

Social Security and the Interactions Between Aggregate and Idiosyncratic Risk*

Daniel Harenberg[†] Alexander Ludwig[‡]

This Version: March 20, 2014

Abstract

We ask whether a PAYG-financed social security system is welfare improving in an economy with idiosyncratic and aggregate risk. We argue that interactions between the two risks are important for this question. One is a direct interaction in the form of a countercyclical variance of idiosyncratic income risk. The other indirectly emerges over a household's life-cycle because retirement savings contain the history of idiosyncratic and aggregate shocks. We show that this leads to risk interactions, even when risks are statistically independent. In our quantitative analysis, we find that introducing social security with a contribution rate of two percent leads to welfare gains of 2.2% of lifetime consumption in expectation, despite substantial crowding out of capital. This welfare gain stands in contrast to the welfare losses documented in the previous literature, which studies one risk in isolation. We show that jointly modeling both risks is crucial: 60% of the welfare benefits from insurance result from the interactions of risks.

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

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

*We thank Klaus Adam, Alan Auerbach, Martin Barbie, Tino Berger, Antoine Bommier, Johannes Brumm, Georg Dürnecker, Ayse Imrohoroglu, Philip Jung, Tom Krebs, Dirk Krueger, Felix Kubler, Michele Tèrtilt, Fabrizio Zilibotti, and various seminar participants at the XVI Workshop on Dynamic Macroeconomics in Vigo, the University of Cologne, University of California at Berkeley, University of Mannheim, University of Frankfurt, University of Munich, University of St. Gallen, Toulouse School of Economics, ETH Zürich, University of Zürich and Stockholm School of Economics 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 of the Land North Rhine-Westfalia.

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

[‡]CMR & FiFo, University of Cologne; MEA; Netspar; ludwig@wiso.uni-koeln.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 that features aggregate wage and asset return risk as well as idiosyncratic productivity risk. We follow the literature and assume that insurance markets for both forms of risk are incomplete. In such a setting, social security can increase economic efficiency by partially substituting for missing markets. The analysis is embedded in a general equilibrium framework to account for the costs of crowding out. Our analysis differs from the previous literature in that, thus far, prior studies have only considered models with one type of risk. One strand of the literature has examined social security when only aggregate risk is present (e.g., the quantitative analysis of Krueger and Kubler (2006)). In that setting, social security can improve intergenerational risk sharing. The other strand has only considered idiosyncratic risk (e.g., İmrohoroğlu, İmrohoroğlu, and Joines (1995, 1998)). There, social security is valuable because of intragenerational insurance. However, households face both types of risk over their lifetimes. We emphasize that jointly modeling both risks is not only relevant for obtaining a complete picture of how much insurance social security can provide. More important, there are interactions between the two types of risk, which amplify the welfare benefits of social security. We show that the quantitative role of these interactions is large so that introducing social security increases welfare.

The first interaction is an interaction over the life-cycle, and we accordingly call it the life-cycle interaction (*LCI*). This interaction arises even when the two risks are statistically independent. To the best of our knowledge, this effect has not been discussed in the related literature. To understand it, consider a standard life-cycle model with idiosyncratic wage risk as well as aggregate wage and return risk. In the absence of social security, retirement consumption is financed by private savings. Private savings are accumulated over the life-cycle and contain the history of idiosyncratic and aggregate shocks. We show analytically that the variance of retirement consumption therefore contains an interaction term between the two types of risk, called *LCI*. As households face both risks for many years before they enter retirement, this interaction term can become large. Since social security reduces the variance of retirement consumption, this interaction of risks amplifies the welfare benefits from social security.

The second interaction is a countercyclical cross-sectional variance of idiosyncratic productivity risk (*CCV*). This means that the variance of idiosyncratic shocks is higher in a downturn than in a boom.¹ 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). We add this channel to understand how benefits from insurance provided by social security are affected when idiosyncratic risk varies over the cycle.²

We begin our analysis with a simplified two-generations model with incomplete markets, yielding analytical results. We expose *LCI* and demonstrate how it amplifies the welfare benefits of social security. We then build a large-scale overlapping generations (OLG) model in the tradition of Auerbach, Kotlikoff, and Skinner (1983) and Auerbach and Kotlikoff (1987), extended with various risks, to evaluate how much the two interactions, *LCI* and *CCV*, matter quantitatively. Aggregate wage risk is introduced through a standard technology shock. Aggregate return risk is driven by the same technology shock and an additional shock to the depreciation rate of physical capital, as in Gomes and Michaelides (2008), Krueger and Kubler (2006), and Storesletten, Telmer, and Yaron (2007), among others. The two shocks enable us to calibrate the model in such a way that it produces realistic fluctuations of wages and returns, both of which are central to the welfare implications of social security.

The social security system is a pure pay-as-you-go (PAYG) system. In every period, all contributions are paid out as a lump-sum to all retirees. Households can also save privately by investing in a risk-free bond and a risky stock. Including this portfolio choice in the quantitative model is important, as social security can be regarded as an asset with a low risk and a low return. Therefore, the risk-return structure of the bond and the stock directly affect the value of social security and an appropriate calibration is crucial. 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, *ceteris paribus*.

When we calibrate the quantitative model to the U.S. economy, we find that a marginal introduction of social security with a contribution rate of 2% leads to a strong welfare gain of approximately 2% in terms of consumption equivalent variation. This welfare improvement is obtained despite substantial welfare losses from crowding out of capital in

¹In contrast to *LCI*, *CCV* is a direct interaction between aggregate and idiosyncratic risk.

²Guvenen, Ozkan, and Song (2013) show that the skewness of idiosyncratic shocks, not their variance, is countercyclical. From an ex-ante perspective, both specifications capture a direct interaction between aggregate and idiosyncratic risk. We chose the *CCV* specification to remain in line with the previous asset pricing literature. Our expectation is that a countercyclical left-skewness would reinforce our results.

general equilibrium. These findings stand in contrast to the previous literature. Crucially, we find that the interactions between the two risks are very important: approximately 60% of the welfare benefits from insurance can be attributed to the interactions, *LCI* and *CCV*. These numbers are robust for a range of common calibration strategies. Furthermore, we also consider economies with only one risk. Consistent with the previous literature, we observe welfare losses in these cases. This reemphasizes the importance of jointly modeling both risks.

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 and thereby increase 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 if labor income and capital returns are imperfectly correlated.

Gordon and Varian (1988), Matsen and Thøgersen (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.

There are not many quantitative papers with aggregate risk and social security. Krueger and Kubler (2006) is the most similar to our work.³ They also examine a marginal introduction of a 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 analysis differs substantially because we also include idiosyncratic risk and analyze interactions between risks.

In contrast, there are many quantitative papers that 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).⁵ Generally speaking, one conclusion from this literature is that welfare in a stationary economy

³Ludwig 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, which in that context is welfare improving.

⁴The recent work by Hasanhodzic and Kotlikoff (2013), who demonstrate that generational risk is small in a calibrated overlapping generations economy with aggregate risk, mirrors these findings.

⁵This literature also includes idiosyncratic survival risk, which we abstract from for two reasons. First, it would distract from the main point of the paper, which focuses on the interactions. Second, it is not straightforward to model both financial risk and survival risk when households have Epstein-Zin preferences (Epstein and Zin 1989; Epstein and Zin 1991) as in our quantitative model, see, e.g., Cordoba and Ripoll (2013).

without social security is higher than in one with a PAYG system. That is, the losses from crowding out dominate the gains from completing insurance markets. More recent work by Nishiyama and Smetters (2007) and Fehr and Habermann (2008) are examples of papers that focus on modeling the institutional features of existing social security systems in detail. We abstract from such details. Our results demonstrate the benefits of a flat minimum pension.

Huggett and Parra (2010) argue that it is important to consider simultaneous reforms to both the social security system and the general tax system. They report strong welfare gains from joint reforms of both systems. The question of how much redistribution to provide through the general tax system and how much life-cycle risk social security should insure is an interesting one. Redistribution through social security is sensible because total lifetime resources are observed at retirement and moral hazard frictions limit optimal redistribution via other social insurance institutions such as unemployment insurance.⁶ Here, however, we do not address the optimal joint design of both systems. Instead, we follow the more standard approach and take the general income tax system as given. Consequently, we calibrate our model to an income process after taxation.

Finally, Gomes, Michaelides, and Polkovichenko (2012) use a very similar model to study how changes in fiscal policy and government debt affect asset prices and capital accumulation. Our work also relates to the literature on the welfare costs of aggregate fluctuations initiated by Lucas (1978). There are two papers that highlight how interactions between idiosyncratic and aggregate risk can increase these costs substantially: Santis and De Santis (2007) in a partial equilibrium, and Krebs (2007) in a general equilibrium model without trade. In contrast to those papers, the present model features trade in general equilibrium.

The remainder of this paper is structured as follows. We derive our analytical results in Section 2. Section 3 describes the quantitative model and Section 4 presents the calibration. The main results of our quantitative analysis are presented in Section 5. Both the choice of calibration targets in Section 4 and the interpretation of the results in Section 5 heavily draw on the insights derived in the analytical section. We conclude in Section 6. Proofs, computational details, and details on the calibration are relegated to separate appendices.

⁶Werning (2007)'s result on the optimality of inter-temporal tax smoothing is suggestive for optimality of taxation based on lifetime earnings. An early proponent of such a scheme is Vickrey (1947), although for horizontal equity rather than efficiency considerations. More recently, Shourideh and Troshkin (2012) and Golosov, Shourideh, Troshkin, and Tsyvinski (2013), who extend the work by Huggett and Parra (2010), also optimize over the parameters of a redistributive pension system.

2 A Two-Generations Model

The analytical model developed here provides useful insights for our quantitative analysis. 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 only consume in the second period of life. This literature only considers aggregate risk. We extend it by adding idiosyncratic risk, which enables us to demonstrate that there is an important interaction mechanism.

2.1 Model

In each period t , a continuum of households is born. Households live for two periods only. A household has preferences over consumption in the second period.⁷ In the first period of life, the household experiences an idiosyncratic productivity shock, which we denote 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 $E_t [u(c_{i,2,t+1})]$. We assume a CCRA per period utility function with coefficient of relative risk aversion θ ,

$$u(c_{i,2,t+1}) = \begin{cases} \frac{c_{i,2,t+1}^{1-\theta}}{1-\theta} & \text{for } \theta \neq 1 \\ \ln(c_{i,2,t+1}) & \text{for } \theta = 1. \end{cases}$$

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 , and hence net wage income is given by $(1 - \tau)\eta_{i,1,t}w_t$. 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 in the two periods of life are given by

$$s_{i,2,t+1} = (1 - \tau)\eta_{i,1,t}w_t \tag{1a}$$

$$c_{i,2,t+1} = s_{i,2,t+1}R_{t+1} + y_{t+1}^{ss} \tag{1b}$$

⁷In our quantitative model, we consider many periods with consumption and saving decisions occurring in each period.

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. The amount paid into the social security system depends on the realization of η , and hence more fortunate households have to pay a larger amount. However, everyone receives the same lump-sum pension payments in retirement, y_{t+1}^{ss} . Thus, there is substantial redistribution, which constitutes an intragenerational sharing of idiosyncratic risk.

Aggregate wages and interest rates are stochastic. In this section, we limit the analysis to a partial equilibrium, and hence wages and returns are exogenous.⁸ 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:

$$w_t = \bar{w}_t \zeta_t = \bar{w}_{t-1} (1 + \lambda) \zeta_t \quad (2a)$$

$$R_t = \bar{R} \varrho_t, \quad (2b)$$

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

Social security is a pure PAYG system with lump-sum pension benefits. It is operated by the government, which must obey the budget constraint

$$\tau w_t N_{1,t} = y_t^{ss} N_{2,t},$$

where $N_{j,t}$ is population in period t of age j , i.e., $N_{j,t} = \int N_{i,j,t} di$. We here abstract from population growth,⁹ hence

$$\tau w_t = y_t^{ss}. \quad (3)$$

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 less than perfectly correlated, then y_t^{ss} and R_t are less than perfectly correlated. Thus, social security pensions provide a hedge against volatile asset income at retirement. This constitutes an intergenerational sharing of aggregate risk.

⁸In the quantitative analysis in the next section, we will consider a general equilibrium.

⁹Our quantitative analysis features population growth, which is required to correctly calibrate the implicit return of social security.

2.2 Analysis

To derive our results, we make the following assumptions.

Assumption 1. a) All shocks are bounded below by zero: $\zeta, \varrho, \eta > 0$.

b) All shocks have a mean of one: $E\zeta = E\varrho = E\eta = 1$.

c) All shocks are uncorrelated over time.

d) All shocks are statistically independent from each other.

Assumption 1a is a common restriction on economic shock processes, and Assumption 1b an innocuous normalization. Despite the observed positive serial correlation of wages and asset returns in annual data, Assumption 1c can be justified on the grounds of the long factual periodicity of each period in a two-period OLG model, which is approximately 30 to 40 years. Assumption 1d is crucial. It greatly simplifies our analytical solutions. More important, it is at the heart of one of the contributions of this paper: although the shocks are independent, we will see below that they indirectly interact and thereby drive up the value of social security.

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, $E[u(c_{i,2,t+1})]$. We can now state our first result.

Proposition 1. *Under Assumption 1, a marginal introduction of social security increases ex-ante expected utility iff*

$$(1 + \lambda) \frac{E\left[\frac{1}{\varrho_{t+1}^\theta}\right] E\left[\frac{1}{\zeta_t^\theta}\right] E\left[\frac{1}{\eta_{i,1,t}^\theta}\right]}{E\left[\frac{1}{\varrho_{t+1}^{\theta-1}}\right] E\left[\frac{1}{\zeta_t^{\theta-1}}\right] E\left[\frac{1}{\eta_{i,1,t}^{\theta-1}}\right]} > \bar{R}. \quad (4)$$

Proof. Relegated to Appendix A. □

The right-hand side of equation (4) reflects the costs of introducing social security, which here is simply the expected return on savings, \bar{R} . The left-hand side reflects the benefit of introducing social security, which is a risk-adjusted implicit return. Obviously, the implicit return increases if λ increases. If we set all shocks to their respective means of unity, we obtain the well-known Aaron (1966) condition, namely that $(1 + \lambda) > \bar{R}$. That is, if there were no risk, then social security would only be welfare enhancing if its return exceeded the return on private savings.

