_{i,1,t}w_t$, where w_t is the aggregate and $\eta_{i,1,t}$ is the idiosyncratic wage component. Wage income is subject to social security contributions at rate τ, hence net wage income is $(1-\tau)\eta_{i,1,t}w_t$. During retirement, the household receives a lump-sum pension income, y_{t+1}^{ss}. As the household only cares about second period consumption and as there is neither satiation nor a bequest motive, the household consumes all resources in the second period of life. Accordingly, the budget constraints are given by

$$s_{i,2,t+1} = (1-\tau)\eta_{i,1,t}w_t \quad \text{and} \quad c_{i,2,t+1} = s_{i,2,t+1}R_{t+1} + y_{t+1}^{ss}, \tag{1}$$

where $s_{i,2,t+1}$ denotes gross savings and $R_{t+1} = 1 + r_{t+1}$ is a risky gross interest rate. From these two equations, one can see how social security can partially insure against idiosyncratic risk. While contributions to social security depend on η, all retirees receive the same lump-sum pension payments, y_{t+1}^{ss}. This constitutes an intragenerational sharing of idiosyncratic risk from the ex-ante perspective.

Aggregate wages and interest rates are stochastic.[6] We denote by ζ_t the shock to wages and by ϱ_t the shock to returns. We further assume that wages grow deterministically at rate λ. We therefore have:

$$w_t = \bar{w}_t\zeta_t = \bar{w}_{t-1}(1+\lambda)\zeta_t \quad \text{and} \quad R_t = \bar{R}\varrho_t, \tag{2}$$

where $\bar{R}$ and $\bar{w}_t$ are the deterministic components of returns and wages.

Social security is a pure PAYG system with lump-sum pension benefits. It is operated by the government, which is required to run a balanced budget every

[6]In this section, we limit the analysis to a partial equilibrium, and hence wages and returns are exogenous.

period. We abstract from population growth,[7] hence

$$\tau w_t = y_t^{ss}. \tag{3}$$

From equations (2) and (3), one can see how social security can provide partial insurance against aggregate risk. If ζ_t and ϱ_t are imperfectly correlated, then y_t^{ss} and R_t are imperfectly correlated. This hedge through social security constitutes an intergenerational sharing of aggregate risk from the ex-ante perspective.

2.2 Analysis

We analyze the welfare effects of introducing a marginal social security system in the two-generations model. That is, starting from a situation with zero contributions, $\tau = 0$, we study a marginal increase, $d\tau$, under the following assumptions:

Assumption 1. *All shocks $\eta_{i,1,t}, \zeta_t, \varrho_t$: (a) are distributed log-normal with means $\mu_{\ln\eta}$, $\mu_{\ln\zeta}$, $\mu_{\ln\varrho}$ and variances $\sigma^2_{\ln(\eta)}$, $\sigma^2_{\ln(\zeta)}$, $\sigma^2_{\ln(\varrho)}$, (b) have a mean of one: $E\zeta = E\varrho = E\eta = 1$, (c) are uncorrelated over time, and (d) are statistically independent from each other.*

Assumptions 1a-b are frequently employed for analytical tractability. Assumption 1c can be justified by the long periodicity of each period in a two-period overlapping generations model of approximately $30 - 40$ years. Assumption 1d is important to illustrate the CWG. We later relax it to extend the model by the CCV.

To evaluate welfare, we adopt an ex-ante perspective. The social welfare function of a cohort born in period t is the unconditional expected utility of a generation, $E\left[u(c_{i,2,t+1})\right]$. In our main results we look at the consumption equivalent variation (CEV) from a marginal introduction of social security. The CEV is the percentage increase in consumption, g_c, required to make the household indifferent between being born into an economy without social security ($\tau = 0$) and with a small social security system ($\tau = d\tau > 0$). We include a superscript PE for "partial equilibrium" to remain consistent with the subsequent quantitative analysis, which considers a general equilibrium. We also index the CEV by AR

[7] Our quantitative model instead also features population growth.

and IR to indicate presence of aggregate and idiosyncratic risk. This way, we can distinguish the CEV in an economy with both risks, $g_c^{PE}(AR, IR)$, from the CEV in an economy with only aggregate or only idiosyncratic risk, $g_c^{PE}(AR, 0)$ and $g_c^{PE}(0, IR)$, and the CEV in a deterministic economy, $g_c^{PE}(0, 0)$. We can now state our first proposition. A sketch of the proof is provided in Appendix A, while the complete proof is relegated to Supplementary Appendix B.

Proposition 1. *Under Assumption 1, the consumption equivalent variation from a marginal introduction of social security is given by*

$$g_c^{PE}(AR, IR) = \left(\frac{1+\lambda}{\bar{R}} \cdot \Psi(\sigma_{\ln AR}, \sigma_{\ln \eta}) - 1\right) d\tau \tag{4}$$

where $\sigma_{\ln AR} \equiv \sqrt{\sigma^2_{\ln \zeta} + \sigma^2_{\ln \varrho}}$ *and*

$$\Psi(\sigma_{\ln AR}, \sigma_{\ln \eta}) \equiv \exp\left(\theta\left(\sigma^2_{\ln \eta} + \sigma^2_{\ln AR}\right)\right) \geq \exp\left(\theta\sigma^2_{\ln \eta}\right) + \exp\left(\theta\sigma^2_{\ln AR}\right), \tag{5}$$

with the inequality being strict for $\sigma^2_{\ln \eta} > 0 \wedge \sigma^2_{\ln AR} > 0$.

To interpret this proposition, first consider a deterministic economy, where $g_c^{PE}(0, 0) = \left(\frac{1+\lambda}{\bar{R}} - 1\right) d\tau$. This reflects the well-known Aaron (1966) condition, namely that an expansion of the social security system increases welfare in a deterministic economy if (and only if) its implicit return exceeds the market rate of return, i.e., if and only if $(1+\lambda) > \bar{R}$.

In the non-degenerate stochastic case where $\sigma^2_{\ln \eta} > 0 \wedge \sigma^2_{\ln AR} > 0$, term Ψ captures an additional risk adjustment reflecting the inter-generational and intra-generational (partial) insurance provided by the system. We make the following important observations: First, Ψ is increasing in risk aversion θ, reflecting the standard intuition that more risk-averse households value insurance more. Second, Ψ increases in σ^2_{η} and σ^2_{ϱ}. This is due to the specific system with lump-sum benefits which, by construction, pools histories of idiosyncratic earnings risk in the cross-section and does not feature any return risk. Third, Ψ increases in σ^2_{ζ}. The reason is a standard hedging argument: social security reduces exposure to the wage shock, ζ, when young and increases it when old. Under independence, there consequently exist welfare gains from mixing both shocks as long as $\tau \in (0, 1)$. Fourth, Ψ is convex in total risk, $\sigma^2_{\ln AR} + \sigma^2_{\ln \eta}$. This mirrors an important result from the literature on the welfare costs of aggregate fluctuations, namely that the welfare gain of insuring against aggregate risk is a convex function of risk, cf. Lucas (1978), De Santis (2007), and Krebs (2007). Relative to this literature we study the effects of joint insurance and therefore total risk is the sum of the risk components. As a consequence of convexity, the whole welfare gain is greater than the sum of the gains from insurance against individual risk components, as in the inequality in (5). This difference in welfare terms is a consequence of the CWG, denoted as Δ_{CWG} in the sequel.

In order to further characterize the Δ_{CWG}, we next make it explicit by providing a formal definition of the individual components of the CEV. We use this definition to derive an exact and intuitive characterization of Δ_{CWG} for the case of log utility ($\theta = 1$) which constitutes a lower bound for preferences with $\theta > 1$.

Definition 1 (Components of CEV)**.** *Let $dg_c^{PE}(IR)$ and $dg_c^{PE}(AR)$ be the contributions to the CEV that are attributable to idiosyncratic and aggregate risk, defined by $g_c^{PE}(AR, IR) = g_c^{PE}(0,0) + dg_c^{PE}(AR) + dg_c^{PE}(IR) + \Delta_{CWG}$.*

Next rewrite Ψ in terms of variances of levels instead of variances of logs, i.e., $\Psi(\sigma_{AR}, \sigma_\eta)$ instead of $\Psi(\sigma_{\ln AR}, \sigma_{\ln \eta})$. Under Assumption 1a (log-normality of shocks), we get $\Psi(\sigma_{AR}, \sigma_\eta) \equiv \left(1 + \sigma_\eta^2 + \sigma_{AR}^2 + \sigma_\eta^2 \sigma_{AR}^2\right)^\theta$, where $\sigma_{AR} \equiv \sqrt{\sigma_\zeta^2 + \sigma_\varrho^2 + \sigma_\zeta^2 \sigma_\varrho^2}$. Employing Definition 1 for logarithmic utility ($\theta = 1$), the CEV writes as

$$
\begin{aligned}
&g_c^{PE}(AR, IR)\Big|_{\theta=1} = \\
&\quad \underbrace{\left(\frac{1+\lambda}{\bar{R}} - 1\right) d\tau}_{\equiv g_c^{PE}(0,0)|_{\theta=1}} + \underbrace{\frac{1+\lambda}{\bar{R}} \sigma_{AR}^2 d\tau}_{\equiv dg_c^{PE}(AR)|_{\theta=1}} + \underbrace{\frac{1+\lambda}{\bar{R}} \sigma_\eta^2 d\tau}_{\equiv dg_c^{PE}(IR)|_{\theta=1}} + \underbrace{\frac{1+\lambda}{\bar{R}} \sigma_{AR}^2 \sigma_\eta^2 d\tau}_{\equiv \Delta_{CWG}|_{\theta=1}} . \qquad (6)
\end{aligned}
$$

For logarithmic utility, the Δ_{CWG} is accordingly directly related to the product of variances of aggregate and idiosyncratic risk. By providing a lump-sum transfer, social security reduces the variance of retirement consumption[8] thereby reducing exposure to the total risk the household faces over the life-cycle, i.e., it reduces exposure to each risk component as well as their multiplicative interaction.

[8]Retirement consumption in the absence of social security is given by $\bar{w}_t \bar{R} \eta_{i,1,t} \zeta_t \varrho_{t+1}$. Its variance is $(\bar{w}_t \bar{R})^2 var(\eta_{i,1,t} \zeta_t \varrho_{t+1}) = (\bar{w}_t \bar{R})^2 (\sigma_\eta^2 + \sigma_{AR}^2 + \sigma_\eta^2 \sigma_{AR}^2)$, because the shocks are independent and have a mean of one, cf. Goodman (1960).

For standard random variables, a product of variances—as it enters the expression for the $\Delta_{CWG}|_{\theta=1}$—would be small and is usually ignored. However, we here deal with life-cycle earnings risk, hence with long horizons so that the single variance terms may well be large. To see this, let us make a rough back of the envelope calculation. Suppose a household works for 40 years, which in this two-generations model corresponds to the first period of a household's life. Assume further that each year the household receives a permanent idiosyncratic income shock with a log variance of 1 percent, corresponding to standard empirical estimates. Then, the bias captured by the $\Delta_{CWG}|_{\theta=1}$ adds approximately $\sigma_\eta^2 \cdot \sigma_{AR}^2 = \frac{\exp(40 \cdot 0.01)-1}{100} \cdot \sigma_{AR}^2 \approx 50\% \cdot \sigma_{AR}^2$ to the CEV, see equation (6).[9] Whatever the exact size of σ_{AR} is, this is clearly a non-negligible effect. Of course, our findings from this very stylized model should be interpreted with caution. Nevertheless, this calculation indicates that the effects may be large in a more realistically specified and appropriately calibrated quantitative model.

Finally, as convexity of the welfare gain is increasing in risk aversion, the contribution of each risk component to the CEV in equation (6) and the scaling due to the Δ_{CWG} constitute lower bounds for preferences with risk aversion above one.[10]

Modification: The Countercyclical Cross-sectional Variance

We modify the two period model slightly in order to illustrate the CCV mechanism. We alter Assumption 1 by conditioning the variance of idiosyncratic productivity risk on the aggregate state of the economy while its unconditional variance remains equal to $\sigma^2_{\ln \eta}$. We also focus on log utility and extend Definition 1 by the CCV:

Assumption 2. *(a) Let* $\zeta_t \in [\zeta_-, \zeta_+]$ *for all* t *where* $\zeta_+ > \zeta_- > 0$ *with* $\zeta_\pm = \chi \cdot \exp(1 \pm \sigma_{\ln_\zeta})$ *and probability* $\pi(\zeta_t = \zeta_+) = \pi(\zeta_t = \zeta_-) = \frac{1}{2}$ *with normalizing constant* $\chi = \frac{2}{\exp(1-\sigma_{\ln \zeta})+\exp(1+\sigma_{\ln \zeta})}$.[11] $\eta_{i,1,t}$ *is distributed as log-normal with a conditional variance given by* $\sigma^2_{\ln \eta \pm} = \sigma^2_{\ln \eta} \pm \Delta_\eta$ *for* $\zeta_t = \zeta_\pm$*, for some variance shifter* $\Delta_\eta \in (0, \sigma^2_{\ln \eta})$*, where it is understood that the rest of Assumption 1 continues to hold. (b) Utility is logarithmic, i.e.,* $\theta = 1$.

[9] By the random walk property of the income process, we have $\sigma_{\ln \eta} = 40 \cdot 0.01$. Under log-normality, we have $\sigma_\eta^2 = \exp(\sigma_{\ln \eta}) - 1$.

[10] This is formally shown in Supplementary Appendix B.

[11] It is straightforward to verify that $\mathbb{E}\zeta = 1$ and var$(\ln \zeta) = \sigma^2_{\ln \zeta}$.

Definition 2 (Components of CEV with CCV)**.** *The components of the total CEV with CCV can be isolated by $g_c^{PE}(AR, IR, CCV) = g_c^{PE}(0,0) + dg_c^{PE}(AR) + dg_c^{PE}(IR) + \Delta_{CWG} + \Delta_{CCV}$.*

We can now state our next result. A sketch of the proof is provided in Appendix A, a complete proof in Supplementary Appendix B.

Proposition 2. *Under Assumption 2 and using Definition 2 we get*

$$g_c^{PE}(AR, IR, CCV) = \left(\frac{1+\lambda}{\bar{R}} \cdot \exp\left(\sigma^2_{\ln \varrho}\right)\left(\frac{1}{\zeta_-}\exp\left(\sigma^2_{\ln \eta l}\right) + \frac{1}{\zeta_+}\exp\left(\sigma^2_{\ln \eta h}\right)\right) - 1\right) d\tau \tag{7a}$$

$$\textit{and} \quad \Delta_{CCV} = \frac{1+\lambda}{\bar{R}}\exp\left(\sigma^2_{\ln \varrho}\right)\Delta_\eta\left(\frac{1}{\zeta_-} - \frac{1}{\zeta_+}\right) d\tau. > 0 \tag{7b}$$

Equation (7a) is the analogue to equation (4) for discrete ζ and including CCV. Equation (7b) shows the increase of welfare gains through the CCV mechanism. This is due to the fact that the CCV raises (reduces) the variance of idiosyncratic productivity risk in states where the payoff in terms of consumption tends to be low (high) and is a consequence of concavity of the utility function. The amplification is therefore stronger the larger aggregate risk ($\sigma^2_{\ln \zeta}$ and $\sigma^2_{\ln \varrho}$) and the larger the variance shifter, Δ_η.

2.3 Discussion and Extensions

The preceding analysis abstracts from a number of important aspects that will be present in the quantitative model of the next section. First, we ignore first period consumption and thereby any means of self insurance through precautionary savings. Hence, the welfare benefits from insurance are overestimated. Second, this also implies that we miss the crowding out of savings through social security in general equilibrium which affects welfare negatively (in a dynamically efficient economy). The trade-off between the welfare gains from insurance and the welfare costs of crowing out is a central part of the quantitative analysis in the next sections. In Harenberg and Ludwig (2015) we demonstrate analytically that the bias in the welfare assessment of crowding out is ambiguous. By contrast, the previous section has shown that the two biases on the insurance side, Δ_{CWG}

and Δ_{CCV}, are unambiguously positive. One may therefore expect that jointly analyzing both risks induces a net positive welfare difference, but this can only be answered with a quantitative model. Third, the setup in our simple model is a situation in which a household faces additive and multiplicative background risk as studied in Franke, Schlesinger, and Stapleton (2011). Such uninsured background risk makes the agent behave as if he were more risk-averse.[12] In our analysis above we shut down the additive background risk by focusing on second period consumption only. In our full-blown quantitative analysis both additive and multiplicative background risk are at work and their impact on effective risk aversion will be captured by term Δ_{CWG}.

3 The Quantitative Model

We make several extensions to the two-generations model, but the main mechanisms remain the same. First, the periodicity is now one calendar year, and there are $J > 2$ overlapping generations. Consumption and savings decisions take place every period. Population grows at a constant rate, which acts as an additional implicit return to social security. Second, we consider a general equilibrium, which allows us to account for the costs of crowding out of capital. Third, we add a one-period, risk-free bond as a second asset. A household thereby has an additional asset to self-insure against idiosyncratic and aggregate risk. Ceteris paribus, this reduces the beneficial effects of social security. Fourth, labor income has a deterministic life-cycle component and idiosyncratic income shocks are allowed to be autocorrelated. Finally, we employ Epstein-Zin (Epstein and Zin 1989; Epstein and Zin 1991) preferences to partially disentangle risk aversion and the elasticity of inter-temporal substitution.

3.1 Time, Risk, and Demographics

Time is discrete and runs from $t = 0, \ldots, \infty$. At the beginning of each period t, an aggregate shock z_t hits the economy. For a given initial z_0, a date-event is uniquely identified by the history of shocks $z^t = (z_0, z_1, \ldots, z_t)$ where the z_t follow a Markov chain with finite support $\mathcal{Z}$ and nonnegative transition matrix

[12]The background risk literature started by considering additive background risk in portfolio choice problems, cf., e.g., Gollier and Pratt (1996). See also Harenberg and Ludwig (2015).

π_z. Thus, $\pi_z(z_{t+1}|z_t)$ represents the probability of z_{t+1} given z_t.

At every point in time t, the economy is populated by J overlapping generations indexed by $j = 1, \dots, J$. We denote the size of a generation by $N_j(z^t)$. Each generation consists of a continuum of households. We normalize the initial population size to unity, i.e., $\sum_{j=1}^{J} N_j(z_0) = 1$. Population grows at the exogenous rate of n, and there is no survival risk. Total population at t is therefore $N(z^t) = (1+n)^t$.