The proposition demonstrates that with risk, the introduction of social security may also be welfare enhancing for $(1 + \lambda) < \bar{R}$. To see this, observe that for $\theta \geq 1$, both the numerator and the denominator in the term representing the risk adjustment represent convex functions of the respective random variables. Convexity is stronger in the numerator, so that a mean-preserving spread, an increase in the variance of the respective random variables, increases the numerator more substantially. This is easiest to see for $\theta = 1$, in which case the denominator simply equals 1. To exhibit this more clearly, we simplify the expression for the risk adjustment by making the following conventional assumption regarding the distribution of shocks:

Assumption 2. $\eta_{i,1,t}, \zeta_t, \varrho_t$ are distributed jointly log-normal with means $\mu_{\ln \eta}, \mu_{\ln \zeta}, \mu_{\ln \varrho}$ and variances $\sigma_{\ln(\eta)}^2, \sigma_{\ln(\zeta)}^2, \sigma_{\ln(\varrho)}^2$.

With this additional assumption, condition (4) becomes intuitive and easy to interpret.

Proposition 2. *Under Assumptions 1 and 2, a marginal introduction of social security increases ex-ante expected utility iff*

$$(1 + \lambda) \cdot (1 + TR)^\theta > \bar{R}, \quad (5)$$

where

$$TR \equiv \text{var}(\eta_{i,1,t} \zeta_t \varrho_{t+1}) = \underbrace{\sigma_\eta^2}_{IR} + \underbrace{\sigma_\zeta^2 + \sigma_\varrho^2 + \sigma_\zeta^2 \sigma_\varrho^2}_{AR} + \underbrace{\sigma_\eta^2 (\sigma_\zeta^2 + \sigma_\varrho^2 + \sigma_\zeta^2 \sigma_\varrho^2)}_{LCI=IR \cdot AR}. \quad (6)$$

Proof. Relegated to Appendix A. □

The risk adjustment on the left-hand side of (5) is now $(1 + TR)^\theta > 0$. It can be interpreted as the value of insurance that social security provides. This value of insurance drives up the implicit return of social security, $1 + \lambda$, thereby modifying the standard Aaron condition. Naturally, the value of insurance is increasing in the coefficient of risk aversion, θ . It is also increasing in TR , which stands for 'total risk', as it contains the variances of all shocks, cf. equation (6). Idiosyncratic risk, abbreviated IR , is simply the variance of idiosyncratic shocks, σ_η^2 . Likewise, aggregate risk, AR , contains the variances of aggregate shocks. The important term is the third, in which IR is multiplied by AR , which constitutes the life-cycle interaction, LCI .

To understand the source of this term, notice that retirement consumption in the absence of social security is given by $\bar{w}_t \bar{R} \eta_{i,1,t} \zeta_t \varrho_{t+1}$. This shows that accumulated savings

for retirement consumption inherit the history of aggregate risks—as captured by the component $\zeta_t \varrho_{t+1}$ —and idiosyncratic risk—component $\eta_{i,1,t}$. By providing a lump-sum transfer, social security jointly reduces exposure to these risks.¹⁰ Evaluating the variance of the product $\eta_{i,1,t} \zeta_t \varrho_{t+1}$ gives rise to the product of variances of idiosyncratic and aggregate risks, as captured by LCI in equation (6).¹¹

From an ex-ante perspective, the variance of assets at retirement may well be large, because in the context of a life-cycle economy, we are dealing with long horizons. Therefore, the interaction, LCI , may be large. We pursue this question in the quantitative part of the paper.

We next address how a marginal introduction of social security translates into utility in terms of consumption equivalent variation (CEV). 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 ($\tau = d\tau > 0$). If $g_c > 0$, then the household prefers the economy with social security. To remain consistent with the subsequent quantitative analysis, we include a superscript PE to indicate partial equilibrium, as we subsequently oppose that with the CEV in a general equilibrium.

Proposition 3. *A logarithmic approximation of the consumption equivalent variation is*

$$g_c^{PE} = \left(\frac{1 + \lambda}{\bar{R}} (1 + TR)^\theta - 1 \right) d\tau. \quad (7)$$

A first-order Taylor series expansion of (7) around $TR = 0$ yields

$$g_c^{PE} \approx \left(\frac{1 + \lambda}{\bar{R}} - 1 + \theta \frac{1 + \lambda}{\bar{R}} AR + \theta \frac{1 + \lambda}{\bar{R}} IR + \theta \frac{1 + \lambda}{\bar{R}} LCI \right) d\tau \quad (8)$$

Proof. Relegated to Appendix A. □

In equation (7) we see that g_c^{PE} is approximately linear in $d\tau$, and hence under condition (2), the welfare gains are linearly increasing in the contribution rate. Moreover, g_c^{PE} increases exponentially in risks.

¹⁰At the same time, as social security is financed by period $t + 1$ wage income, it increases exposure to period $t + 1$ aggregate wage risk, ζ_{t+1} . However, in our marginal consideration, ζ_{t+1} appears linearly in the first-order condition, so that it has no effect because of independence as postulated in Assumption 1.

¹¹Although our proof of Proposition 2 does not make explicit use of the product formula of variances, term TR can be derived by applying that rule to the variance of term $\eta_{i,1,t} \zeta_t \varrho_{t+1}$, cf. Goodman (1960).

Equation (8) helps to understand what happens, if we shut down a risk. If the economy has no risk, i.e., if $IR = AR = 0$, then $g_c^{PE} \approx \left(\frac{1+\lambda}{\bar{R}} - 1\right) d\tau$. In a deterministic economy in which $1 + \lambda < \bar{R}$, the welfare losses from introducing social security increase approximately linearly in the contribution rate. If instead the economy only has aggregate risk, i.e., $IR = 0$, then $g_c^{PE} \approx \left(\frac{1+\lambda}{\bar{R}} - 1 + \theta \frac{1+\lambda}{\bar{R}} AR\right) d\tau$. That is, the welfare losses in the deterministic economy are reduced by the welfare gain attributable to the insurance against aggregate risk. We can therefore think of the total g_c^{PE} as consisting of the effects attributable to each risk minus the welfare loss in the deterministic economy with $1 + \lambda < \bar{R}$. Equation (8) motivates us to consider a linear decomposition of the form

$$g_c^{PE}(IR, AR) \approx g_c^{PE}(0, 0) + dg_c(IR) + dg_c(AR) + dg_c(LCI), \quad (9)$$

where $g_c^{PE}(0, 0)$ denotes the welfare in an economy without risk and $dg_c(X)$ the welfare effect attributable to risk type $X \in \{IR, AR, LCI\}$. This is the decomposition approach we follow in the quantitative part. In addition, equations (7)-(8) will guide us when we calibrate the quantitative model. They clearly reveal which parameters are important for the welfare effects. Finally, it is noteworthy that in both equations, the interaction LCI enters with the same weight as the pure risks. From this analysis, it can be expected that the life-cycle interaction can substantially increase the welfare gains of social security.

2.3 Discussion and Extensions

Our simple setup only provides a partial characterization of the total welfare effects of social security. First, we have excluded the counter-cyclical, cross-sectional variance of idiosyncratic risk (CCV) that was mentioned in the introduction and that we will include in the following quantitative model. CCV constitutes an explicit interaction between idiosyncratic and aggregate risk, and it can be shown that it amplifies the welfare consequences of social security. In contrast, LCI is present despite the orthogonality of the risks, which makes LCI somewhat more surprising.

Second, idiosyncratic income risk in our model can be understood as an uninsurable, additive background risk (Gollier and Pratt 1996). Loosely speaking, this makes the agent behave as if he were more risk-averse. In Harenberg and Ludwig (2014), we present an extension to a multi-period model and show that, in such a more general framework, LCI appears as a multiplicative background risk (Franke, Schlesinger, and Stapleton 2006), which comes on top of the additive background risk. The intuition follows from the dynamic budget constraint in the working period, in which total income is given by $sR + w\eta$. In this

budget constraint, $w\eta$ is an additive background risk. But in a dynamic setting, savings s represent a multiplicative background risk, because they are endogenously determined from past labor income, $w\eta$. Thus they inherit the multiplicative, stochastic nature of past labor income. The previous theoretical literature on background risk fails to address this point, as it essentially considers at a static optimization problem.¹²

Finally, as we do not consider consumption in the first period of life, the simple model omits the effects of taxation on the reallocation of consumption and savings. We also ignore any feedback in general equilibrium. In Harenberg and Ludwig (2014), we incorporate both channels in a standard Diamond (1965) model with risk. Thus, in contrast to the simple model presented here, consumption and savings decisions occur in the first period and wages and returns are determined in general equilibrium. There, we conclude that interactions between risks reduce precautionary savings and simultaneously increase crowding out of capital. Our quantitative model developed in the next section accounts for all of these additional channels.

3 The Quantitative Model

Our quantitative model extends the two-generations model presented in the previous section along several dimensions. Nonetheless, the main mechanisms remain the same. All of the extensions are standard and well understood. First, we include the counter-cyclical, cross-sectional variance of the idiosyncratic risk, CCV , as a direct, explicit interaction between idiosyncratic and aggregate risk. Second, we consider a general equilibrium, which allows us to account for the costs of crowding out of capital. Third, the periodicity is one calendar year, and there are J overlapping generations instead of two. Consumption and savings decisions take place every period. Population grows at a constant rate, which acts as an additional implicit return to social security. Fourth, we add a one-period, risk-free bond as a second asset. As the propositions in the previous section show, the interest rate on private savings is an important determinant of the value of social security. The bond imposes discipline in that respect and allows us to match asset prices. As a side effect, the household has an additional asset to self-insure against idiosyncratic and aggregate risk. *Ceteris paribus*, this reduces the beneficial effects of social security. Fifth, labor income features, as a deterministic component, a hump-shaped productivity profile over

¹²This literature is chiefly concerned with the choice of exposure to aggregate risk when background risk is increased. Relative to this literature, a key difference here is that social security jointly reduces exposure to both components of risk.

the life-cycle. Idiosyncratic income shocks are allowed to be autocorrelated. Finally, we employ Epstein-Zin (Epstein and Zin 1989; Epstein and Zin 1991) preferences. In contrast to CRRA preferences, they can partially disentangle risk aversion and the elasticity of inter-temporal substitution. We extensively exploit this feature in our sensitivity analyses.

3.1 Demographics and Risk

Time is discrete and runs from $t = 0, \dots, \infty$. 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 $z^t = (z_0, z_1, \dots, z_t)$. The shocks z_t follow a Markov chain with finite support \mathcal{Z} and nonnegative transition matrix π_z . Thus $\pi_z(z_{t+1}|z_t)$ represents the probability of the shock next period given the current shock.

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. Without loss of generality, 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.¹³ Consequently, the size of the population at t is $N(z^t) = (1 + n)^t$. We assume that at z_0 , the population structure is consistent with this population growth.

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. Thus, households within each generation differ with respect to their history of idiosyncratic shocks, e^j . Like the aggregate shock, 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 $\pi_e(e^j)$. We assume that a Law of Large Numbers applies, so that $\pi_e(e^j)$ represents both the individual probability for e^j and the fraction of the population with that shock history. Likewise, $\pi_e(e_{j+1}|e_j)$ represents both the individual transition probability and the population counterpart. Finally, $\pi_e(e_j)$ denotes the unconditional probability of shock e_j .

¹³See Footnote 5 in the Introduction for the important reasons for not including survival risk.

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 .¹⁴ Preferences over consumption c are represented by a recursive utility function $u_j(c, \cdot)$ of the Epstein-Zin form (Kreps and Porteus 1978; Epstein and Zin 1989; Epstein and Zin 1991; Weil 1989)¹⁵:

$$\begin{aligned}
 u_j(c, e^j, z^t) &= \left(\left[c_j(e^j, z^t) \right]^{\frac{1-\theta}{\gamma}} \right. \\
 &\quad \left. + \beta \left(\sum_{z^{t+1}} \sum_{e_{j+1}} \pi_z(z^{t+1}|z^t) \pi_e(e_{j+1}|e_j) \left[u_{j+1}(c, e^{j+1}, z^{t+1}) \right]^{1-\theta} \right)^{\frac{1}{\gamma}} \right)^{\frac{\gamma}{1-\theta}}, \\
 u_J(c, e^J, z^t) &= c_J(e^J, z^t), \\
 c &> 0,
 \end{aligned} \tag{10}$$

where β is the discount factor and θ controls risk aversion. The parameter γ is defined as $\gamma = \frac{1-\theta}{1-\frac{\theta}{\psi}}$ 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. This is the only channel through which idiosyncratic risk enters the model. The dependence of $\eta(e_j, z_t)$ on the aggregate shock is necessary to model the *CCV*. To shut *CCV* down, we can simply make this dependence ineffective. Without loss of generality, we assume that $E(\eta(e_j, z_t)|z_t) = 1$. Labor income $y_j(e_j, z^t)$ is

$$y_j(e_j, z^t) = w(z^t) \epsilon_j \eta(e_j, z_t), \tag{11}$$

where $w(z^t)$ is the real aggregate wage in terms of the consumption good at z^t . Insurance markets against 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. They wish to save due to a life-cycle

¹⁴In the two-generations model, we used indexes to denote a household for the sake of exposition. As we now have more states and ages, a more elaborate, state-dependent notation is required. However, for the two-generations model, they are equivalent, e.g., $u(c_{i,2,t+1}) \equiv u(c, e^2, z^{t+1})$.