Households within a cohort are ex-ante identical but receive an idiosyncratic shock e_j each period so that there is ex-post intragenerational heterogeneity. We denote by e^j the history of idiosyncratic shocks. e_j follows a Markov chain with finite support $\mathcal{E}$ and strictly positive transition matrix π_e. The transition probabilities are $\pi_e(e_{j+1}|e_j)$, and the probability of a specific idiosyncratic shock history is $\pi_e(e^j)$. By a Law of Large Numbers $\pi_e(e^j)$ represents both the individual probability for e^j and the fraction of the population with that shock history.[13] Finally, $\pi_e(e_j)$ denotes the unconditional probability of shock e_j.

3.2 Households

At any date-event z^t, a household is fully characterized by its age j and its history of idiosyncratic shocks e^j. Denote by $u_j(c, e^j, z^t)$ the expected remaining life-time utility from consumption allocation c at age j, history e^j, and date-event z^t. Preferences are represented by a recursive utility function $u_j(c, \cdot)$ of the Epstein-Zin kind (Kreps and Porteus 1978; Epstein and Zin 1989; Epstein and Zin 1991; Weil 1989):[14]

$$u_j(c, e^j, z^t) =$$

$$\left(\left[c_j(e^j, z^t)\right]^{\frac{1-\theta}{\gamma}} + \beta \left(\sum_{z_{t+1}} \sum_{e_{j+1}} \pi_z(z_{t+1}|z_t) \pi_e(e_{j+1}|e_j) \left[u_{j+1}(c, e^{j+1}, z^{t+1})\right]^{1-\theta} \right)^{\frac{1}{\gamma}} \right)^{\frac{\gamma}{1-\theta}},$$

$$u_J(c, e^J, z^t) = c_J(e^J, z^t) , \qquad c > 0 ,$$

[13]Likewise, $\pi_e(e_{j+1}|e_j)$ represents both the individual transition probability and its population counterpart.

[14]In a slight abuse of notation, we use letter u to denote remaining lifetime utility in this recursive formulation, which was used in Section 2 to denote the per-period utility function.

where β is the discount factor and θ controls risk aversion. Parameter γ is defined as $\gamma \equiv \frac{1-\theta}{1-\frac{1}{\psi}}$ with ψ denoting the inter-temporal elasticity of substitution (IES). The CRRA utility specification is nested for $\theta = \frac{1}{\psi}$, which yields $\gamma = 1$.

Households inelastically supply one unit of labor until they retire at the fixed retirement age j_r. They are endowed with a deterministic life-cycle productivity profile ϵ_j. The idiosyncratic, stochastic component of income, $\eta(e_j, z_t)$, depends on the realization of idiosyncratic and aggregate shocks. The dependence of $\eta(e_j, z_t)$ on the aggregate shock is necessary to model the CCV. We assume that $E\left(\eta(e_j, z_t)|z_t\right) = 1$. Labor income is $y_j(e_j, z^t) = w(z^t)\epsilon_j\eta(e_j, z_t)$, where $w(z^t)$ is the real aggregate wage in terms of the consumption good at z^t. Insurance markets for labor income risk are closed by assumption.

Households can transfer wealth between periods by holding stocks and bonds in amounts $s_{j+1}(e^j, z^t)$ and $b_{j+1}(e^j, z^t)$, respectively. The stock has a risky return $r_s(z^{t+1})$ that depends on the realization of the aggregate shock in the following period, whereas the bond pays a one-period risk-free interest rate $r_b(z^t)$. The sequential budget constraint is standard:

$$
\begin{aligned}
&c_j(e^j, z^t) + s_{j+1}(e^j, z^t) + b_{j+1}(e^j, z^t) = (1 + r_s(z^t))s_j(e^j, z^t) \\
&\quad + (1 + r_b(z^{t-1}))b_j(e^j, z^t) + (1-\tau)y_j(e_j, z^t)I(j) + y^{ss}(z^t)(1 - I(j)),
\end{aligned}
$$

where τ is a fixed social security contribution rate, $y^{ss}(z^t)$ is pension income, and $I(j)$ is an indicator function that takes the value 1 if $j < j_r$ and 0 otherwise.[15] Households cannot die in debt, $s_{J+1}(e^J, z^t) + b_{J+1}(e^J, z^t) \geq 0$. Since there are no bequests, households are born with zero assets, i.e., $s_1(e^1, z^t) = b_1(e^1, z^t) = 0$.

3.3 Firms

There is a representative firm that produces $Y(z^t)$ using capital, $K(z^t)$, and labor, $L(z^t)$. The production technology is Cobb-Douglas with capital share α and deterministic labor-augmenting productivity growth λ. At each date-event, it is subject to a multiplicative shock to total factor productivity $\zeta(z_t)$,

[15] We do not consider an exogenous borrowing constraint. This may bias results in favor of social security because income (and asset) poor households can relax their budget constraint. With an exogenous borrowing constraint it would be natural to modify the social security system to have a progressive contribution rate with an exemption for income poor households so that negative welfare effects of social security contributions would be avoided for these households.

which depends only on the current aggregate shock, so that we have $Y(z^t) = \zeta(z_t)K(z^t)^\alpha((1+\lambda)^t L(z^t))^{1-\alpha}$.

Assuming a stochastic depreciation rate $\delta(z_t)$,[16] the capital stock evolves according to $K(z^t) = I(z^{t-1}) + K(z^{t-1})(1-\delta(z_{t-1}))$. Because of perfect competition, the firm remunerates the factors of production according to their marginal productivities. Thus wages, $w(z^t)$, and the return on capital, $r(z^t)$, are given by

$$w(z^t) = (1+\lambda)^t(1-\alpha)\zeta(z_t)\left(\frac{K(z^t)}{(1+\lambda)^t L(z^t)}\right)^\alpha, \tag{8a}$$

$$r(z^t) = \alpha\zeta(z_t)\left(\frac{(1+\lambda)^t L(z^t)}{K(z^t)}\right)^{1-\alpha} - \delta(z_t). \tag{8b}$$

The capital stock, $K(z^t)$, is financed by issuing stocks and bonds in quantities S and B, so that $K(z^t) = S(z^t) + B(z^t) = S(z^t)(1+\overline{\kappa}_f)$. The debt-equity ratio, $\overline{\kappa}_f$, is exogenous and constant. Therefore, the firm only decides on aggregate capital and not on the capital structure.[17] This mechanical leverage allows us to keep the depreciation shocks, which drive stock return volatility, small in the calibration. This is desirable, because large depreciation shocks imply unrealistically large fluctuations on the real side of the economy.[18] As demonstrated in Supplementary Appendix B, the leveraged stock return is

$$r_s(z^t) = r(z^t) + \overline{\kappa}_f\left(r(z^t) - r_b(z^{t-1})\right), \tag{9}$$

which shows that leverage increases mean and variance of stock returns.

3.4 Social Security

Social security works just like in the two-generations model of Section 2. The government organizes a PAYG system with a fixed contribution rate τ that is levied on labor income. Lump-sum pension income $y^{ss}(z^t)$—that does not

[16]The same assumption is employed by Storesletten, Telmer, and Yaron (2007), Gomes and Michaelides (2008), Krueger and Kubler (2006), among others.

[17]Leverage is frequently modeled this way in the finance literature to increase the volatility of stock returns, cf., e.g., Boldrin, Christiano, and Fisher (1995) and Croce (2014).

[18]In the baseline calibration, the depreciation shocks are so small that $\overline{\kappa}_f = 0$ delivers the same results. Only in the calibration for sensitivity analyses where we match the volatility of stock returns, larger depreciation shocks would be needed when $\overline{\kappa}_f = 0$.

depend on the idiosyncratic history—adjusts to ensure that the social security budget is balanced at every date-event. Denoting by $P(z^t)$ the number of pensioners, $P(z^t) = \sum_{j=j_r}^{J} N_j(z^t)$, the budget constraint reads as $\tau w(z^t)L(z^t) = y^{ss}(z^t)P(z^t)$.

3.5 Equilibrium

We study a competitive general equilibrium, where households and firms maximize and all markets clear. In the computational solution, we focus on recursive Markov equilibria. We express all aggregate variables in terms of labor efficiency units, i.e., we divide aggregate variables by $(1+\lambda)^t L(z^t) = (1+\lambda)^t \sum_{j=1}^{j_r-1} \epsilon_j N_j(z^t)$. The corresponding normalized variable is written in lower case, e.g., $k(z^t) = \frac{K(z^t)}{(1+\lambda)^t L(z^t)}$. Individual variables are detrended only by the level of technology, and the corresponding variables are denoted with a tilde, e.g., $\tilde{c}_j(\cdot) = \frac{c_j(\cdot)}{(1+\lambda)^t}$. Accordingly, the monotone transformation of utility is denoted by $\tilde{u}_j(\cdot)$. Since the model has (ex-post) heterogeneous households and aggregate uncertainty, the distribution of households becomes part of the state space. We denote by Φ the distribution of households over age, current income state, stocks, and bonds. The corresponding equilibrium law of motion, $\Phi' = H(\Phi, z)$, is induced by household's optimal choices and the exogenous shock processes. Every period there are five markets that clear: consumption good, capital, labor, stocks, and bonds. A precise definition of the recursive Markov equilibrium is relegated to Supplementary Appendix B.

3.6 Computational Solution

We compute an equilibrium of our model by applying the Krusell and Smith (1998) method.[19] To approximate the law of motion of the distribution, $H(\Phi, z)$, we consider various forecast functions, $\hat{H}$, of the average capital stock and the ex-ante equity premium and select the one with the best fit. The average goodness of fit of the selected approximate law of motion is in the range of $R^2 = 0.99$ for all of the calibrations. The state space is further reduced by one dimension by recasting the problem in terms of cash-on-hand. To speed up the solution, we employ a variant of the endogenous grid method (Carroll 2006) that allows for

[19] Also see, e.g., Gomes and Michaelides (2008) and Storesletten, Telmer, and Yaron (2007).

two continuous choices. Details of the computational solution are provided in Supplementary Appendix C.

3.7 Welfare Criterion

We employ the same welfare concept as in the two-generations economy of Section 2, namely ex-ante expected utility of a household at the start of economic life. As explained in Davila, Hong, Krusell, and Ríos-Rull (2012), in an economy with ex-ante identical but ex-post heterogeneous agents, this concept represents a natural objective for a social planner who is behind the Rawlsian veil of ignorance. It is a Utilitarian welfare criterion, which weighs lifetime utilities with their respective probabilities.

A household's welfare of being born into an economy with policy A can be written as $E\left[\tilde{u}_1(\tilde{c}^A, e^1, z^t)\right]$, where the expectation is taken over all date-events z^t. Consequently, it is an expectation over all possible equilibrium values of aggregate capital and prices. As before, we express the welfare difference when comparing policy A to policy B in terms of a consumption equivalent variation, g_c. As we prove in Supplementary Appendix B, it is given by

$$g_c = \frac{E\left[\tilde{u}_1(\tilde{c}^B, e^1, z^t)\right]}{E\left[\tilde{u}_1(\tilde{c}^A, e^1, z^t)\right]} - 1. \tag{10}$$

A positive number indicates the percentage of lifetime consumption a household would be willing to give up under policy A in order to be born into an economy with policy B. By adopting an ex-ante perspective, we compare the long-run welfare effects of such a reform. While this does not include the transition between the two economies, it is important to understand that for the experiment described below (an introduction of social security), including the welfare effects along the transition would increase g_c. The reason is that moving from policy A to policy B implies a gradual decrease in capital. Thus, generations that live through the transition experience the full benefits from insurance but are spared some of the long-run welfare costs of crowding out. Therefore, by ignoring the transition, we calculate a lower bound on the welfare effects.

3.8 Experiment and Decomposition Analysis

In terms of the previous section, our computational experiment consists of comparing policy A, which has a social security contribution rate of $\tau = 0\%$, to policy B, which has $\tau = 2\%$. This is the experiment performed by Krueger and Kubler (2006). It can be interpreted as the introduction of a marginal social security system, or of a minimum pension, as in Section 2.

In general equilibrium, this experiment unambiguously leads to a lower capital stock because private savings are crowded out. Naturally, the reduction in aggregate capital leads to changes in relative prices—wages decrease and returns increase. We call the economy dynamically inefficient if the reduction in capital and the induced price changes per se lead to a welfare gain. When calibrating the model, we always make sure that the economy is dynamically efficient to avoid that welfare gains stem from a mitigation of overaccumulation of capital.

To separate the welfare gains of insurance from the welfare losses of crowding out, we perform a partial equilibrium experiment. In this partial equilibrium, the social security system changes, but prices, i.e., wages and returns, remain unaffected. Conceptually, it corresponds to a small open economy with free movement of the factors of production. To formalize this, denote by $\mathcal{P}_A = \left\{\{z^t, r(z^t), r_s(z^t), r_b(z^t), \tilde{w}(z^t)\}_{t=1}^{T} | \tau = 0\%\right\}$ the sequence of shocks and prices obtained from the general equilibrium of the economy without a social security system, i.e., under policy A ($\tau = 0\%$). Likewise, denote by $\hat{H}_A$ the approximate law of motion of this equilibrium. We compute the partial equilibrium under the old price sequence $\mathcal{P}_A$ and the old laws of motion $\hat{H}_A$, but with policy B ($\tau = 2\%$). The welfare gains stemming from insurance are then:

$$g_c^{PE} = \frac{E\left[\tilde{u}_1(\tilde{c}^B, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 2\%\right]}{E\left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 0\%\right]} - 1. \tag{11}$$

Analogously, the corresponding general equilibrium number is

$$g_c^{GE} = \frac{E\left[\tilde{u}_1(\tilde{c}^B, e^1, z^t) | \mathcal{P}_B, \hat{H}_B, \tau = 2\%\right]}{E\left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 0\%\right]} - 1, \tag{12}$$

where the crucial difference is that in the new equilibrium with policy B ($\tau = 2\%$), households choose consumption optimally given the new general equilibrium prices and laws of motion, $\mathcal{P}_B, \hat{H}_B$. The welfare costs of crowding out (CO) are given by the difference $g_c^{CO} = g_c^{GE} - g_c^{PE}$. In a dynamically efficient economy as defined above, g_c^{CO} is negative by definition.

The final step is the decomposition of g_c^{PE} into insurance against aggregate risk, idiosyncratic risk, as well as the two biases, i.e., the differences in overall welfare that are attributable to CWG and CCV, respectively. Recalling our decomposition of the CEV in Section 2, Definitions 1 and 2, we get:

$$\begin{aligned} g_c^{PE}(AR, IR, CCV) &= g_c^{PE}(0,0) \quad + dg_c(AR) + dg_c(IR) + \Delta_{CWG} + \Delta_{CCV} \\ g_c^{PE}(AR, IR) &= g_c^{PE}(0,0) \quad + dg_c(AR) + dg_c(IR) + \Delta_{CWG} \\ g_c^{PE}(0, IR) &= g_c^{PE}(0,0) \quad \qquad\qquad\quad + dg_c(IR) \\ g_c^{PE}(AR, 0) &= g_c^{PE}(0,0) \quad + dg_c(AR) \end{aligned}$$

The right-hand side of the first line shows all of the components. To isolate those, we compute $g_c^{PE}(AR, 0)$ and $g_c^{PE}(0, 0)$, as in equation (11), but for an economy with only aggregate risk and one without risk, respectively.[20] With those numbers at hand, we can back out the welfare effect attributable to aggregate risk, $dg_c(AR)$. Likewise, we compute $g_c^{PE}(0, IR)$ for an economy featuring only idiosyncratic risk to back out $dg_c(IR)$. Next, we compute $g_c^{PE}(AR, IR)$. As we already know $dg_c(AR)$ and $dg_c(IR)$, we can back out how much of the welfare effects are attributable to the CWG, the Δ_{CWG}. In the same manner, we obtain Δ_{CCV}.

In summary, this decomposition procedure allows us to isolate the welfare effects in a very consistent manner, because the procedure is performed in partial equilibrium, and hence prices, shocks, and model parameters are identical in all computations. These welfare numbers are consistent with the general equilibrium results, because that is where the original equilibrium sequences and laws of motion come from.

4 Calibration

The selection of targets and parameters to be calibrated is informed by our theoretical insights, in particular Propositions 1 and 2. These indicate that the coefficient of relative risk aversion, θ, the variances of the shocks, and the returns

[20] As shown in the Supplementary Appendix B.4, $g_c(0, 0)$ can be calculated from the present discounted value of lifetime income, independent of preference parameters.

on savings are crucial in determining the value of social security. Guided by this, our baseline calibration takes a very conservative approach, in the sense that it features a low θ and small aggregate shocks. In the sensitivity analysis, we then first increase θ to match the Sharpe ratio, $\varsigma = \frac{E\left[r_{s,t}-r_{b,t}\right]}{\sigma\left[r_{s,t}-r_{b,t}\right]}$, and then aggregate shocks to match the equity premium, $\mu = E\left[r_{s,t} - r_{b,t}\right]$, see Section 5.3.

One set of parameters is determined exogenously by either taking its value from other studies or measuring its value in the data. We refer to these parameters as first-stage parameters. The second set of parameters is jointly calibrated by matching the model-simulated moments to their corresponding moments in the data. Accordingly, we refer to those parameters as second-stage parameters.[21]

Table 1 summarizes our conservative baseline calibration, described next. Additional information on our empirical approach to measure calibration targets and on the numerical implementation of the procedure is provided in Supplementary Appendices D and C, respectively.

4.1 Demographics

Households begin working at the biological age of 21, which corresponds to $j = 1$. We set $J = 58$, implying a life expectancy at birth of 78 years, which is computed from the Human Mortality Database (HMD) for year 2007. We set $j_r = 45$, corresponding to a statutory retirement age of 65. Population grows at a rate of 1.1%, which reflects the current growth trend of the U.S. population.

4.2 Households

In our baseline calibration, we treat the coefficient of risk aversion as a fist-stage parameter, setting it to 3, which is well within the standard range of $[2, 4]$. Given this choice, our model produces a Sharpe ratio of $\varsigma = 0.079$ and an equity premium of $\mu = 0.75\%$. These are by a factor of 4.2 and 7.4 lower than their respective empirical estimates of 0.33 and 5.60%.[22]

The inter-temporal elasticity of substitution (IES) is set to 0.5. This is at the lower end of the range of values used in the literature, as reviewed, e.g., by Bansal and Yaron (2004). A higher value of the IES means that households

[21] The second-stage parameters jointly determine all targeted moments. When we note that we calibrate a parameter to a target, we mean that it has the strongest impact on that target.