¹⁵In a slight abuse of notation, we use the lower case letter u to denote remaining lifetime utility in this recursive formulation, which was used in Section 2 to denote the per-period utility function.

motive and a precautionary motive, because markets for idiosyncratic and aggregate risks are assumed to be incomplete. Both assets constitute a claim on the firm's capital in the following period. They only differ in their returns: 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-ahead risk-free interest rate $r_b(z^t)$. The sequential budget constraint is standard:

$$c_j(e^j, z^t) + s_{j+1}(e^j, z^t) + b_{j+1}(e^j, z^t) = (1 + r_s(z^t))s_j(e^j, z^t) + (1 + r_b(z^{t-1}))b_j(e^j, z^t) + (1 - \tau)y_j(e_j, z^t)I(j) + y^{ss}(z^t)(1 - I(j)), \quad (12)$$

where τ is a fixed social security contribution rate, $y^{ss}(z^t)$ is pension income from social security, and $I(j)$ is an indicator function that takes the value 1 if $j < j_r$ and 0 otherwise. Households cannot die in debt:

$$s_{J+1}(e^J, z^t) + b_{J+1}(e^J, z^t) \geq 0. \quad (13)$$

Together with the requirement of positive consumption $c > 0 \forall e^j, z^t$, equation (13) implies a natural debt limit at each date-event. Loosely speaking, at each z^t , a household with e^j can borrow up to the present value of its worst future income stream. As households do not have a bequest motive and utility is strictly increasing in consumption, equation (13) implies zero bequests. Consequently, 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 the final good $Y(z^t)$ 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 $\zeta(z_t)$, which depends only on the current aggregate shock:

$$Y(z^t) = \zeta(z_t)K(z^t)^\alpha((1 + \lambda)^t L(z^t))^{1-\alpha}. \quad (14)$$

Assuming a stochastic depreciation rate $\delta(z_t)$ as in Storesletten, Telmer, and Yaron (2007), Gomes and Michaelides (2008), Krueger and Kubler (2006) and others, the capital stock

evolves according to:

$$K(z^t) = I(z^{t-1}) + K(z^{t-1})(1 - \delta(z_{t-1})). \quad (15)$$

The firm rents capital and labor from households by paying factor prices equal to their respective marginal productivities. Thus wages, $w(z^t)$, and the return on capital, $r(z^t)$, are given by

$$w(z^t) = (1 + \lambda)^t (1 - \alpha) \zeta(z_t) \left(\frac{K(z^t)}{(1 + \lambda)^t L(z^t)} \right)^\alpha, \quad (16a)$$

$$r(z^t) = \alpha \zeta(z_t) \left(\frac{(1 + \lambda)^t L(z^t)}{K(z^t)} \right)^{1-\alpha} - \delta(z_t). \quad (16b)$$

The capital stock, $K(z^t)$, is financed by issuing stocks and bonds in quantities S and B , respectively:

$$K(z^t) = S(z^t) + B(z^t) = S(z^t)(1 + \bar{\kappa}_f). \quad (17)$$

The debt-equity ratio, $\bar{\kappa}_f$, is exogenous and constant. Therefore, the firm only decides on aggregate capital and does not make a decision regarding the capital structure. Thus, the firm's problem is just like in the standard model, which is nested for $\bar{\kappa}_f = 0$. Modeling leverage in this manner is a common practice in the finance literature to mechanically increase the volatility of stock returns, see, for example, Boldrin, Christiano, and Fisher (1995) or Croce (2010). This enables us to keep the depreciation shocks small in the calibration. Note that all of our results go through for $\bar{\kappa}_f = 0$.¹⁶

As shown in the Supplementary Appendix D, the bond and stock returns are

$$r_b(z^{t-1}) = \frac{1}{\bar{\kappa}_f} E \left[r(z^t)(1 + \bar{\kappa}_f) - r_s(z^t) | z^{t-1} \right] \quad (18a)$$

$$r_s(z^t) = r(z^t)(1 + \bar{\kappa}_f) - \bar{\kappa}_f r_b(z^{t-1}). \quad (18b)$$

Equation (18b) demonstrates how the debt-equity-ratio, $\bar{\kappa}_f$, leverages the stock return, thereby increasing both its expected value and variance.

¹⁶In the baseline calibration, the depreciation shocks are so small that $\bar{\kappa}_f = 0$ delivers the same results. Only in the calibration that matches the volatility of stock returns, $\bar{\kappa}_f = 0$ implies a volatility of aggregate investment that is too high, and a higher $\bar{\kappa}_f$ improves on that. It is well known that, in the standard model, it is difficult to match the volatilities of both the financial and the real side of the economy.

3.4 Social Security

Social security works just like in the two-generations model of Section 2. The government organizes a pay-as-you-go system with a fixed contribution rate τ that is levied on labor income. Pension income $y^{ss}(z^t)$ adjusts to ensure that the social security budget is balanced at every date-event. As before, $y^{ss}(z^t)$ does not depend on the idiosyncratic history, which means that every household receives the same pension income. Denoting by $P(z^t)$ the number of pensioners, $P(z^t) = \sum_{j=j_r}^J N_j(z^t)$, the budget constraint now reads $\tau w(z^t)L(z^t) = y^{ss}(z^t)P(z^t)$. The case $\tau = 0.0$, i.e., an economy without a social security system, characterizes the initial situation.

3.5 Equilibrium

To define equilibrium, we remove population and technology growth and express all aggregate variables in terms of labor efficiency units. This is achieved by dividing all 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 variable is written in lower case, e.g., $k(z^t) = \frac{K(z^t)}{(1+\lambda)^t L(z^t)}$. Population statistics used in the equilibrium definition are the stationary economic dependency ratio, $p = \frac{P(z^t)}{L(z^t)} = \frac{\sum_{j=j_r}^J (1+n)^{J-j}}{\sum_{j=1}^{j_r-1} (1+n)^{J-j}}$, and the stationary labor-to-population ratio, $\ell = \frac{L(z^t)}{N(z^t)} = \frac{\sum_{j=1}^{j_r-1} (1+n)^{J-j} \epsilon_j}{\sum_{j=1}^J (1+n)^{J-j}}$. Individual variables are only detrended 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 $\tilde{u}_j(\cdot)$.

We proceed with a recursive definition of equilibrium. To this end, we define a state space that is sufficient for solving the households' maximization problem. Let $\mathcal{E} = \{e_1, e_2, \dots, e_{max}\}$ and $\mathcal{J} = \{1, 2, \dots, J\}$, and let \mathcal{M} be a sigma-algebra over $\{[\tilde{s}, \infty] \times [\tilde{b}, \infty] \times \mathcal{E} \times \mathcal{J}\}$, where \tilde{s} and \tilde{b} are implied by the detrended version of equation (13). The measure Φ is defined over \mathcal{M} , and the set of all such measures is denoted by \mathcal{Q} . We follow the applied literature and define the state space to consist of Φ , the current idiosyncratic state $(\tilde{s}, \tilde{b}, e)$, and the current aggregate shock z . As a recursive equilibrium does not depend on the date-event, we drop time index t and use a prime for next period's variables.

Definition 1. *For any initial $(z_0, \Phi_0) \in \mathcal{Z} \times \mathcal{Q}$, a recursive competitive equilibrium consists of a measure Φ , measurable functions for household choices $\{\tilde{c}_j(\tilde{s}, \tilde{b}, e; \Phi, z), \tilde{s}'_j(\tilde{s}, \tilde{b}, e; \Phi, z), \tilde{b}'_j(\tilde{s}, \tilde{b}, e; \Phi, z)\}$ and an associated value function $\tilde{u}_j(\tilde{s}, \tilde{b}, e; \Phi, z)$, firm choices $k(\Phi, z)$,*

social security settings $\{\tau, \tilde{y}^{ss}(\Phi, z)\}$, factor prices $\{\tilde{w}(\Phi, z), r(\Phi, z)\}$, asset returns $\{r_b(\Phi), r_s(\Phi, z)\}$, and a law of motion $H(\Phi, z)$ such that:

a) given functions for prices and returns and the law of motion, the households' policy functions $\{\tilde{c}_j(\tilde{s}, \tilde{b}, e; \Phi, z), \tilde{s}'_j(\tilde{s}, \tilde{b}, e; \Phi, z), \tilde{b}'_j(\tilde{s}, \tilde{b}, e; \Phi, z)\}$ solve

$$\begin{aligned} & \max_{\tilde{c} > 0, \tilde{s}', \tilde{b}'} \tilde{u}_j(\tilde{s}, \tilde{b}, e; \Phi, z) \\ & = \begin{cases} \left(\tilde{c}^{\frac{1-\theta}{\gamma}} + \tilde{\beta} \left(\sum_{z'} \sum_{e'} \pi_z(z'|z) \pi_e(e'|e) \tilde{u}_{j+1}^{1-\theta}(\tilde{s}', \tilde{b}', e'; H(\Phi, z), z') \right)^{\frac{1}{\gamma}} \right)^{\frac{\gamma}{1-\theta}} \\ \tilde{c} & \text{if } j = J \end{cases} \\ & \text{s. t.} \quad \tilde{c} + \tilde{s}' + \tilde{b}' = (1 + r_s(\Phi, z))\tilde{s} + (1 + r_b(\Phi))\tilde{b} \\ & \quad \quad \quad + (1 - \tau)\tilde{y}_j(e, \Phi, z)I(j) + \tilde{y}^{ss}(\Phi, z)(1 - I(j)), \\ & \quad \quad \tilde{y}_j(e, \Phi, z) = \tilde{w}(\Phi, z)\epsilon_j\eta(e, z), \\ & \quad \quad \tilde{s}' + \tilde{b}' \geq 0 \quad \quad \quad \text{if } j = J. \end{aligned} \tag{19}$$

where $\tilde{\beta} = \beta(1 + \lambda)^{\frac{1-\theta}{\gamma}}$.

b) functions for prices and for firm choices are related by

$$\begin{aligned} \tilde{w}(\Phi, z) &= (1 - \alpha)\zeta(z)k(\Phi, z)^\alpha \\ r(\Phi, z) &= \alpha\zeta(z)k(\Phi, z)^{\alpha-1} - \delta(z). \end{aligned}$$

c) functions for asset returns are given by

$$\begin{aligned} r_b(\Phi) &= \frac{1}{\bar{\kappa}_f} E[r(\Phi, z)(1 + \bar{\kappa}_f) - r_s(\Phi, z)] \\ r_s(\Phi, z) &= r(\Phi, z)(1 + \bar{\kappa}_f) - \bar{\kappa}_f r_b(\Phi) \end{aligned}$$

d) the pension system budget constraint holds, i.e.,

$$\tau\tilde{w}(\Phi, z) = \tilde{y}^{ss}(\Phi, z)p \tag{20}$$

where p is the economic dependency ratio defined above.

e) all markets clear:

$$\begin{aligned}
\zeta(z)k(\Phi, z)^\alpha + (1 - \delta(z))k(\Phi, z) &= \frac{1}{\ell} \sum_{j=1}^J \sum_e \int_{\tilde{b}} \int_{\tilde{s}} \tilde{c}_j(\tilde{s}, \tilde{b}, e; \Phi, z) \Phi(\tilde{s}, \tilde{b}, e, j) d\tilde{b} d\tilde{s} \\
&\quad + k(H(\Phi, z), z')(1 + \lambda)(1 + n) \\
k(H(\Phi, z), z')(1 + \lambda)(1 + n) &= \frac{1}{\ell} \sum_{j=1}^J \sum_e \int_{\tilde{b}} \int_{\tilde{s}} (\tilde{s}'_j(\tilde{s}, \tilde{b}, e; \Phi, z) \\
&\quad + \tilde{b}'_j(\tilde{s}, \tilde{b}, e; \Phi, z)) \Phi(\tilde{s}, \tilde{b}, e, j) d\tilde{b} d\tilde{s} \\
\frac{k(H(\Phi, z), z')(1 + \lambda)(1 + n)}{(1 + \bar{\kappa}_f)} &= \frac{1}{\ell} \sum_{j=1}^J \sum_e \int_{\tilde{b}} \int_{\tilde{s}} \tilde{s}'_j(\tilde{s}, \tilde{b}, e; \Phi, z) \Phi(\tilde{s}, \tilde{b}, e, j) d\tilde{b} d\tilde{s}.
\end{aligned}$$

By Walras' Law, the bond market also clears.

f) the law of motion H is generated by the policy functions and the Markov transition matrix π_e , so that

$$\Phi' = H(\Phi, z)$$

with the initialization at $j = 1$ of $\tilde{s} = \tilde{b} = 0$.

3.6 Computational Solution

Like Gomes and Michaelides (2008) and Storesletten, Telmer, and Yaron (2007), among others, we compute an equilibrium of our model by applying the Krusell and Smith (1998) method. 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 two continuous choices. We parallelize the solution and simulation, which enables us to achieve a high degree of accuracy. Details of the computational solution are provided in the Supplementary Appendix B.

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 weights lifetime utilities with their respective probabilities.

A household's welfare of being born into an economy with policy A can be written as

$$E \left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) \right], \quad (21)$$

where the expectation is taken over all date-events z^t and \tilde{c}^A denotes the optimal, detrended consumption attainable under this policy. 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 .

Observation 1. *The consumption equivalent variation, i.e., the percentage increase in consumption required in regime A to make the household as well off as under regime B, is given by*

$$g_c = \frac{E \left[\tilde{u}_1(\tilde{c}^B, e^1, z^t) \right]}{E \left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) \right]} - 1. \quad (22)$$

Proof. Relegated to the Supplementary Appendix D. □

A positive number then indicates the percentage of lifetime consumption a household would be willing to give up 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.

We can provide a closed-form solution for g_c for an economy without risk. Following the discussion in Section 2.2, we denote the consumption equivalent variation in an economy

without risk by $g_c(0, 0)$.

Observation 2. Denote by pvi_1^A (pvi_1^B) the present discounted value of lifetime income in policy A (B). The consumption equivalent variation in the risk-free economy is given by

$$g_c(0, 0) = \frac{\tilde{u}_1(\tilde{c}^B)}{\tilde{u}_1(\tilde{c}^A)} - 1 = \frac{\widetilde{pvi}_1^B}{\widetilde{pvi}_1^A} - 1,$$

i.e., it is not affected by preference parameters.

Proof. Relegated to the Supplementary Appendix D. □

We make use of this result in the decomposition analysis, which is described next.