[22] Calculated from Robert Shiller's website, see http://aida.wss.yale.edu/$\sim$shiller/data.htm.

Table 1: Summary of the Baseline Calibration

Parameter	Value	Target (source)	Stage
Demographics			
Biological age at $j = 1$	21		1^{st}
Model age at retirement, j_r	45		1^{st}
Model age maximum, J	58	Life expectancy at birth	1^{st}
Population growth, n	0.011	U.S. Social Sec. Admin. (SSA)	1^{st}
Households			
Discount factor, β	0.981	Capital output ratio, 2.65 (NIPA)	2^{nd}
Coefficient of relative risk aversion, θ	3.0		1^{st}
Inter-temporal elasticity of substitution, ψ	0.5		1^{st}
Age productivity, $\{\epsilon_j\}$	Cf. Appendix D	Earnings profiles (PSID)	1^{st}
CCV, $\sigma_\nu^2(z)$	$\{0.0445, 0.0156\}$	Storesletten, et al. (2007)	1^{st}
Autocorrelation of log income, ρ	0.952	Storesletten, et al. (2007)	1^{st}
Firms			
Capital share, α	0.32	Wage share (NIPA)	1^{st}
Leverage ratio, $\overline{\kappa}_f$	0.66	Rajan and Zingales (1995)	1^{st}
Technology growth, λ	0.018	TFP growth (NIPA)	1^{st}
Mean depreciation rate of capital, δ_0	0.102	Bond return, 0.023 (Shiller)	1^{st}
Aggregate Risk			
Standard deviation of depreciation, $\bar{\delta}$	0.078	Std. of consumption growth, 0.030 (Shiller)	2^{nd}
Aggregate productivity states, $1 \pm \bar{\zeta}$	0.029	Std. of TFP, 0.029 (NIPA)	1^{st}
Transition probabilities of productivity, π^ζ	0.941	Autocorrelation of TFP, 0.88 (NIPA)	1^{st}
Conditional prob. of depreciation shocks, π^δ	0.886	Corr.(TFP, returns), 0.50 (NIPA, Shiller)	2^{nd}

Notes: 1^{st} stage parameters are set exogenously, 2^{nd} stage parameters are jointly calibrated to the targets.

react more strongly to price changes, with the consequence that welfare losses from crowding out are lower which we investigate in our sensitivity analysis.

In our baseline calibration, the discount factor β is calibrated to match the capital output ratio of 2.65, which we calculate from NIPA data.[23] We obtain $\beta = 0.981$.

The age-specific productivity profile ϵ_j is computed from PSID data by applying the method of Huggett, Ventura, and Yaron (2011). It is displayed in Supplementary Appendix D. The stochastic component of labor income, $\eta(e_j, z_t)$, is based on Storesletten, Telmer, and Yaron (2007), who also provide an estimate of the countercyclical cross-sectional variance of income risk, CCV. These comprise an autocorrelation coefficient of log income of $\rho = 0.952$ and a conditional variance of innovations, $\sigma^2_\nu(z_t)$, of 0.0445 in recessions and 0.0156 in booms. We approximate this process through a discrete, four-state Markov process using the Rouwenhorst method, cf. Kopecky and Suen (2010).

4.3 Firms

We set the value of the capital share parameter to $\alpha = 0.32$. This is directly estimated from NIPA data on total compensation as a fraction of GDP. Our estimate of the deterministic trend growth rate is based on data on total factor productivity (TFP). The point estimate is $\lambda = 0.018$, which is in line with other studies. Leverage in the firm sector is set to $\overline{\kappa}_f = 0.66$, based on Rajan and Zingales (1995).

The mean depreciation rate of capital, δ_0, is a second-stage parameter. We calibrate it so as to match the average bond return of 2.3%.[24] In economies without aggregate risk we calibrate δ_0 to produce a risk-free return of 4.2%, which corresponds to the empirical estimate of Siegel (2002).[25]

4.4 Aggregate Risk

Aggregate risk is driven by a four-state Markov chain with support $\mathcal{Z} = \{z_1, ..., z_4\}$ and transition matrix π_z. Each aggregate state maps into a combination of technology shock and depreciation shock, $(\zeta(z), \delta(z))$. Both shocks can take a high

[23] Our estimate is in line with the estimates of, e.g., Fernández-Villaverde and Krueger (2011).
[24] Again calculated from the data on Robert Shiller's website.
[25] Cf. our discussion on the decomposition of welfare effects in Subsection 3.8.

and a low value, given by $\zeta(z) = 1 \pm \bar{\zeta}$ and $\delta(z) = \delta_0 \pm \bar{\delta}$. We denote the transition probability of remaining in the low technology state by π^{ζ}. To govern the correlation between technology and depreciation shocks, we let the probability of being in the high depreciation state conditional on being in the low technology state be π^{δ}. Assuming symmetry of the transition probabilities, the Markov chain of aggregate shocks is characterized by four parameters, $(\bar{\zeta}, \bar{\delta}, \pi^{\zeta}, \pi^{\delta})$, see Supplementary Appendix D.3 for details. We set $\bar{\zeta}$ and π^{ζ} to match the standard deviation and autocorrelation of TFP in the data of 0.029 and 0.88, both estimated using NIPA data. The remaining parameters, $\bar{\delta}$ and π^{δ}, are calibrated as second-stage parameters to jointly match the standard deviation of aggregate consumption growth of 0.03 and the correlation of the cyclical component of TFP with risky returns of 0.5. We get $\bar{\zeta} = 0.029$, $\bar{\delta} = 0.078$, $\pi^{\zeta} = 0.941$, and $\pi^{\delta} = 0.886$.

5 Results

5.1 Baseline Calibration

Aggregate Effects. The effects of introducing social security at a contribution rate of 2% on capital accumulation, prices and welfare are documented in Table 2. Our experiment leads, on average, to a long-run reduction in the capital stock of 11.97%, which is accompanied by a 3.93% reduction in wages, an increase in the return on stocks of 1.02 percentage points, and an increase in the return on bonds of 1.04 percentage points. The average return on bonds increases to a greater extent, because the insurance provided through social security leads households to rebalance their portfolios towards stocks. This reduces relative demand for bonds, decreasing their price and increasing their return.

Table 2 also reports the consumption equivalent variation, g_c^{GE}, as defined in equation (12). The reform yields a CEV of 2.21% despite the sizeable crowding out of capital. This constitutes a substantial welfare gain from a minimum pension at a contribution rate of 2%. Below, we decompose this welfare gain into its various components.

Finally, Table 2 reports a small increase in aggregate consumption. This does not mean that the economy is dynamically inefficient. The reduction in aggregate capital, per se, leads to lower consumption and a welfare loss, as

Table 2: Aggregate Effects of The Social Security Experiment

Variable	Change
Aggregate capital, K	$\Delta K/K = -11.97\%$
Aggregate wage, w	$\Delta w/w = -3.93\%$
Stock return, r_s	$\Delta r_s = +1.02pp$
Bond return, r_b	$\Delta r_b = +1.04pp$
Aggregate consumption, C	$\Delta C/C = +0.24\%$
Consumption equivalent variation	$g_c^{GE} = +2.21\%$

Notes: $\Delta X/X$ is the expected percent change in variable X between two steady states, i.e., $\Delta X/X = \frac{E(X_t|\tau=2\%)-E(X_t|\tau=0\%)}{E(X_t|\tau=0\%)}$. Δx is the change in variable x expressed in percentage points (*pp*), i.e., $\Delta x = E(x_t|\tau = 2\%) - E(x_t|\tau = 0\%)$. g_c^{GE} is the consumption equivalent variation in general equilibrium, cf. Subsection 3.8.

we will see in the following paragraph. The reason for the mildly increased aggregate consumption is that, even though we operate in a dynamically efficient economy, the implicit return of social security exceeds the expected one-period risk-free rate.[26] While such an increase in aggregate consumption contributes to the welfare gain, the contribution is small. In fact, in the sensitivity analysis of Subsection 5.3, we frequently find a decrease in aggregate consumption with accompanying welfare gains that are larger than those reported here.

Benefits from Insurance versus Costs from Crowding Out. To further understand the nature of the welfare gain, we decompose it into the benefits from insurance and the losses from crowding out by conducting the partial equilibrium (PE) experiment described in Subsection 3.8. Accordingly, the sequences of wages and returns before and after the introduction of social security are identical. As a consequence, the CEV in this experiment reflects purely the benefits from insurance. Subtracting this number from the overall welfare gain reported in Table 2 yields the losses from crowding out. As Table 3 reveals, the net welfare gains attributable to the total insurance provided by social security amount to $+5.48\%$ and the losses from crowding out stand at -3.27%. The fact that the reduction in aggregate capital leads to a welfare loss means that the equilibrium

[26]This is consistent with Proposition 1 of Krueger and Kubler (2006), which states that for dynamic efficiency, it is sufficient that, at every point in time, there are states of the world with positive probability in which the bond return is larger than the implicit social security return.

is dynamically efficient, cf. Subsection 3.8.

Table 3: Benefits from Insurance versus Costs from Crowding Out

Variable	GE	PE	CO
CEV, g_c	+2.21%	+5.48%	-3.27%

Notes: For the definition of the CEV, g_c, in general equilibrium (GE) and partial equilibrium (PE), cf. Subsection 3.8. The difference of the two CEVs is the loss from crowding out (CO).