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, which means that it is close to our analytical results 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. That is, we isolate the welfare consequences of purely changing capital without introducing social security. If choosing a lower capital stock leads to a welfare gain, then there was over-accumulation of capital, which implies dynamic inefficiency. Conversely, if the reduction in capital leads to welfare losses, we say that the economy is dynamically efficient. Our notion of dynamic inefficiency precisely corresponds to the over-accumulation of capital as defined by Davila, Hong, Krusell, and Ríos-Rull (2012).¹⁷ In a dynamically inefficient economy, introducing social security could be welfare

¹⁷In heterogeneous agent models, the golden rule of capital formation cannot be applied to discriminate between efficient and inefficient steady states. This is emphasized by Davila, Hong, Krusell, and Ríos-Rull (2012) in a model with idiosyncratic income shocks and by Kuhle (2012) in a model with heterogeneous types. In our economy, the same opposing effects as in Davila, Hong, Krusell, and Ríos-Rull (2012) are at work. On the one hand, idiosyncratic risk induces precautionary savings thereby creating a negative pecuniary externality because households do not internalize the effects on aggregate capital when making their savings decision. As a consequence, a social planner would choose a lower capital stock. On the

improving just because it cures or mitigates the over-accumulation of capital. We avoid this situation in our calibrated model variants.

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. Recall that households use \hat{H}_A to forecast prices and solve their optimization problem. We now 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 ($\tau = 2\%$). That is, households optimize under the old laws of motion, but under the new social security system. Consistently, the simulations are performed with the old shocks and prices and the new optimal policy functions. Again, this is like a small open economy, in which the domestic social security system has changed, but world prices remain unaffected. The welfare gains stemming from insurance are then simply given by

$$g_c^{PE} = \frac{E \left[\tilde{u}_1(\tilde{c}^B, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 2\% \right]}{E \left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 0\% \right]} - 1. \quad (23)$$

As long as there is sufficient risk to insure, this number will generally be positive, because prices do not change and there are thus no costs from the crowding out of capital. Analogously, the corresponding general equilibrium number is

$$g_c^{GE} = \frac{E \left[\tilde{u}_1(\tilde{c}^B, e^1, z^t) | \mathcal{P}_B, \hat{H}_B, \tau = 2\% \right]}{E \left[\tilde{u}_1(\tilde{c}^A, e^1, z^t) | \mathcal{P}_A, \hat{H}_A, \tau = 0\% \right]} - 1, \quad (24)$$

where the crucial difference is that in the new equilibrium with policy B ($\tau = 2\%$), the utility of a newborn is computed under the new general equilibrium prices and laws of

other hand, idiosyncratic wage shocks generate wealth inequality that propagates over time and age. A social planner may then choose a higher capital stock, because the implied lower return may reduce inequality. These mechanisms come on top of the standard result that a deterministic OLG model may give rise to over-accumulation of capital.

motion, \mathcal{P}_B, \hat{H}_B . The welfare costs of crowding out are then given by the difference

$$g_c^{CrowdingOut} = g_c^{GE} - g_c^{PE}. \quad (25)$$

In a dynamically efficient economy as defined above, $g_c^{CrowdingOut}$ will be negative by definition.

The final step is the decomposition of g_c^{PE} into insurance against aggregate risk, idiosyncratic risk, and the two interactions, LCI and CCV . Recall from Proposition 3 that a Taylor-series approximation of g_c is linear in risks and interactions. For convenience, we restate equation (9):

$$g_c^{PE}(IR, AR) \approx g_c^{PE}(0, 0) + dg_c(IR) + dg_c(AR) + dg_c(IR) + dg_c(LCI),$$

where $g_c^{PE}(0, 0)$ denotes the welfare in an economy without risk and $dg_c(X)$ the welfare effect attributable to risk type $X \in \{IR, AR, LCI\}$. We now extend this by CCV and write it as a system of equations:

$$g_c^{PE}(AR, IR, CCV) = g_c^{PE}(0, 0) + dg_c(AR) + dg_c(IR) + dg_c(LCI) + dg_c(CCV) \quad (26a)$$

$$g_c^{PE}(AR, IR) = g_c^{PE}(0, 0) + dg_c(AR) + dg_c(IR) + dg_c(LCI) \quad (26b)$$

$$g_c^{PE}(IR) = g_c^{PE}(0, 0) + dg_c(IR) \quad (26c)$$

$$g_c^{PE}(AR) = g_c^{PE}(0, 0) + dg_c(AR) \quad (26d)$$

The right-hand side of the first line shows all of the components we want to isolate. To do so, we begin in the last line, where we compute $g_c^{PE}(AR)$ and $g_c^{PE}(0, 0)$, as in equation (23), but for an economy with only aggregate risk and one without risk, respectively.¹⁸ With those numbers at hand, we can back out the welfare effect attributable to aggregate risk, $dg_c(AR)$. Likewise, we compute $g_c^{PE}(IR)$ for an economy featuring only idiosyncratic risk to back out $dg_c(IR)$.¹⁹ 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 LCI , $dg_c(LCI)$. In the same manner, we obtain $dg_c(CCV)$. 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

¹⁸To compute $g_c(0, 0)$ we employ Observation 2.

¹⁹Notice that in models without aggregate risk, there will no longer be two assets with different returns. Thus, there will be no portfolio choice.

the general equilibrium results, because that is where the original equilibrium sequences and laws of motion come from.

4 Calibration

The selection of the targets and parameters to be calibrated is informed by our theoretical insights, in particular by Proposition 2. It indicates that the coefficient of relative risk aversion, θ , the variances of the shocks, and the returns 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 ratio²⁰, $\varsigma = \frac{E[r_{s,t} - r_{b,t}]}{\sigma[r_{s,t} - r_{b,t}]}$, and then aggregate shocks to match the equity premium, $\mu = E[r_{s,t} - r_{b,t}]$, 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, which we now describe in detail. Additional information on our empirical approach to measure calibration targets is provided in the Supplementary Appendix C. The numerical implementation of the calibration procedure is described in the Supplementary Appendix B.

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 the 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 US population.

²⁰The Sharpe ratio is the market price of risk, which in our model is given by $\varsigma = \frac{E[r_s(z^t) - r_b(z^t)]}{\sigma[r_s(z^t) - r_b(z^t)]}$. We divide by the standard deviation of the excess return, instead of the standard deviation of the stock return, because our bond returns are not constant over time.

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

Table 1: Summary of the Baseline Calibration

Parameter	Value	Target (source)	Stage
<i>Demographics</i>			
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
<i>Households</i>			
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 C	Earnings profiles (PSID)	1 st
$CCV, \sigma_v^2(z)$	{0.0445, 0.0156}	Storesletten, et al. (2007)	1 st
Autocorrelation of log income, ρ	0.952	Storesletten, et al. (2007)	1 st
<i>Firms</i>			
Capital share, α	0.32	Wage share (NIPA)	1 st
Leverage ratio, $\bar{\kappa}_f$	0.66	Croce (2010)	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
<i>Aggregate Risk</i>			
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

4.2 Households

In our baseline calibration, we treat the coefficient of risk aversion as a first-stage parameter, i.e., it is not set to match a target. We set 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, respectively 7.4, lower than their respective empirical estimates of 0.33 and 5.60%, taken from Robert Shiller's website.²² Below, we describe alternative calibrations in which risk aversion is calibrated as a second-stage parameter to either match the Sharpe ratio or the equity premium.

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 react more strongly to price changes, with the consequence that the welfare losses from crowding out are lower. We investigate how our results depend on this choice in our sensitivity analysis.

In our baseline calibration, we determine the discount factor β to match the capital output ratio of 2.65, which we calculate from NIPA data and which is in line with the estimates of, e.g., Fernández-Villaverde and Krueger (2011). Thus β is a second-stage parameter, and 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 the Supplementary Appendix C. The stochastic component of labor income, $\eta(e_j, z_t)$, is based on the estimates of labor earnings risk in Storesletten, Telmer, and Yaron (2004), who also provide an estimate of the countercyclical cross-sectional variance of income risk, *CCV*. Specifically, we employ the same estimates that they use in their life-cycle paper (Storesletten, Telmer, and Yaron 2007). These comprise an autocorrelation coefficient of log income of $\rho = 0.952$ and a conditional variance of innovations, $\sigma_v^2(z_t)$, 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 as proposed by Kopecky and Suen (2010).

4.3 Firms

We set the value of the capital share parameter, a first-stage parameter, to $\alpha = 0.32$, a standard choice. This is directly estimated from NIPA data on total compensation as a fraction of (appropriately adjusted) GDP. Our estimate of the deterministic trend growth rate, also a first-stage parameter, is based on data on total factor productivity (TFP).

²²See <http://aida.wss.yale.edu/~shiller/data.htm>.

The point estimate is $\lambda = 0.018$, which is in line with other studies. Leverage in the firm sector is set to $\bar{\kappa}_f = 0.66$, based on the empirical estimates in Rajan and Zingales (1995) and the calibration in Croce (2010).

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%, again taken from Robert Shiller's website. In economies without aggregate risk, cf. our discussion on the decomposition of welfare effects in Subsection 3.8, we calibrate δ_0 to produce a risk-free return of 4.2%, which corresponds to the empirical estimate of Siegel (2002).

4.4 Aggregate Risk

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

$$\zeta(z) = \begin{cases} 1 - \bar{\zeta} & \text{for } z \in z_1, z_2 \\ 1 + \bar{\zeta} & \text{for } z \in z_3, z_4 \end{cases} \quad \text{and } \delta(z) = \begin{cases} \delta_0 + \bar{\delta} & \text{for } z \in z_1, z_3 \\ \delta_0 - \bar{\delta} & \text{for } z \in z_2, z_4. \end{cases} \quad (27)$$

Set up in this way, z_1 corresponds to a low wage and a low return, while z_4 corresponds to a high wage and a high return.

To calibrate the entries in the transition matrix, denote the transition probability of remaining in the low technology state by $\pi^\zeta = \pi(\zeta' = 1 - \bar{\zeta} \mid \zeta = 1 - \bar{\zeta})$. Assuming that the transition of technology shocks is symmetric, we then have $\pi(\zeta' = 1 + \bar{\zeta} \mid \zeta = 1 + \bar{\zeta}) = \pi^\zeta$ and $1 - \pi^\zeta = \pi(\zeta' = 1 - \bar{\zeta} \mid \zeta = 1 + \bar{\zeta}) = \pi(\zeta' = 1 + \bar{\zeta} \mid \zeta = 1 - \bar{\zeta})$.

To govern the correlation between technology and depreciation shocks, let the probability of being in the high (low) depreciation state conditional on being in the low (high) technology state be $\pi^\delta = \pi(\delta' = \delta_0 + \bar{\delta} \mid \zeta' = 1 - \bar{\zeta}) = \pi(\delta' = \delta_0 - \bar{\delta} \mid \zeta' = 1 + \bar{\zeta})$, where the second equality follows from assuming symmetry of the matrix. We then have that the transition matrix of aggregate states follows from the corresponding assignment of

states in (27) as

$$\pi_z = \begin{bmatrix} \pi^\zeta \cdot \pi^\delta & \pi^\zeta \cdot (1 - \pi^\delta) & (1 - \pi^\zeta) \cdot (1 - \pi^\delta) & (1 - \pi^\zeta) \cdot \pi^\delta \\ \pi^\zeta \cdot \pi^\delta & \pi^\zeta \cdot (1 - \pi^\delta) & (1 - \pi^\zeta) \cdot (1 - \pi^\delta) & (1 - \pi^\zeta) \cdot \pi^\delta \\ (1 - \pi^\zeta) \cdot \pi^\delta & (1 - \pi^\zeta) \cdot (1 - \pi^\delta) & \pi^\zeta \cdot (1 - \pi^\delta) & \pi^\zeta \cdot \pi^\delta \\ (1 - \pi^\zeta) \cdot \pi^\delta & (1 - \pi^\zeta) \cdot (1 - \pi^\delta) & \pi^\zeta \cdot (1 - \pi^\delta) & \pi^\zeta \cdot \pi^\delta \end{bmatrix}.$$

In summary, the Markov chain process of aggregate shocks is characterized by four parameters, $(\bar{\zeta}, \bar{\delta}, \pi^\zeta, \pi^\delta)$. We calibrate $\bar{\zeta}$ and π^ζ as first-stage parameters to match the standard deviation and autocorrelation of TFP in the data of 0.029 and 0.88, respectively, 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. To match these moments, our calibration routine yields $\bar{\delta} = 0.078$, $\bar{\zeta} = 0.029$, $\pi^\zeta = 0.941$, 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. Recall that we compare two long-run stochastic equilibria (steady states). As the table shows, the 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 (24). 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%.

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

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.

se, leads to lower consumption and a welfare loss, as 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 is larger than the expected one-period risk-free rate.²³ 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 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, we consider a small open economy in which aggregate prices are determined on world markets and social security is introduced in the small home country. That is, in the PE experiment, 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 the benefits from insurance, because there are no price changes. 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

²³This is consistent with Proposition 1 of Krueger and Kubler (2006), which—loosely speaking—states that for dynamic efficiency, it is sufficient that, at every point in time, there are states of the world in which the bond return is larger than the implicit social security return and which are reached with positive probability.

crowding out stand at -3.27% . The fact that the reduction in aggregate capital leads to a welfare loss means that the equilibrium is dynamically efficient; see the discussion in Subsection 3.8.

Table 3: Benefits from Insurance versus Costs from Crowding Out

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

Notes: For the definition of consumption equivalent variation, g_c , in general equilibrium (GE) and partial equilibrium (PE), cf. Subsection 3.8.

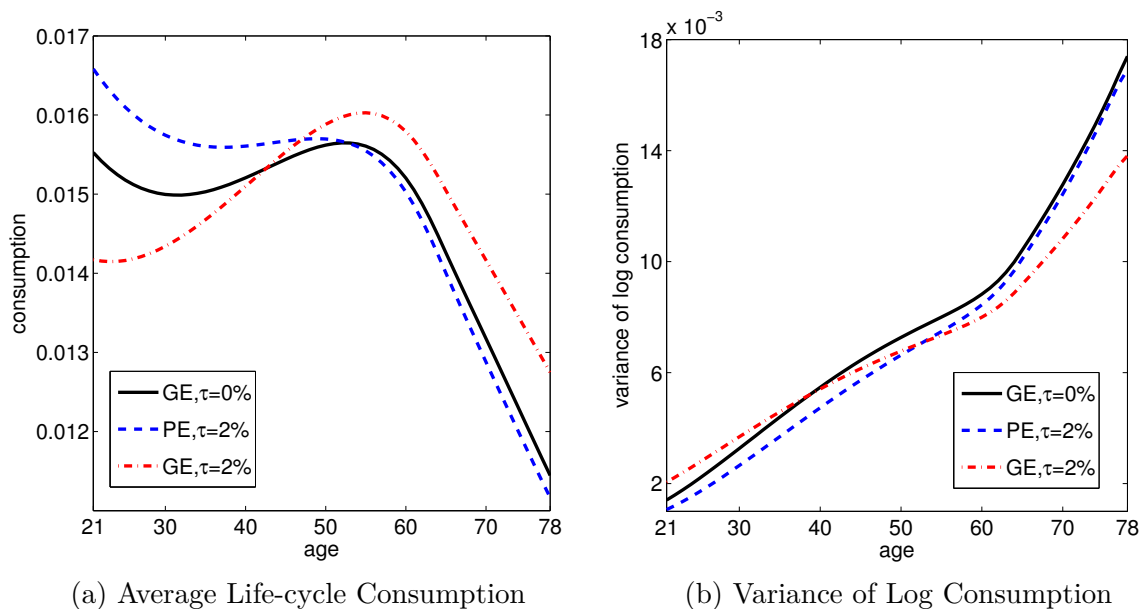
Distributional Effects. The welfare effects reported in Table 3 represent aggregate effects, which encompass important inter- and intragenerational distributional changes. Our two-generations model in Section 2 demonstrated that, from an ex-ante perspective, households value the intra-generational redistribution from the income (and asset) rich to the income (and asset) poor, which results from the lump-sum payments made by social security. In addition, a redistribution over a household's life-cycle might enhance welfare because of market incompleteness. 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).²⁴ The solid line represents the pre-experiment ($\tau = 0$) economy, the dashed line represents the post-experiment ($\tau = 0.02$) partial equilibrium economy, and the dash-dotted line represents the post-experiment ($\tau = 0.02$) general equilibrium economy. 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, see Panel (a) of Figure 1. Due to discounting, the early consumption gains are weighted more strongly than the later consumption losses. Simultaneously, the variance of log consumption decreases at all ages, as shown in Panel (b). 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

²⁴We look at the variance of log consumption, rather than the variance of consumption, because it is a relative inequality measure, invariant to the mean.

higher average returns, as reported previously in Table 2. To profit from higher returns, young households increase their life-cycle savings, which raises consumption at higher ages. As can be seen from Panel (a) of Figure 1, consumption remains below its pre-experiment, general equilibrium level until age 43. Panel (b) shows that, 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. In response to a shock, they hence adjust consumption rather than savings. After the age of 40, the variance of log consumption is smaller than in the pre-experiment economy, and indeed the gap widens considerably with age. This reduced volatility at older ages constitutes the major source of the welfare gains. Thus when comparing the two general equilibria, households expect to lose from the reform when young, because consumption will be lower on average and have a higher volatility at young ages, while they expect to gain on both fronts when old. The fact that we find an overall welfare gain of the reform, as reported in Table 2, means that the gains when old outweigh the losses when young from an expected ex-ante perspective, even though there is discounting.

Figure 1: Life-cycle 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.²⁵ The very close match substantially strengthens our analysis of the distributional 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.

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).

Decomposition into risks and interactions. We now investigate how much of the welfare gains in partial equilibrium of +5.48% can be attributed to insurance against aggregate risk and idiosyncratic wage risk, as well as to the interactions *CCV* and *LCI*. This is a key aspect of our contribution. The procedure for this decomposition is described in Subsection 3.8. Essentially, it performs the same partial equilibrium experiment as described above, but now in addition the risks are turned off. Recall from Section 4 that we set the interest rate exogenously to 4.2% when solving economies without aggregate risk.

The results of this decomposition are 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 risk-free interest rate of $r_b = 0.042$.²⁶ As shown in the third and fourth column, the welfare gains

²⁵See, e.g., Castañeda, Díaz-Giménez, and Ríos-Rull (2003) and De Nardi (2004).

²⁶We arrive at this welfare number both when we solve the model computationally without risks and when we use the analytical formula presented in Observation 2.

from insurance against idiosyncratic risk amount to 0.84% and against aggregate risk to 2.02% in terms of consumption equivalent variations. Thus, in welfare terms, the role played by aggregate risk is approximately twice as important as the role played by idiosyncratic risk. Turning to the interactions, the welfare gains attributable to *LCI* stand at 1.65% and those attributable to *CCV* at 1.60%. Hence, the interactions account for $\frac{dg_c(CC\bar{V})+dg_c(LCI)}{g_c^{PE}} \cdot 100[\%] = \frac{1.65+1.60}{5.48} \cdot 100[\%] \approx 60\%$ of the total insurance gains through social security. Since the previous literature on social security has only included one risk, it has completely missed the important role played by these large interactions between idiosyncratic and aggregate risk.

Table 5: Decomposition of Welfare Benefits in Partial Equilibrium

g_c^{PE}	$g_c^{PE}(0,0)$	$dg_c(IR)$	$dg_c(AR)$	$dg_c(LCI)$	$dg_c(CC\bar{V})$
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 of interactions is large. We next address 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. To this end, 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, $\mu = E[r_{s,t} - r_{b,t}]$, and the volatility of stock returns. Specifically, we target an equity premium of $\mu = 5.60\%$ and a standard deviation of stock returns of $\sigma(r_s) = 16.8\%$, again based on data taken from Rob Shiller's webpage. For the economy without aggregate risk and the deterministic economy, we again target an interest rate of 4.2%. Throughout these experiments, we target a capital output ratio of 2.65 by adjusting the discount factor, β . Table 6 summarizes the second-stage parameters, i.e., the parameters that are jointly calibrated. The remaining first-stage parameters take the same value as in the baseline, see Table 1. Table 7 displays the targeted moments for these economies. For comparison, both tables include the corresponding values of the baseline (*BL*).

Let us first turn to the economy with only aggregate risk, denoted *AR-only*. As mentioned above, this economy is calibrated to match the equity premium, as shown in Table 7. It also matches the volatility of stock returns. However, it overshoots the Sharpe ratio and the volatility of aggregate consumption, the simulated moments of which are $\varsigma = \frac{E[r_{s,t}-r_{b,t}]}{\sigma[r_{s,t}-r_{b,t}]} = 0.352$ and $\sigma\left(\frac{\Delta C_{t+1}}{C_t}\right) = 0.041$, respectively, as compared to 0.33 and 0.03 in the data. The reason is that this economy requires a very high coefficient of risk aversion of $\theta = 14.65$ and, to match the high volatility of stock returns, also large depreciation shocks of $\bar{\delta} = 0.111$, see Table 6.²⁷

Table 6: The Role of Both Risks: Parameters

Scenario	θ	β	δ_0	$\bar{\delta}$	π^δ
<i>BL</i>	3.00	0.981	0.102	0.078	0.886
<i>AR-only</i>	14.65	0.994	0.077	0.111	0.833
<i>IR-only</i>	3.00	0.967	0.079	0.000	NA
<i>No-risk</i>	3.00	0.998	0.079	0.000	NA

Notes: *BL*: baseline calibration with $\theta = 3$; *AR-only*: economy with only aggregate risk, calibrated to match equity premium; *IR-only*: economy with only idiosyncratic risk; *No-risk*: deterministic economy.

Table 7: The Role of Both Risks: Moments

Scenario	ς	μ	$E\left[\frac{K}{Y}\right]$	$E[r_b]$	$\sigma\left(\frac{\Delta C_{t+1}}{C_t}\right)$	$\sigma(r_s)$	$\sigma(\zeta, r_s)$
<i>BL</i>	0.078	0.008	2.65	0.023	0.030	0.104	0.500
<i>AR-only</i>	0.352	0.056	2.65	0.023	0.041	0.168	0.500
<i>IR-only</i>	0.000	0.000	2.65	0.042	NA	NA	NA
<i>No-risk</i>	0.000	0.000	2.65	0.042	NA	NA	NA

Notes: *BL*: baseline calibration with $\theta = 3$; *AR-only*: economy with only aggregate risk, calibrated to match equity premium; *IR-only*: economy with only idiosyncratic risk; *No-risk*: deterministic economy. $\varsigma = \frac{E[r_{s,t}-r_{b,t}]}{\sigma[r_{s,t}-r_{b,t}]}$: Sharpe ratio; $\mu = E[r_{s,t} - r_{b,t}]$: equity premium; $E\left[\frac{K}{Y}\right]$: average capital-output ratio; $E[r_b]$: average bond return; $\sigma\left(\frac{\Delta C_{t+1}}{C_t}\right)$: standard deviation of aggregate consumption; $\sigma(r_s)$: standard deviation of stock returns; $\sigma(\zeta, r_s)$: correlation of TFP shocks and stock returns.

Propositions 2 and 3 show that welfare from introducing social security increases in risk aversion and the volatility of aggregate risk. With respect to these two, the *AR-only*

²⁷However, both the coefficient of risk aversion and the depreciation shocks are smaller than those required by Krueger and Kubler (2006). This is because of the leverage in the stock returns as described in Subsection 3.3.

Table 8: The Role of Both Risks: Benefits and Costs

Scenario	Consumption equivalent variation, g_c		
	GE	PE	Crowding Out
<i>AR-only</i>	-0.54%	-0.38%	-0.16%
<i>IR-only</i>	-1.91%	0.27%	-2.18%
<i>No-risk</i>	-1.13%	-0.62%	-0.51%

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

calibration is an extreme case. Nevertheless, Table 8 documents welfare losses for this case. In general equilibrium, i.e., accounting for the welfare losses from the crowding-out of capital, they stand at -0.54% , again expressed as 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 welfare gains. Finally, introducing social security in the no-risk economy leads to welfare losses in both general and partial equilibrium.

These findings reemphasize our main contribution: To correctly assess the welfare effects of social security, it is crucial to include both idiosyncratic and aggregate risks together in a life-cycle model.

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 finding of long-run welfare gains and sizeable interactions is robust when we consider economies with alternative levels of risk and risk aversion. To isolate the effects of risk and risk aversion, we hold the discount factor β constant at its baseline value. The advantage is that a crucial preference parameter—which has a strong impact on welfare—remains unchanged, making the comparison and interpretation of the results much easier. The disadvantage is that the model will overshoot the capital-output ratio when we increase risk or risk aversion. Despite the too high capital-output ratio, 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 . As shown in

the propositions of Section 2, the expected return on savings is crucial in determining the welfare of introducing social security. The level of the capital-output ratio, on the other hand, is only of second-order importance for our welfare results, because the consumption equivalent variation is defined in terms of a percentage deviation.

The calibrated parameters and targeted moments for the various scenarios we consider are summarized in Tables 9 and 10, respectively. Recall that our baseline calibration features a coefficient of risk aversion of $\theta = 3$. With this value, both the implied equity premium, μ , and the implied Sharpe ratio, ς , are too low when compared to the data. To understand how strongly θ affects the value of social security, we calibrate the model to match a Sharpe ratio of $\varsigma = 0.33$, which we term scenario *SR* in Table 9. This yields a coefficient of risk aversion of $\theta = 10.88$. In this calibration, we continue to match the volatility of consumption growth by an appropriate choice of $\bar{\delta}$.

Table 9: Sensitivity Analysis: Parameters

Scenario	θ	β	δ_0	$\bar{\delta}$	π^δ
<i>IES</i> = 0.5					
<i>BL</i>	3.00	0.981	0.102	0.078	0.886
<i>SR</i>	10.88	0.981	0.022	0.044	0.830
<i>EP</i>	5.51	0.981	-0.002	0.114	0.832
<i>IES</i> = 1.5					
<i>BL</i> _{<i>IES</i>=1.5}	3.00	0.974	0.099	0.040	0.887
<i>SR</i> _{<i>IES</i>=1.5}	11.72	0.974	0.022	0.039	0.831
<i>EP</i> _{<i>IES</i>=1.5}	5.59	0.974	0.003	0.116	0.837

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

In a next step, we instead calibrate the model to match the equity premium and the volatility of stock returns, while retaining the same low bond return. We refer to this calibration as scenario *EP*. *Ceteris paribus*, matching a higher equity premium would require an even 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, at $\theta = 5.51$, is lower than in the *SR*-calibration. Like in the *AR-only* scenario of the previous subsection, the Sharpe ratio and consumption volatility are too high in the *EP* scenario, see Table 10.

We repeat this sensitivity analysis with a higher *IES*. To this end, we first proceed

Table 10: Sensitivity Analysis: Moments

Scenario	ς	μ	$E \left[\frac{K}{Y} \right]$	$E [r_b]$	$\sigma \left(\frac{\Delta C_{t+1}}{C_t} \right)$	$\sigma(r_s)$	$\sigma(\zeta, r_s)$
<i>IES</i> = 0.5							
<i>BL</i>	0.078	0.008	2.65	0.023	0.030	0.104	0.500
<i>SR</i>	0.333	0.020	5.80	0.023	0.030	0.066	0.500
<i>EP</i>	0.357	0.056	7.67	0.023	0.066	0.168	0.500
<i>IES</i> = 1.5							
<i>BL</i> _{<i>IES</i>=1.5}	0.051	0.003	2.65	0.023	0.030	0.054	0.500
<i>SR</i> _{<i>IES</i>=1.5}	0.333	0.018	5.90	0.023	0.030	0.057	0.500
<i>EP</i> _{<i>IES</i>=1.5}	0.356	0.056	6.57	0.023	0.089	0.168	0.500

Notes: *BL*: baseline calibration with $\theta = 3$; *SR*: scenario matching the Sharpe ratio; *EP*: scenario matching the equity premium. $\varsigma = \frac{E[r_{s,t} - r_{b,t}]}{\sigma[r_{s,t} - r_{b,t}]}$: Sharpe ratio; $\mu = E[r_{s,t} - r_{b,t}]$: equity premium; $E \left[\frac{K}{Y} \right]$: average capital-output ratio; $E[r_b]$: average bond return; $\sigma \left(\frac{\Delta C_{t+1}}{C_t} \right)$: standard deviation of aggregate consumption; $\sigma(r_s)$: standard deviation of stock returns; $\sigma(\zeta, r_s)$: correlation of TFP shocks and stock returns.