Distributional Effects. The welfare effects reported in Table 3 represent aggregate effects, which encompass important inter- and intragenerational distributional changes. To see how these distributional changes realize in the quantitative model, Figure 1 displays the average life-cycle consumption in Panel (a) and the variance of log consumption over the life-cycle in Panel (b). The introduction of social security in partial equilibrium leads to better consumption insurance and therefore reduces precautionary savings. Consequently, the consumption profile is pivoted clockwise such that households consume more on average in early stages of the life-cycle at the expense of slightly reduced average consumption when older than 52 years. Due to discounting, the early consumption gains are weighted more strongly than the later consumption losses. Simultaneously, the variance of log consumption decreases over the life-cycle. Both effects underlie the strong partial equilibrium aggregate welfare gain reported in Table 3.

In the post-experiment general equilibrium, the consumption profile is pivoted counter-clockwise. The reason is that the crowding out of capital now leads to lower wages and higher average returns, as reported previously in Table 2. In response to higher returns, households increase their life-cycle savings when young and increase consumption when old. Consumption remains below its pre-experiment, general equilibrium level until age 43. At the same time, the variance of log consumption is slightly higher than in the pre-experiment economy until age 40. This is a consequence of the increased life-cycle savings, which leads households to also bear more volatile consumption when young. After the age of 40, the variance of log consumption is smaller than in the pre-experiment economy, and the gap widens considerably with age. This reduced volatility at older ages constitutes the dominating source of the welfare gain reported in Table 2.

Figure 1: Life-cycle Consumption

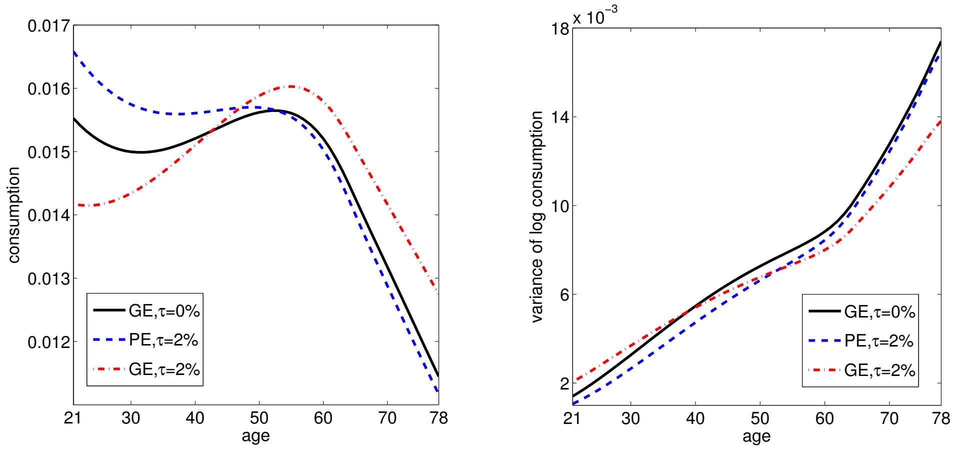

(a) Average Life-cycle Consumption (b) Variance of Log Consumption

Notes: Average consumption in Panel (a) and variance of log consumption in Panel (b) at each age for the economy without social security (GE, $\tau = 0\%$), with social security (GE, $\tau = 2\%$), and the partial equilibrium with social security and old prices (PE, $\tau = 2\%$).

To conclude the discussion of the distributional consequences, Table 4 reports the Gini coefficients for assets, labor earnings, and consumption. We make three observations. First, the simulated Gini coefficients closely align with the data reported by, e.g., Hintermaier and Koeniger (2011). This is notable because they were not a target in the calibration, and it is not easy to match them.[27] The very close match substantially strengthens our analysis of the welfare consequences of the reform. Second, the Gini coefficient for assets increases. This is so because households take on more risky portfolio compositions in response to the introduction of social security and because of higher average returns, see Table 2. Third, improved consumption insurance leads to a lower degree of consumption dispersion in the economy. The Gini coefficient for consumption decreases slightly.

Decomposition into Risks. We investigate how much of the welfare gains in partial equilibrium of $+5.48\%$ can be attributed to insurance against aggregate risk and against idiosyncratic risk, to the direct interaction between risks in form of the CCV, and to the convexity of the welfare gain, CWG.[28] Results are

[27] See, e.g., Castañeda, Díaz-Giménez, and Ríos-Rull (2003) and De Nardi (2004).

[28] See Subsection 3.8 for the decomposition procedure.

Table 4: Distributional Consequences: Gini Coefficients

Variable	$\tau = 0.00$	$\tau = 0.02$	Change	Data
Assets	0.764	0.806	4.21pp	0.809
Earnings	0.437	0.437	0.00pp	0.439
Consumption	0.262	0.260	-0.20pp	

Notes: pp stands for percentage points. Estimates in column "Data" are taken from Hintermaier and Koeniger (2011).

summarized in Table 5. The consumption equivalent variation in a deterministic environment is negative at -0.62%, because the implicit return of social security of $\lambda + n = 0.018 + 0.01 = 0.028$ is below the interest rate of $r_b = 0.042$, our target in the risk-free economy, cf. the discussion in Section 4.[29] The welfare gains from insurance against idiosyncratic risk amount to 0.84% and against aggregate risk to 2.02% in terms of consumption equivalent variations. Hence, the role played by aggregate risk is approximately twice as important as the role played by idiosyncratic risk. Turning to the direct interactions between risks, the difference in welfare attributable to the CCV, the Δ_{CCV}, is at 1.60%. The Δ_{CWG} is of similar size. Recall the back-of-the-envelope calculation in our simple model of Section 2.2, where we found that Δ_{CWG} scales the gains due to insurance against aggregate risk by approximately 50%. The corresponding ratio stands at 82% for our baseline scenario. The reason it is larger is related to the fact that the convexity of the welfare gain increases in risk aversion, which is $\theta = 3$ in our baseline calibration, versus $\theta = 1$ in the back-of-the-envelope calculation.

A crucial point is that the two welfare differences, the Δ_{CWG} and the Δ_{CCV}, jointly account for approximately 60% of the total insurance gains through social security, calculated as $\frac{\Delta_{CCV}+\Delta_{CWG}}{g_c^{PE}} \cdot 100[\%]$. Combining the findings from the previous literature which focuses only on one risk therefore leads to substantial quantitative biases in the welfare assessments of social security.

[29]This welfare number is computed analytically, cf. Corollary 1 in Supplementary Appendix B.4.

Table 5: Decomposition of Welfare Benefits in Partial Equilibrium

g_c^{PE}	$g_c^{PE}(0,0)$	$dg_c(IR)$	$dg_c(AR)$	Δ_{CWG}	Δ_{CCV}
5.48% =	-0.62%	+0.84%	+2.02%	+1.65%	+1.60%

Notes: This table presents the decomposition of the welfare gain expressed as consumption equivalent variation, g_c^{PE}, into various sources, cf. Subsection 3.8.

5.2 On the Importance of Modeling both Risks

The analysis of our baseline scenario suggests that the role played by the two biases is large. To investigate whether it is indeed the joint presence of both risks (aggregate and idiosyncratic risk) as well as their interactions that lead us to conclude that social security is beneficial in the long run, we compute the general equilibria of economies that feature only aggregate risk, only idiosyncratic risk, or no risk. We calibrate each economy to standard targets in the literature. For the economy without idiosyncratic risk, we adopt the targets of Krueger and Kubler (2006) and match the equity premium, $\mu = E\left[r_{s,t} - r_{b,t}\right]$, and the volatility of stock returns. Specifically, we target an equity premium of $\mu = 5.60\%$ and a standard deviation of stock returns of $\sigma(r_s) = 16.8\%$.[30] For the economy without aggregate risk and the deterministic economy, we target an interest rate of 4.2%, the same rate used in the PE decomposition procedure for economies without aggregate risk. Throughout these experiments, we target a capital output ratio of 2.65 by adjusting the discount factor, β. Details of the calibration are described in Supplementary Appendix D.

Our analysis of Section 2 shows that welfare gains from introducing social security increase in risk aversion and the volatility of aggregate risk. With respect to these two, the AR-$only$ calibration is an extreme case in that it features very high aggregate risk and high risk aversion. But even with such an extreme calibration, Table 6 documents welfare losses for this case. In general equilibrium, i.e., accounting for the welfare losses from crowding-out of capital, they stand at -0.54%, again expressed as a consumption equivalent variation. Even in the short-run, the benefits from insurance through social security do not dominate, as the welfare losses are -0.38% in partial equilibrium. Also in the economy featuring only idiosyncratic risk, denoted IR-$only$, we find large welfare losses in general equilibrium of -1.92%. In partial equilibrium, there are modest

[30]Again based on data taken from Rob Shiller's webpage.

Table 6: The Role of Both Risks: Benefits and Costs

	Consumption equivalent variation, g_c		
Scenario	GE	PE	CO
AR-$only$	-0.54%	-0.38%	-0.16%
IR-$only$	-1.91%	0.27%	-2.18%
No-$risk$	-1.13%	-0.62%	-0.51%

Notes: CO: crowding out; AR-$only$: economy with only aggregate risk, calibrated to match equity premium; IR-$only$: economy with only idiosyncratic risk; No-$risk$: deterministic economy.

welfare gains. Finally, introducing social security in the no-risk economy leads to welfare losses in both general and partial equilibrium.

5.3 Sensitivity Analysis

We now investigate the sensitivity of our results with respect to the calibration targets. Specifically, we are interested in whether our key findings of long-run welfare gains and sizeable interactions are robust when we consider economies with alternative levels of risk and risk aversion. In one variant we calibrate the model to match the equity premium and the volatility of stock returns, the same targets as in the AR-only economy of Subsection 5.2. This scenario is referred to as EP. It implies a consumption volatility and a Sharpe ratio which are both to high relative to the data. We therefore also consider an intermediate case where we match the Sharpe ratio and the volatility of consumption. This scenario is referred to as SR. To isolate the effects of risk and risk aversion, we hold the discount factor β constant at its baseline value in the experiments. Therefore, a crucial preference parameter—which has a strong impact on welfare—remains unchanged, making the comparison and interpretation of the results much easier. The expected, risk-free bond return, r_b, is always kept at the same level of the baseline through an appropriate calibration of δ_0.

We repeat this sensitivity analysis with a higher IES. To this end, we first proceed as in our baseline calibration, i.e., for our choice of risk aversion of $\theta = 3$, we define a modified baseline ($BL_{IES=1.5}$) in which we set the IES to 1.5 and recalibrate all parameters. Starting from this modified baseline, we then repeat the analogues to the SR and EP calibrations, referred to as $SR_{IES=1.5}$

and $EP_{IES=1.5}$, respectively. Details on the calibration are described in Supplementary Appendix D.

The welfare results in general equilibrium are presented in Table 7, together with the benefits from insurance and the losses from crowding out of capital formation. Throughout, our result from the baseline scenario (BL) is confirmed: there are large welfare gains ranging from 2 to 5 percent in terms of consumption equivalent variation when losses from crowding out are fully taken into account. Welfare gains increase in risk aversion: BL features $\theta = 3$, SR has $\theta = 11$ and EP has $\theta = 5.6$.[31] Our two baseline scenarios, which have an IES of 0.5 and 1.5 and a reasonable degree of risk aversion of 3, deliver the smallest welfare numbers with total welfare gains of 2.2% and 2.5%, respectively.[32]

Table 7: Sensitivity Analysis: Benefits, Costs, and Bias

Scenario	Consumption equivalent variation, g_c GE	PE	CO	$\frac{\Delta_{CCV}+\Delta_{CWG}}{g_c^{PE}}$
$IES = 0.5$				
BL	+2.21%	+5.48%	-3.27%	0.59
SR	+4.16%	+8.32%	-4.16%	0.65
EP	+2.78%	+6.53%	-3.75%	0.67
$IES = 1.5$				
$BL_{IES=1.5}$	+2.56%	+3.65%	-1.10%	0.60
$SR_{IES=1.5}$	+5.08%	+8.09%	-3.01%	0.65
$EP_{IES=1.5}$	+4.45%	+7.34%	-2.89%	0.71

Notes: CO: crowding out; BL: baseline calibration with $\theta = 3$; SR: scenario matching Sharpe ratio; EP: scenario matching equity premium.

Losses from crowding out are smaller in all the $IES = 1.5$ scenarios than their $IES = 0.5$ counterparts. This is because households react more strongly to the change in interest rates induced by lower capital formation. Higher interest rates make them save more, thus dampening the crowding out.[33] This is the main

[31] Ceteris paribus, matching a higher equity premium would require a higher degree of risk aversion. However, as we simultaneously increase the variance of risky returns—by an appropriate choice of $\bar{\delta}$—, we also introduce more risk into the economy. As a consequence, the coefficient of risk aversion is lower in the EP than in the SR-calibration.

[32] We also conducted an experiment with a risk aversion of 2 (an IES of 0.5). Then welfare gains in general equilibrium are at 1.47%, confirming the monotonous impact of risk aversion.

[33] For the same reason, a higher IES leads to a smaller volatility of real aggregates. The

reason why the general equilibrium welfare gains in the $IES = 1.5$ scenarios are always larger than in their $IES = 0.5$ counterparts.

Finally, to examine the role played by the welfare biases attributable to the CWG and the CCV across scenarios, we compute the ratio $\frac{\Delta_{CCV}+\Delta_{CWG}}{g_c^{PE}}$. It amounts to approximately 60 percent in our two baseline scenarios and reaches 71 percent in scenario $EP_{IES=1.5}$. Therefore, our finding that roughly 60% of the total welfare gains would be missed from adding up the isolated benefits is robust across calibrations.[34]

6 Conclusion

This paper extends the previous literature on the welfare effects of social security by evaluating its benefits and costs when households face multiple risks. We include both idiosyncratic and aggregate risk in a life-cycle model and consider a pay-as-you-go (PAYG) financed social security system which partially insures both risks through a minimum lump-sum pension. We show that the whole gain from insurance is greater than the sum of the insurance benefits attributable to the isolated risk components. One source for this welfare difference is a direct interaction of risks, in the form of a countercyclical, cross-sectional variance of idiosyncratic income risk. The other is due to the convexity of the welfare gain in total risk.

Based on a calibrated, large-scale overlapping generations model, we find that introducing a PAYG financed social security system with a contribution rate of 2% leads to strong long-run welfare gains of over 2% in terms of a consumption equivalent variation despite significant crowding out of capital. Our finding of strong welfare gains is very robust across different calibrations. Such positive net welfare gains are contrary to standard findings in the related literature. We document that considering both risks jointly is crucial for this finding. By examining only one risk in isolation, the previous literature missed two amplifying mechanisms that turn out to have quantitatively important welfare implications. In fact, when considering only one type of risk—either aggregate business cycle or idiosyncratic productivity risk—in isolation, we document

fluctuations caused by depreciation shocks are counteracted by households' savings, so that, for higher IES, volatility of capital is smaller, leading to smaller volatility of aggregate output.

[34] With a risk aversion of 2 (an IES of 0.5), cf. Footnote 32, the share is approximately 50%.

welfare losses, which is in line with the previous literature.

There is an interesting parallel to the literature on the welfare costs of fluctuations. In his seminal contribution, Lucas (1987) demonstrated that the costs of business cycles are negligible. However, when business cycle risk is analyzed in conjunction with idiosyncratic income risk, then welfare costs can become very large, see De Santis (2007) and Krebs (2007).

While our analysis uncovers important biases in the welfare assessment and documents that they matter quantitatively, some aspects are not taken into account. We abstract from endogenous labor supply. This may bias results in favor of social security for two reasons. First, we do not account for self-insurance against risk through endogenous labor supply adjustments. Second, a higher contribution rate would distort labor supply decisions and thereby crowd out aggregate labor supply. However, when taking labor market frictions into account and considering small policy changes, as in this paper, a calibrated model would only lead to second-order effects of endogenous labor reactions.

Another important extension would be to include survival risk. When annuity markets are missing, social security can be beneficial because it partially substitutes for these markets.[35] We abstract from this risk to focus the analysis on earnings risks. Also, it is not straightforward to jointly model survival risk and financial risk with Epstein-Zin preferences, see Cordoba and Ripoll (2013).

Furthermore, in our economy, intergenerational sharing of aggregate risk is limited to generations alive at the same point in time. From a social planner's perspective, it would be desirable to also share this risk with future, unborn generations. This could be achieved by allowing the government to take on debt or to manage a pension fund in order to smooth shocks over time. That would open up an additional insurance channel, which would increase the welfare gains of introducing social security.

Finally, our analysis restricts attention to redistribution within the social security system, taking redistribution through taxes and transfers during the working period as given. Yet, Huggett and Parra (2010) argue that it is important to simultaneously analyze the optimal design of both systems. This is an interesting aspect that we plan to address, in addition to the other questions discussed above, in our future research.

[35] Caliendo, Guo, and Hosseini (2014) demonstrate that this does not hold in a stylized general equilibrium with accidental bequests.

References

Aaron, H. J. (1966). The Social Insurance Paradox. *Canadian Journal of Economics and Political Science 32*, 371–374.

Auerbach, A. J. and L. J. Kotlikoff (1987). *Dynamic Fiscal Policy*. Cambridge, MA: Cambridge University Press.

Ball, L. and N. G. Mankiw (2007). Intergenerational Risk Sharing in the Spirit of Arrow, Debreu, and Rawls, with Applications to Social Security Design. *Journal of Political Economy 115*(4), 523–547.

Bansal, R. and A. Yaron (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance 59*(4), 1481–1509.

Bohn, H. (2001). Social Security and Demographic Uncertainty: The Risk Sharing Properties of Alternative Policies. In J. Campbell and M. Feldstein (Eds.), *Risk Aspects of Investment Based Social Security Reform*, Number 7030, pp. 203–241. Chicago: University of Chicago Press.

Bohn, H. (2009). Intergenerational Risk Sharing and Fiscal Policy. *Journal of Monetary Economics 56*, 805–816.

Boldrin, M., L. J. Christiano, and J. D. M. Fisher (1995). Asset Pricing Lessons for Modeling Business Cycles. *NBER Working Paper Series* (5262), 1–52.

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

Carroll, C. D. (2006). The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems. *Economics Letters 91*(3), 312–320.