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.

The welfare results in general equilibrium are presented in Table 11, together with the decomposition into 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: recall that *BL* features $\theta = 3$, *SR* has $\theta = 11$ and *EP* has $\theta = 5.6$. Our two baseline scenarios, which have an *IES* of 0.5 and 1.5, respectively, 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.²⁸

It is interesting that the 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

²⁸We also conducted an experiment with a risk aversion of 2 and an *IES* of 0.5. In this experiment, welfare gains in general equilibrium are 1.47%, confirming the monotonous impact of risk aversion.

Table 11: Sensitivity Analysis: Benefits and Costs

Scenario	Consumption equivalent variation, g_c		
	GE	PE	Crowding Out
<i>IES</i> = 0.5			
<i>BL</i>	+2.21%	+5.48%	-3.27%
<i>SR</i>	+4.16%	+8.32%	-4.16%
<i>EP</i>	+2.78%	+6.53%	-3.75%
<i>IES</i> = 1.5			
<i>BL</i> _{<i>IES</i>=1.5}	+2.56%	+3.65%	-1.10%
<i>SR</i> _{<i>IES</i>=1.5}	+5.08%	+8.09%	-3.01%
<i>EP</i> _{<i>IES</i>=1.5}	+4.45%	+7.34%	-2.89%

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

rates make them save more, thus dampening the crowding out.²⁹ This is the main 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 compare the role played by the interactions of risks across scenarios, we compute the ratio $\frac{dg_c(CCV)+dg_c(LCI)}{g_c^{PE}}$, displayed in Table 12. It amounts to approximately 60 percent in our two baseline scenarios and reaches 71 percent in scenario *EP*_{*IES*=1.5}. Therefore, a share of interactions of roughly 60% of total welfare gains is a robust (lower bound) finding.³⁰

²⁹For the same reason, a higher *IES* leads to smaller volatility of real aggregates in general equilibrium. The fluctuations introduced by the depreciation shocks are counteracted by households' savings, so that, for a higher *IES*, the volatility of capital is smaller, leading to smaller volatility of aggregate output and consumption.

³⁰In the scenario with risk aversion of 2 and an *IES* of 0.5, cf. Footnote 28, the share of interactions is approximately 50%.

Table 12: Sensitivity Analysis: Welfare Share of Interactions

Scenario	Welfare share, $\frac{dg_c(LCI)+dg_c(CCV)}{g_c^{PE}}$
<i>IES</i> = 0.5	
<i>BL</i>	0.59
<i>SR</i>	0.65
<i>EP</i>	0.67
<i>IES</i> = 1.5	
<i>BL</i> _{<i>IES</i>=1.5}	0.60
<i>SR</i> _{<i>IES</i>=1.5}	0.65
<i>EP</i> _{<i>IES</i>=1.5}	0.71

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

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 show that it is crucial to take interactions between the two into account. One interaction is explicit, in the form of a countercyclical, cross-sectional variance of idiosyncratic income risk. The other interaction emerges from the economic structure of a household's life-cycle, even when the two risks are statistically independent. This second interaction, which has not been previously studied, is exposed analytically. In our quantitative analysis, 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 consumption equivalent variation despite significant crowding out of capital. Such welfare gains contrast with standard findings in the related literature. We document that jointly modeling both risks and their interactions is crucial for this finding. Of course, by only examining one risk in isolation, it was not possible to address this in 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 and a countercyclical cross-sectional variance of the income risk, then the welfare costs can become very large, see e.g., Krebs (2007).

While our analysis reveals an important channel and documents that it matters quantitatively, some aspects are not taken into account. We abstract from endogenous labor supply. This may bias the 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 we do in this paper, a realistically calibrated model would only lead to second-order effects of endogenous labor supply 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.³¹ Preliminary results reveal that this makes the analysis much more complex,

³¹Caliendo, Guo, and Hosseini (2013) demonstrate that this does not hold in a stylized general equilibrium with accidental bequests.

because survival risk seems to interact with the other risks. Moreover, it is not straightforward to jointly model survival risk and financial risk while using Epstein-Zin preferences, see Cordoba and Ripoll (2013).

Furthermore, in our economy, intergenerational sharing of aggregate risks 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 manage a pension fund 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.

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.
- Auerbach, A. J., L. J. Kotlikoff, and J. Skinner (1983). The Efficiency Gains from Dynamic Tax Reform. *International Economic Review* 24 (2), 81–100.
- 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., L. Guo, and R. Hosseini (2013). Social Security is NOT a Substitute for Annuities. *Society for Economic Dynamics 2013 Meeting Papers* 680, 1–28.
- 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. (2010). Long-Run Productivity Risk: A New Hope for Production-Based Asset Pricing? *SSRN Electronic Journal*.
- 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.
- Diamond, P. A. (1965). National Debt in a Neoclassical Growth Model. *American Economic Review* 55(5), 1126–1150.
- 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 (2006). Multiplicative Background Risk. *Management Science* 52, 146–153.
- Gollier, C. and J. W. Pratt (1996). Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica* 64, 1109–1123.
- Golosov, M., A. Shourideh, M. Troshkin, and A. Tsyvinski (2013). Optimal Pension Systems with Simple Instruments. *American Economic Review, Papers and Proceedings* 103(3), 502–507.

- 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 (2013). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, forthcoming.
- Harenberg, D. and A. Ludwig (2014). On the Welfare Effects of Social Security in an Analytically Tractable Overlapping Generations Model with Aggregate and Idiosyncratic Risk, *Working Paper*.
- 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. *The 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). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106(5), 867–896.
- Kuhle, W. (2012, December). Dynamic efficiency and the two-part golden rule with heterogeneous agents. *Journal of Macroeconomics* 34(4), 992–1006.
- 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, January). 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. Thøgersen (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.
- Santis, M. D. and M. De Santis (2007). Individual Consumption Risk and the Welfare Cost of Business Cycles. *American Economic Review* 97(4), 1488–1506.
- 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.
- Shourideh, A. and M. Troshkin (2012). Providing Efficient Incentives to Work: Retirement Ages and the Pension System. *SSRN Electronic Journal*, 1–43.
- 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.
- Vickrey, W. S. (1947). *Agenda for Progressive Taxation*. New York: Ronald.
- Weil, P. (1989). The Equity Premium Puzzle and the Risk-Free Rate Puzzle. *Journal of Monetary Economics* 24, 401–421.
- Werning, I. (2007). Optimal Fiscal Policy with Redistribution. *Quarterly Journal of Economics* 122(3), 925–967.

Appendix

A Proofs

Proof of Proposition 1. Maximize

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

which is equivalent to

$$\max \frac{1}{1-\theta} E R_{p,t,t+1}^{1-\theta}$$

where $R_{p,t,t+1} \equiv \eta_{i,1,t} \zeta_t \bar{R} \varrho_{t+1} + \tau \left((1+\lambda) \zeta_{t+1} - \bar{R} \eta_{i,1,t} \zeta_t \varrho_{t+1} \right)$ is a consumption (or portfolio) return. Increasing ex-ante utility for a marginal introduction of social security requires the first-order condition w.r.t. τ to exceed zero, hence:

$$E \left[R_{p,t,t+1}^{-\theta} \frac{\partial R_{p,t,t+1}}{\partial \tau} \right] \Bigg|_{\tau=0} > 0. \quad (28)$$

Taking the according partial derivatives we get:

$$\begin{aligned} (1+\lambda) E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{-\theta} \zeta_{t+1} \right] &> \bar{R} E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{1-\theta} \right] \\ \Leftrightarrow (1+\lambda) E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{-\theta} \right] E \left[\zeta_{t+1} \right] &> \bar{R} E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{1-\theta} \right] \\ \Leftrightarrow (1+\lambda) E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{-\theta} \right] &> \bar{R} E \left[(\eta_{i,1,t} \zeta_t \varrho_{t+1})^{1-\theta} \right] \end{aligned} \quad (29)$$

where the transformations follow from Assumption 1. Rewriting gives equation (4). \square

Proof of Proposition 2. Define

$$\begin{aligned} Z_1 &\equiv (\eta_{i,1,t} \zeta_t \varrho_{t+1})^{-\theta} \\ Z_2 &\equiv (\eta_{i,1,t} \zeta_t \varrho_{t+1})^{1-\theta}. \end{aligned}$$

By log-normality we have that $EZ_i = \exp(E \ln Z_i + \frac{1}{2} \sigma_{\ln Z_i}^2)$, $i = 1, 2$. Turning first to Z_1

observe that

$$\begin{aligned} E \ln Z_1 &= -\theta (E \ln \eta_{i,1,t} + E \ln \tilde{\varrho} + E \ln \zeta) \\ \sigma_{\ln Z_1}^2 &= \theta^2 (\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2). \end{aligned}$$

Therefore:

$$\begin{aligned} E[Z_1] &= \exp \left(-\theta \left(E \ln \eta_{i,1,t} + \frac{\sigma_{\ln \eta}^2}{2} \right) \right) \cdot \exp \left(-\theta \left(E \ln \varrho + \frac{\sigma_{\ln \varrho}^2}{2} \right) \right) \\ &\quad \cdot \exp \left(-\theta \left(E \ln \zeta + \frac{\sigma_{\ln \zeta}^2}{2} \right) \right) \exp \left(\frac{1}{2} \theta (1 + \theta) (\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2) \right) \\ &= (E[\eta_{i,1,t}])^{-\theta} (E[\varrho])^{-\theta} (E[\zeta])^{-\theta} \cdot \exp \left(\frac{1}{2} \theta (1 + \theta) (\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2) \right) \\ &= \exp \left(\frac{1}{2} \theta (1 + \theta) (\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2) \right), \end{aligned}$$

where the last line follows from Assumption 1b.

Next, observe that for $x \in \{\eta, \zeta, \varrho\}$, log-normality implies that

$$\begin{aligned} \sigma_x^2 &= \exp \left(2E \ln x + \sigma_{\ln x}^2 \right) \left(\exp \left(\sigma_{\ln x}^2 \right) - 1 \right) \\ &= (Ex)^2 \left(\exp \left(\sigma_{\ln x}^2 \right) - 1 \right) \\ &= \left(\exp \left(\sigma_{\ln x}^2 \right) - 1 \right), \end{aligned}$$

where the last line again follows from Assumption 1b. Therefore:

$$\exp \left(\frac{1}{2} \theta (1 + \theta) (\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2) \right) = \left((1 + \sigma_{\eta}^2)(1 + \sigma_{\zeta}^2)(1 + \sigma_{\varrho}^2) \right)^{\frac{1}{2} \theta (1 + \theta)}.$$

We consequently have

$$E[Z_1] = \left((1 + \sigma_{\eta}^2)(1 + \sigma_{\zeta}^2)(1 + \sigma_{\varrho}^2) \right)^{\frac{1}{2} \theta (1 + \theta)}.$$

As to Z_2 observe that

$$\begin{aligned} E \ln Z_2 &= (1 - \theta) (E \ln \eta_{i,1,t} + E \ln \zeta + E \ln \tilde{\varrho}) \\ \sigma_{\ln Z_2}^2 &= (1 - \theta)^2 (\sigma_{\ln \eta}^2 + \sigma_{\ln \zeta}^2 + \sigma_{\ln \varrho}^2). \end{aligned}$$

Therefore:

$$\begin{aligned}
E[Z_2] &= \exp\left((1-\theta)\left(E \ln \eta_{i,1,t} + \frac{\sigma_{\ln \eta}^2}{2}\right)\left(E \ln \varrho + \frac{\sigma_{\ln \varrho}^2}{2}\right)\left(E \ln \zeta + \frac{\sigma_{\ln \zeta}^2}{2}\right)\right) \\
&\quad \cdot \exp\left(\frac{1}{2}\theta(\theta-1)\left(\sigma_{\ln \eta}^2 + \sigma_{\ln \varrho}^2 + \sigma_{\ln \zeta}^2\right)\right) \\
&= \left((1 + \sigma_{\eta}^2)(1 + \sigma_{\zeta}^2)(1 + \sigma_{\varrho}^2)\right)^{\frac{1}{2}\theta(\theta-1)}
\end{aligned}$$

and hence

$$\frac{E[Z_1]}{E[Z_2]} = \left((1 + \sigma_{\eta}^2)(1 + \sigma_{\zeta}^2)(1 + \sigma_{\varrho}^2)\right)^{\theta}.$$

□

Proof of Proposition 3. To evaluate the CEV between two scenarios, i.e., comparing $Eu(c_{i,2,t+1}^{\tau > 0})$ with $Eu(c_{i,2,t+1}^{\tau = 0})$. We use that

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

and evaluate this expression at $\tau = 0$.

1. **Case $\theta \neq 1$.** We have that, evaluated at $\tau = 0$,

$$\begin{aligned}
\frac{\partial Eu(c_{i,2,t+1}^{\tau = 0})}{\partial \tau} &= \bar{w}_t^{1-\theta} E \left[\left(\bar{R} \eta \zeta_t \varrho_{t+1} \right)^{-\theta} \cdot \left((1 + \lambda) \zeta_{t+1} - \bar{R} \eta \zeta_t \varrho_{t+1} \right) \right] \\
&= \bar{w}_t^{1-\theta} \left(\bar{R}^{-\theta} (1 + \lambda) E \left[(\eta \zeta_t \varrho_{t+1})^{-\theta} \zeta_{t+1} \right] - \bar{R}^{1-\theta} E \left[(\eta \zeta_t \varrho_{t+1})^{1-\theta} \right] \right) \\
&= \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} \left(\frac{1 + \lambda}{\bar{R}} E Z_1 - E Z_2 \right)
\end{aligned}$$

where Z_1, Z_2 are defined in our proof of Proposition 2.

We also have that

$$\begin{aligned}
Eu(c_{i,2,t+1}^{\tau = 0}) &= \frac{1}{1 - \theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} E (\eta_{i,1,t} \zeta_t \varrho_{t+1})^{1-\theta} \\
&= \frac{1}{1 - \theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} E Z_2.
\end{aligned}$$

Therefore:

$$\begin{aligned} Eu(c_{i,2,t+1}^{\tau>0}) &= \frac{1}{1-\theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} EZ_2 \\ &\quad + \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} \left(\frac{1+\lambda}{\bar{R}} EZ_1 - EZ_2 \right) d\tau. \end{aligned}$$

The CEV, denoted by g_c , is defined by the relationship:

$$Eu(c_{i,2,t+1}^{\tau=0}(1+g_c)) = Eu(c_{i,2,t+1}^{\tau>0}),$$

from which, using the above formulae, we get

$$\begin{aligned} (1+g_c^{PE})^{1-\theta} \frac{1}{1-\theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} EZ_2 &= \frac{1}{1-\theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} EZ_2 \\ &\quad + \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} \left(\frac{1+\lambda}{\bar{R}} EZ_1 - EZ_2 \right) d\tau, \end{aligned}$$

where we put a superscript PE to indicate that this is the CEV in partial equilibrium.

Hence:

$$\begin{aligned} (1+g_c^{PE})^{1-\theta} &= 1 + \frac{\bar{w}_t^{1-\theta} \bar{R}^{1-\theta} \left(\frac{1+\lambda}{\bar{R}} EZ_1 - EZ_2 \right)}{\frac{1}{1-\theta} \bar{w}_t^{1-\theta} \bar{R}^{1-\theta} EZ_2} d\tau \\ &= 1 + (1-\theta) \left(\frac{1+\lambda}{\bar{R}} \frac{EZ_1}{EZ_2} - 1 \right) d\tau \\ &= 1 + (1-\theta) \left(\frac{1+\lambda}{\bar{R}} (1+TR)^\theta - 1 \right) d\tau \end{aligned}$$

where the last line again follows from the proof of Proposition 2.

Hence,

$$g_c^{PE} = \left(1 + (1-\theta) \left(\frac{1+\lambda}{\bar{R}} (1+TR)^\theta - 1 \right) d\tau \right)^{\frac{1}{1-\theta}} - 1,$$

or, expressed in logs, i.e., using that $g_c^{PE} \approx \ln(1 + g_c^{PE})$ for small g_c^{PE} , we get

$$\begin{aligned} g_c^{PE} &\approx \frac{1}{1-\theta} \cdot \ln \left(1 + (1-\theta) \left(\frac{1+\lambda}{\bar{R}} (1+TR)^\theta - 1 \right) d\tau \right) \\ &\approx \left(\frac{1+\lambda}{\bar{R}} (1+TR)^\theta - 1 \right) d\tau \end{aligned} \quad (30)$$

2. **Case $\theta = 1$.** We have that, evaluated at $\tau = 0$,

$$\begin{aligned} \frac{\partial Eu(c_{i,2,t+1}^{\tau=0})}{\partial \tau} &= E \left[\left(\bar{R} \eta \zeta_t \varrho_{t+1} \right)^{-1} \cdot \left((1+\lambda) \zeta_{t+1} - \bar{R} \eta \zeta_t \varrho_{t+1} \right) \right] \\ &= \left(\bar{R}^{-1} (1+\lambda) E \left[\left(\eta \zeta_t \varrho_{t+1} \right)^{-1} \zeta_{t+1} \right] - 1 \right) \\ &= \frac{1+\lambda}{\bar{R}} EZ_1 - 1. \end{aligned}$$

We also have that

$$Eu(c_{i,2,t+1}^{\tau=0}) = \ln(\bar{w}_t \bar{R}) + E \ln(\eta_{i,1,t} \zeta_t \varrho_{t+1}).$$

Therefore:

$$Eu(c_{i,2,t+1}^{\tau>0}) = \ln(\bar{w}_t \bar{R}) + E \ln(\eta_{i,1,t} \zeta_t \varrho_{t+1}) + \left(\frac{1+\lambda}{\bar{R}} EZ_1 - 1 \right) d\tau.$$

For g_c^{PE} we accordingly get

$$1 + g_c^{PE} = \exp \left(\left(\frac{1+\lambda}{\bar{R}} EZ_1 - 1 \right) d\tau \right).$$

Approximating the above in logs we get:

$$g_c^{PE} \approx \left(\frac{1+\lambda}{\bar{R}} (1+TR) - 1 \right) d\tau$$

which is the same as equation (30) for $\theta = 1$.

Taking a first-order Taylor series expansion of equation (30) round $TR = 0$ we get

$$g_c^{PE} \approx \left(\frac{1+\lambda}{\bar{R}} - 1 + \theta \frac{1+\lambda}{\bar{R}} TR \right) \cdot d\tau.$$

□

Supplementary Appendix

B Supplementary Computational Appendix

B.1 Overview

The numerical solution follows Gomes and Michaelides (2008), Storesletten, Telmer, and Yaron (2007), and Krusell and Smith (1997, 1998). The algorithm consists of the following steps, details of which are given in the next subsections.

1. Choose arguments and a functional form for the approximate law of motion, and make an initial guess for its coefficients.
2. Given the approximate law of motion, solve the household's problem.
3. Simulate the economy using the obtained optimal policy functions. In every period, compute the market clearing prices.
4. Update the coefficients of the approximate law of motion by running a regression on the simulated aggregate statistics.
5. If the coefficients have converged, and the R^2 of the regression is sufficiently high, stop, else go to 2.
6. Repeat steps 1 to 5 for different arguments and functional forms of the law of motion. Select the one with the highest R^2 .
7. Given the functional form for the approximate law of motion that achieved the best fit, calibrate the economy to match the targets.
 - (a) Provide an initial guess for the parameters to be calibrated.
 - (b) Given the parameters, repeat steps 2 to 5.
 - (c) Calculate the target statistics from the simulations. If they are close to the targets in the data, stop, else update the guess for the parameters and go to 7b.
8. Given the calibrated parameters, increase the social security contribution rate and compute the new general equilibrium by repeating steps 2 to 5.
9. Compute the welfare gains of the experiment in general equilibrium from the simulated variables of the first and the second economy.

10. Given the approximate laws of motion and the simulated prices of the first economy, perform the risk decomposition analysis.
 - (a) Given the approximate law of motion of the first economy, solve the household's problem.
 - (b) Given the simulated prices of the first economy, simulate the economy using the obtained optimal policy functions. (Do not compute market clearing prices.)
 - (c) Increase the social security contribution rate and repeat steps 10a and 10b.
 - (d) Compute the welfare gains of the experiment in partial equilibrium (PE) from the simulated variables of the pre-experiment PE and the post-experiment PE.
 - (e) If this was the no-risk, deterministic economy, stop, else turn off a risk and repeat steps 10a to 10d.

The numerical solution is implemented in Fortran, using an object-oriented approach enabled since the Fortran 2003 standard. Given the accuracy described below and running on 16 cores, it takes on average 4 days to get the solution even for a good initial guess.

B.2 Solving for the approximate law of motion

The idea behind the Krusell-Smith-method (1997, 1998) is to approximate the infinite dimensional distribution, Φ , by a finite number of statistics. The household then uses a law of motion of these statistics, $\hat{H}(\cdot)$, as an approximation to the true law of motion of the distribution, $H(\Phi, z)$. The statistics have to enable the household to forecast the prices that it needs to solve its optimization problem. We follow Krusell and Smith (1997), Gomes and Michaelides (2008), and Storesletten, Telmer, and Yaron (2007), and choose mean aggregate capital, k , together with a second variable to forecast the bond return. As this second variable, we choose the expected equity premium, $\mu = \mathbb{E}(r'_s - r'_b)$ (see Storesletten, Telmer, and Yaron (2007)).³² Thus, the approximate law of motion becomes³³

$$\{k'(z'), \mu'(z')\} = \hat{H}(k, \mu, z, z').$$

³²We choose μ instead of the bond price because this enables us to avoid $E(r_b) > E(r_s)$ by construction. This is desirable because such a situation would never arise in equilibrium.

³³ $\hat{H}(\cdot)$ now has z' as an argument because the distribution will be defined on cash-on-hand, as described in the household problem below. Cash-on-hand at the beginning of a period depends on the realization of the aggregate shock in that period through the aggregate stock return and CCV . Therefore, the shock z' enters the law of motion.

The functional form for \hat{H} that gives the best approximation in our baseline economy is

$$\ln k_{t+1} = \psi_{0,z}^k + \psi_{1,z}^k \ln k_t + \psi_{2,z}^k (\ln k_t)^2 \quad (31a)$$

$$\mu_{t+1,z'} = \psi_{0,z'}^\mu + \psi_{1,z'}^\mu \ln k_{t+1} + \psi_{2,z'}^\mu (\ln k_{t+1})^2 \quad (31b)$$

This is similar to the best fit regression found by Storesletten, Telmer, and Yaron (2007). Note that the forecast of capital, $\ln k_{t+1}$, enters as a regressor in eq. (31b). Effectively, the forecast for $\mu_{t+1,z'}$, which is conditional on z' , depends on $\ln k_t$ and z through the forecast of $\ln k_{t+1}$. The discrete, aggregate shock, z , can take four values, so that we estimate eight equations. Therefore, we report eight coefficients of determination, which for the baseline economy are $R_k^2 = \{0.9998, 0.9999, 0.9998, 0.9998\}$ and $R_\mu^2 = \{0.9918, 0.9945, 0.9584, 0.9695\}$. For the other economies, the R^2 are always higher.³⁴

To find the coefficients, we solve $g(\Psi) = \Psi - \tilde{\Psi}(\Psi)$, where Ψ collects all the coefficients, i.e. $\Psi = \{\psi_{l,z}^m\}_{l=\{0,1,2\}, z=\{1,2,3,4\}, m=\{k,\mu\}}$. To solve this nonlinear equation system, a multidimensional Broyden algorithm is used. During the solution, we normalize (and subsequently de-normalize) the coefficients around unity. For these coefficients around unity, the convergence criterion is $\max\{|g(\Psi)|\} < 1.0^{-7}$. The Newton-like update steps are limited to a small length, and backtracking is used to find an update, if the first step was too large.³⁵

B.3 Solving the household's problem

First, we rewrite the household problem in terms of cash-on-hand, \tilde{x} . This reduces the state space by one dimension, so that the idiosyncratic state consists of (\tilde{x}, e) . Second, we recast the two control variables bond, \tilde{b}' , and stock, \tilde{s}' , as total savings, \tilde{a}' , and the portfolio share invested in stock, κ . This enables us to employ the endogenous grid method proposed by Carroll (2006), as detailed below. And third, we replace the distribution, Φ , by the approximation discussed in the previous section, so that the aggregate state consists of (k, μ, z) . With a slight abuse of notation,³⁶ the optimization problem in recursive form

³⁴For example, for the equity premium calibration with $IES = 0.5$, the coefficients of determination are $R_k^2 = \{0.9999, 0.9999, 0.9999, 0.9999\}$ and $R_\mu^2 = \{0.9961, 0.9968, 0.9944, 0.9949\}$. This economy is the closest to Storesletten, Telmer, and Yaron (2007) and Gomes and Michaelides (2008), and the R^2 are very close to the ones reported there.

³⁵The Newton-like update step is $\Psi_{i+1} = \Psi_i - sJ(\Psi)^{-1}g(\Psi)$, where $J(\Psi)$ is a finite-difference approximation to the Jacobi matrix of the system of equations and s determines the maximum step length.

³⁶Technically, some variables would need to be renamed, e.g. \tilde{y} to $\tilde{\tilde{y}}$, because the state space is now different than the one in Definition 1. For sake of readability, we do not change the notation.

then reads

$$\begin{aligned} & \max_{\tilde{c} > 0, \tilde{a}', \kappa} u_j(\tilde{x}, e; k, \mu, z) \\ & = \begin{cases} \left(\tilde{c}^{\frac{1-\theta}{\gamma}} + \tilde{\beta} \left(\sum_{z'} \sum_{e'} \pi_z(z'|z) \pi_e(e'|e) u_{j+1}^{1-\theta}(\tilde{x}', e'; \hat{H}(k, \mu, z, z'), z') \right)^{\frac{1}{\gamma}} \right)^{\frac{\gamma}{1-\theta}} \\ \tilde{c} \end{cases} \quad \text{if } j = J \end{aligned}$$

s. t.

$$\begin{aligned} \tilde{x}' &= \tilde{a}' \frac{(1 + r_b' + \kappa(r_s' - r_b'))}{1 + \lambda} + \tilde{y}'_{j+1}, \\ \tilde{a}' &\geq 0 \quad \text{if } j = J, \end{aligned}$$

where $\tilde{\beta} = \beta(1 + \lambda)^{\frac{1-\theta}{\gamma}}$, $r_s' = r_s(\hat{H}(k, \mu, z, z'), z')$, $r_b' = r_b(\hat{H}(k, \mu, z, z'), z')$, and income in the next period is given by

$$\tilde{y}'_{j+1} = \tilde{y}_{j+1}(e', \hat{H}(k, \mu, z, z'), z') = \begin{cases} (1 - \tau) \tilde{w}(\hat{H}(k, \mu, z, z'), z') \epsilon_{j+1} \eta(e', z') & \text{if } j + 1 < J \\ \tilde{y}_{ss}(\hat{H}(k, \mu, z, z'), z') & \text{else} \end{cases}$$

The budget constraint contains a growth adjustment of $\frac{1}{1+\lambda}$, because x' is cash on hand at the beginning of next period, while a' is the savings choice made in the previous period. In contrast, the budget constraint in the equilibrium definition of Subsection 3.5 contains only contemporaneous variables, i.e., states and choices in the current period, so that no growth adjustment is needed there.