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003). Accounting for the U.S. Earnings and Wealth Inequality. *Journal of Political Economy 111*(4), 818–857.

Conesa, J. C. and D. Krueger (1999). Social Security Reform with Heterogeneous Agents. *Review of Economic Dynamics 2*(4), 757–795.

Constantinides, G. G. M. and D. Duffie (1996). Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy 104*(2), 219–240.

Cordoba, J. and M. Ripoll (2013). Beyond Expected Utility in the Economics of Health and Longevity, *Working Paper*.

Croce, M. M. (2014). Long-run productivity risk: A new hope for production-based asset pricing? *Journal of Monetary Economics 66*, 13–31.

Davila, J., J. H. Hong, P. Krusell, and J.-V. Ríos-Rull (2012). Constrained Efficiency in the Neoclassical Growth Model With Uninsurable Idiosyncratic Shocks. *Econometrica 80*(6), 2431–2467.

De Nardi, M. (2004). Wealth Inequality and Intergenerational Links. *Review of Economic Studies 71*, 743–768.

De Santis, M. (2007). Individual Consumption Risk and the Welfare Cost of Business Cycles. *American Economic Review 97*(4), 1488–1506.

Diamond, P. A. (1977). A Framework for Social Security Analysis. *Journal of Public Economics 8*(3), 275–298.

Epstein, L. G. and S. E. Zin (1989). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica 57*(4), 937–969.

Epstein, L. G. and S. E. Zin (1991). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. *Journal of Political Economy 99*(2), 263–286.

Fehr, H. and C. Habermann (2008). Risk Sharing and Efficiency Implications of Progressive Pension Arrangements. *Scandinavian Journal of Economics 110*(2), 419–443.

Fernández-Villaverde, J. and D. Krueger (2011). Consumption and Saving Over the Life Cycle: How Important Are Consumer Durables? *Macroeconomic Dynamics 15*(05), 725–770.

Franke, G., H. Schlesinger, and R. C. Stapleton (2011). Risk taking with additive and multiplicative background risks. *Journal of Economic Theory 146*(4), 1547–1568.

Gollier, C. and J. W. Pratt (1996). Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica 64*, 1109–1123.

Gomes, F., A. Michaelides, and V. Polkovnichenko (2012). Fiscal Policy and Asset Prices with Incomplete Markets. *Review of Financial Studies 26*(2), 531–566.

Gomes, F. J. and A. Michaelides (2008). Asset Pricing with Limited Risk Sharing and Heterogeneous Agents. *The Review of Financial Studies 21*(1), 415–448.

Goodman, L. A. (1960). On the Exact Variance of Products. *Journal of the American Statistical Association 55*(292), 708–713.

Gordon, R. H. and H. R. Varian (1988). Intergenerational Risk Sharing. *Journal of Public Economics 37(2)*, 185–202.

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Harenberg, D. and A. Ludwig (2015). Social Security in an Analytically Tractable Overlapping Generations Model with Aggregate and Idiosyncratic Risk. *SAFE Working Paper 71*, 1–32.

Hasanhodzic, J. and L. Kotlikoff (2013). Generational Risk–Is It a Big Deal?: Simulating an 80-Period OLG Model with Aggregate Shocks. *NBER Working Paper 19179*, 1–47.

Hintermaier, T. and W. Koeniger (2010). The Method of Endogenous Gridpoints with Occasionally Binding Constraints Among Endogenous Variables. *Journal of Economic Dynamics and Control 34*(10), 2074–2088.

Hintermaier, T. and W. Koeniger (2011). On the Evolution of the US Consumer Wealth Distribution. *Review of Economic Dynamics 14*(2), 317–338.

Huggett, M. and J. C. Parra (2010). How Well Does the U.S. Social Insurance System Provide Social Insurance? *Journal of Political Economy 118*(1), 76–112.

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

Huggett, M., G. Ventura, and A. Yaron (2011). Sources of Lifetime Inequality. *American Economic Review 101*(7), 2923–2954.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (1995). A Life Cycle Analysis of Social Security. *Economic Theory 6*(1), 83–114.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (1998). The Effect of Tax-Favored Retirement Accounts on Capital Accumulation. *American Economic Review 88*(4), 749–768.

Kopecky, K. A. and R. M. Suen (2010). Finite State Markov-Chain Approximations to Highly Persistent Processes. *Review of Economic Dynamics 13*(3), 701–714.

Krebs, T. (2007). Job Displacement Risk and the Cost of Business Cycles. *American Economic Review 97*(3), 664–686.

Kreps, D. M. and E. L. Porteus (1978). Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica 46*(1), 185–200.

Krueger, D. and F. Kubler (2006). Pareto-Improving Social Security Reform when Financial Markets are Incomplete!? *American Economic Review 96*(3), 737–755.

Krusell, P. and J. A. A. Smith (1997). Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns. *Macroeconomic Dynamics 1*(02), 387–422.

Krusell, P. and J. A. A. Smith (1998, October). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Lucas, R. E. (1978). Asset Prices in an Exchange Economy. *Econometrica 46(6)*, 1426–1445.

Lucas, R. E. (1987). *Models of Business Cycles*. Oxford: Blackwell Publishing Ltd.

Ludwig, A. and M. Reiter (2010). Sharing Demographic Risk - Who is Afraid of the Baby Bust? *American Economic Journal: Economic Policy 2(4)*, 83–118.

Ludwig, A., T. Schelkle, and E. Vogel (2012). Demographic Change, Human Capital and Welfare. *Review of Economic Dynamics 15*(1), 94–107.

Mankiw, N. G. (1986). The Equity Premium and the Concentration of Aggregate Shocks. *Journal of Financial Economics 17*, 211–219.

Matsen, E. and O. Thogersen (2004). Designing Social Security - a Portfolio Choice Approach. *European Economic Review 48*, 883–904.

Merton, R. (1983). On the Role of Social Security as a Means for Efficient Risk Sharing in an Economy Where Human Capital is not Tradable. In Z. Bodie and J. B. Shoven (Eds.), *Financial Aspects of the United States Pension System*, pp. 325–358. Chicago: University of Chicago Press.

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

Olovsson, C. (2010). Quantifying the Risk-Sharing Welfare Gains of Social Security. *Journal of Monetary Economics 57*, 364–375.

Rajan, R. and L. Zingales (1995). What Do We Know About Capital Structure? Some Evidence From International Data. *The Journal of Finance 50*(5), 1421–1460.

Ríos-Rull, J.-V. (1999). Computation of Equilibria in Heterogenous Agent Models. In R. Marimón and A. Scott (Eds.), *Computational Models for the Study of Dynamic Economies*, pp. 238–264. Oxford, N.Y.: Oxford University Press.

Shiller, R. J. (1999). Social Security and Institutions for Intergenerational, Intragenerational, and International Risk-Sharing. *Carnegie- Rochester Conference Series on Public Policy 50(99)*, 165–204.

Siegel, J. J. (2002). *Stocks for the Long Run : The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies*. New York: McGraw-Hill.

Storesletten, K., C. I. Telmer, and A. Yaron (1999). The Risk-Sharing Implications of Alternative Social Security Arrangements. *Carnegie-Rochester Conference Series on Public Policy 50*, 213–259.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy 112*(3), 695–717.

Storesletten, K., C. I. Telmer, and A. Yaron (2007). Asset Pricing with Idiosyncratic Risk and Overlapping Generations. *Review of Economic Dynamics 10*(4), 519–548.

Weil, P. (1989). The Equity Premium Puzzle and the Risk-Free Rate Puzzle. *Journal of Monetary Economics 24*, 401–421.

A Appendix: Sketch of Proofs

We here sketch the proofs of the propositions. All analytical details are contained in Supplementary Appendix B.

Proof of Proposition 1. We maximize

$$Eu(c_{i,2,t+1}) = \frac{1}{1-\theta} E\Big(\bar{w}_t\Big(\bar{R}\eta_{i,1,t}\zeta_t\varrho_{t+1} + \tau\left((1+\lambda)\zeta_{t+1} - \bar{R}\eta_{i,1,t}\zeta_t\varrho_{t+1}\right)\Big)\Big)^{1-\theta}.$$

Evaluating the FOC at $\tau = 0$ we get under Assumption 1 the condition:

$$\frac{1+\lambda}{\bar{R}}\frac{E[Z_1]}{E[Z_2]} = \frac{1+\lambda}{\bar{R}} \exp\left(\theta\left(\sigma^2_{\ln\eta} + \sigma^2_{\ln AR}\right)\right) > 0 \tag{13}$$

where $\sigma_{\ln AR} \equiv \sqrt{\sigma^2_{\ln\zeta} + \sigma^2_{\ln\varrho}}$. To evaluate the CEV between two scenarios, i.e., comparing $Eu\left(c^{\tau>0}_{i,2,t+1}\right)$ with $Eu\left(c^{\tau=0}_{i,2,t+1}\right)$, we use that

$$Eu\left(c^{\tau>0}_{i,2,t+1}\right) = Eu\left(c^{\tau=0}_{i,2,t+1}\right) + \frac{\partial Eu\left(c^{\tau=0}_{i,2,t+1}\right)}{\partial\tau} d\tau.$$

and evaluate this expression at $\tau = 0$, using expression (13). This gives (4). □

Proof of Equation (7). We use the discrete state analogue to (13) in (17) to get equation (7a). Equation (7b) then follows from applying Definition 2 to the above. □