Applying the envelope theorem and simplifying we get the two first-order-conditions³⁷

$$\mathbb{E} \left[u_{j+1}(\cdot)^{\frac{(1-\theta)(\gamma-1)}{\gamma}} (\tilde{c}')^{\frac{1-\theta-\gamma}{\gamma}} (r_s' - r_b') \right] = 0, \quad (32a)$$

$$\tilde{c} = \left(\tilde{\beta} \frac{1 + r_b'}{1 + \lambda} \left(\mathbb{E} [u_{j+1}(\cdot)^{1-\theta}] \right)^{\frac{1-\gamma}{\gamma}} \mathbb{E} \left[u_{j+1}(\cdot)^{\frac{(1-\theta)(\gamma-1)}{\gamma}} (\tilde{c}')^{\frac{1-\theta-\gamma}{\gamma}} \right] \right)^{\frac{\gamma}{1-\theta-\gamma}}. \quad (32b)$$

To solve for the optimal choices $(\tilde{c}, \tilde{a}', \kappa)$, we apply a variant of the endogenous grid method proposed first by Carroll (2006). In fact, essentially we follow a simplified version of the two-step procedure of Hintermaier and Koeniger (2010). The exogenous grid is defined on total assets in the next period, \tilde{a}' . For a given grid-point \tilde{a}'_i , we first solve eq. (32a) for the portfolio share κ using Brent's root-finding method. Then, given \tilde{a}'_i and

³⁷See Weil (1989) for the envelope theorem with recursive Epstein-Zin preferences.

the corresponding $\kappa(\tilde{a}'_i)$, we use eq. (32b) to get the optimal consumption, $\tilde{c}_i(\tilde{a}'_i)$. Finally, the budget constraint $\tilde{x} = \tilde{c} + \tilde{a}'$ gives us the endogenous grid-point \tilde{x}_i that corresponds to the optimal choices $(\tilde{a}'_i, \tilde{c}_i)$.

When evaluating the expectations, we interpolate u_{j+1} and \tilde{c}' by multidimensional linear interpolation in the continuous states \tilde{x}, k, μ . The aggregate shock z and the idiosyncratic shock e are both discrete and follow a discrete Markov chain. For both, we construct the Markov transition matrix with the Rouwenhorst method (Kopecky and Suen (2010)). With the Rouwenhorst method, it is straightforward to implement the countercyclical cross-sectional variance, CCV , because the variances affect only the grid and not the transition matrix, which in turn is determined purely by ρ .

As is standard in life-cycle models, we iterate backwards, starting with the last generation J , for which the solution is $\tilde{c}_J = \tilde{x}_J$, since they do not leave bequests. In the backwards iteration, we construct age-dependent, exogenous grids $\{\tilde{a}'_{i,j}\}_{i,j}$ to improve the approximation quality. The solution is parallelized in the dimension k , so that for each generation, the solution for all values of k is computed in parallel.

We discretize the state space in the following way. The continuous state variables cash-on-hand, \tilde{x} , aggregate capital, k , and equity premium, μ , have 20, 16, and 10 grid-points, respectively. The discrete state variables, which are the number of generations, J , the idiosyncratic shock, e , and the aggregate shock, z , have 58, 4, and 4 grid-points, respectively. We check that this is sufficient by doubling each of the grid-points in turn and find no change to our results. The first-order-condition in eq. (32a) is solved to an accuracy of 1.0^{-10} .

B.4 Simulating the economy

We simulate the economy 16 times for 4000 periods each time and throw away the first 1000 periods, so that we are left with a total 48.000 simulation periods.³⁸ In each period, we record the aggregates, the life-cycle profiles, and the distribution. The aggregates are needed to estimate the laws of motion, and to calibrate the economy. Like in the solution of the household problem, the optimal policy functions are interpolated in the dimensions of the aggregate states k, μ by multidimensional linear interpolation.

The distribution over households is normalized to a mass of one. We do not simulate

³⁸We found that a large number of simulation periods is necessary for the distribution to converge in the sense that increasing the number of simulation periods does not change the results. In particular, we found that for less than 30.000 simulation periods, the means and standard deviations of the aggregates as well as the estimates of the laws of motion are still sensitive to the number of periods.

many, discrete household units; instead we keep the continuum of households and approximate the distribution with continuous methods. As described in Subsection 3.1, the Law of Large Numbers implies that $\pi_e(e'|e)$ represents the fraction of the population moving from idiosyncratic state e to e' . Therefore, we get a nearly exact approximation in that dimension. In the cash-on-hand dimension, the distribution is discretized on a much finer grid than the policy functions obtained in the household solution, as proposed by Ríos-Rull (1999). This finer discretization improves the approximation quality substantially and helps in ensuring that no households are stuck on the bounds of the distribution. If the lowest or the highest points of the distribution have positive mass, then the cash-on-hand grid is extended and the discretization is made finer.

In each period, the beginning-of-period distribution is iterated forward by using the computed optimal policy functions and the realizations of the shock. For a given cash-on-hand at the beginning of the period, the implied cash-on-hand in the following period will almost always lie between two grid points. Since we are dealing with a continuum of households, we assign a fraction $(1 - f)$ to the lower grid point and f to the upper grid point of the interval which contains the implied cash-on-hand, where f is the distance to the lower grid point.

In each period t , we calculate the market-clearing prices. The current stock return, $r_s(\Phi_t, z_t)$ is given by the contemporaneous aggregate capital and aggregate shock. The current bond return, $r_b(\Phi_t)$, is determined one period before by the bond market clearing condition. We compute it with a nonlinear equation solver to an accuracy of 1.0^{-8} .

We make sure that the grid for the aggregate states is large enough by checking whether the realized values lie on the bounds of the grid. If they do, the grid is increased. To get good initial guesses for the bounds of the aggregate grids and the distribution over households, we compute a degenerate equilibrium, where the realization of the aggregate shocks in the simulations is always equal to their mean. We call this a mean-shock equilibrium.

To check the accuracy of the solution, we compute in each period the 'aggregation error' and the 'income error'. The aggregation error $e_t^{agg} = \frac{Y_t - C_t - I_t}{Y_t}$ says by how much the aggregate budget constraint is violated due to interpolation and aggregation errors, expressed in percent of output. For all economies, the maximum aggregation error is in the order of 1.0^{-6} and the average is in the order of 2.0^{-9} . The income error comes from Euler's formula, which says that total output must equal total factor income. Again expressed in terms of output, we find that it never exceeds 1.0^{-14} .

B.5 Calibrating the economy


The calibration procedure is cast as a system of nonlinear equations. Let \mathcal{T} denote the target statistics in the data and \mathcal{P} the model parameters to be calibrated. For given \mathcal{P} , $\hat{\mathcal{T}}(\mathcal{P})$ are the model-generated statistics, which we get from the simulations. Then the calibration procedure tries to find a root of $\mathcal{T} - \hat{\mathcal{T}}(\mathcal{P}) = 0$. We use Broyden's multidimensional secant method to solve the system to an accuracy of 1.0^{-4} .

C Supplementary Calibration Appendix

C.1 Households

We compute the age-specific productivity profile ϵ_j from PSID data by applying the method of Huggett, Ventura, and Yaron (2011).³⁹ Figure 2 displays the raw data and the smoothed profile. The graphs shows the data when identifying cohort effects, but it looks very similar for time effects.

Figure 2: Life-cycle Productivity



Notes: Age-specific productivity profile ϵ_j from PSID data.

³⁹The data is taken from Ludwig, Schelkle, and Vogel (2012).

C.2 Firms

To estimate α , we take data on total compensation of employees (NIPA Table 1.12) and deflate it with the GDP deflator (NIPA Table 1.1.4). In the numerator, we adjust GDP (NIPA Table 1.1.5), again deflated by the GDP deflator, by nonfarm proprietors' income and other factors that should not be directly related to wage income. Without these adjustments, our estimate of α would be considerably higher (at $\alpha = 0.43$).

To measure capital, we take the stock of fixed assets (NIPA Table 1.1), appropriately deflated. We relate this to total GDP.

We determine the growth rate of technology λ by estimating the Solow residual from the production function, given our estimate of α , our measure for capital, and a measure of labor supply determined by multiplying all full- and part-time employees in domestic employment (NIPA Table 6.4A) with an index for aggregate hours (NIPA Table 6.4A). Notice that we ignore age-specific productivity which should augment our measure of employment. We then fit a linear trend specification to the Solow residual. Acknowledging the labor augmenting technological progress specification chosen, this gives our point estimate.

C.3 Aggregate Risk

First, we discuss the correlation of TFP and stock returns, $\sigma(\zeta_t, r_{s,t})$, a second-stage calibration target. Linear detrending of the data, as done, e.g., by Krueger and Kubler (2006), results in $\sigma(\zeta_t, r_{s,t}) < 0$ as well as a negative correlation of wages and asset returns, i.e., $\sigma(w_t, r_{s,t}) < 0$. Not only does this seem counter to economic intuition in an annual RBC model, but our estimate for $\sigma(\zeta_t, r_{s,t})$ is also statistically insignificant.

Assuming, on the contrary, a unit root process for the log of TFP and detrending by first differences yields a highly significant positive correlation of $\sigma(\zeta_t, r_{s,t}) = 0.50$ (p-value 0.00). Now also $\sigma(w_t, r_{s,t})$ is positive and significant with $\sigma(w_t, r_{s,t}) = 0.306$ (p-value 0.025), which coincides with our economic intuition as we would expect these variables to co-move over the cycle. To make this estimate consistent with the model setup, we transform the numbers to an equivalent deterministic trend specification in the following way. We compute the implicit average horizon h in the unit root specification so that the unconditional variance over h periods coincides with the specification of Krueger and Kubler (2006). This gives an average horizon of $h = 19.2751$ years.⁴⁰

⁴⁰Observe that the unit root estimates in fact imply even stronger aggregate fluctuations. Adjusting the variance in the linear trend specification such that the average horizon equals the average horizon

D Supplementary Appendix: Additional Proofs

Proof of Observation 1. The property follows from homotheticity of Epstein-Zin preferences. To prove it we proceed by induction. We look at two alternative (expected) consumption streams \tilde{c}^A and \tilde{c}^B . One can think of them as optimal consumption under policy regime A and B . We ask how big the percentage increase of consumption stream \tilde{c}^A in each period has to be to reach the same utility level as reached for consumption stream \tilde{c}^B . For sake of simplicity we drop indices t and i and adopt the notation $\tilde{u}_j^X \equiv \tilde{u}_j(c^X)$ for $X \in \{A, B\}$.

1. *Induction claim:* At each age j we have that

$$\tilde{u}_j^B = (1 + g_c)\tilde{u}_j^A.$$

2. *Induction start:* For our Epstein-Zin utility specification (cf. equation 10), at age J we have that

$$\tilde{u}_J^A = \tilde{c}_J^A \quad \text{and} \quad \tilde{u}_J^B = \tilde{c}_J^B.$$

Hence, by the induction claim, we get

$$\tilde{u}_J^B = (1 + g_c)\tilde{u}_J^A = (1 + g_c)\tilde{c}_J^A$$

and, correspondingly,

$$\begin{aligned} \tilde{u}_{J-1}^B &= \left[(\tilde{c}_{J-1}^B)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_{J-1}(\tilde{u}_J^B)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\ &= \left[((1 + g_c)\tilde{c}_{J-1}^A)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_{J-1}((1 + g_c)\tilde{c}_J^A)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\ &= (1 + g_c) \left[(\tilde{c}_{J-1}^A)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_{J-1}(\tilde{c}_J^A)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\ &= (1 + g_c)\tilde{u}_{J-1}^A. \end{aligned}$$

3. *Induction step:* Using the induction claim for any period $j < J - 1$ we therefore

of households in our model, appropriately adjusted to account for the correlation of TFP innovations, gives an average horizon of 34.88 years. This implies a standard deviation of 0.039. Relative to our baseline calibration this means that the standard deviation of innovations increases by roughly 76 percent. However, the overall effects of this additional increase in risk are small. Results are available upon request.

have:

$$\begin{aligned}
\tilde{u}_j^B &= \left[(\tilde{c}_j^B)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_j(\tilde{u}_{j+1}^B)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\
&= \left[((1+g_c)\tilde{c}_j^A)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_j((1+g_c)\tilde{u}_{j+1}^A)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\
&= (1+g_c) \left[(\tilde{c}_j^A)^{\frac{1-\theta}{\gamma}} + \beta \left(\mathbb{E}_j(\tilde{u}_{j+1}^A)^{1-\theta} \right)^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \\
&= (1+g_c)\tilde{u}_j^A.
\end{aligned}$$

□

Proof of Observation 2. The property follows from linearity of consumption policy functions in initial wealth which we first establish. We again simplify notation and drop the i and t indices. Recursive substitution from $j = J, \dots, 1$, using that $\tilde{u}_J = c_J$ gives

$$\tilde{u}_1 = \left[\sum_{j=1}^J \tilde{\beta}^{j-1} \tilde{c}_j^{\frac{1-\theta}{\gamma}} \right]^{\frac{\gamma}{1-\theta}}.$$

As for the resource constraint, write

$$\sum_{j=1}^J \tilde{y}_j \left(\frac{1}{1+r} \right)^{j-1} - \sum_{j=1}^J \tilde{c}_j \left(\frac{1}{1+r} \right)^{j-1} \geq 0$$

where, in slight abuse of notation, we use \tilde{y}_j to denote labor income during the working period and retirement income thereafter (see main text).

The Lagrangian writes as

$$\mathcal{L} = \left[\sum_{j=1}^J \tilde{\beta}^{j-1} \tilde{c}_j^{\frac{1-\theta}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} + \lambda \left(\tilde{a}_1 + \sum_{j=1}^J \tilde{y}_j \left(\frac{1}{1+r} \right)^{j-1} - \sum_{j=1}^J \tilde{c}_j \left(\frac{1}{1+r} \right)^{j-1} \right).$$

First-order conditions give:

$$\tilde{\beta}^{j-1} \frac{1-\theta}{\gamma} \tilde{c}_j^{\frac{1-\theta-\gamma}{\gamma}} - \tilde{\lambda} \left(\frac{1}{1+r} \right)^{j-1} = 0$$

where $\tilde{\lambda} \equiv \lambda \left(\frac{\gamma}{1-\theta} \left[\sum_{j=1}^J \tilde{\beta}^{j-1} \tilde{c}_j^{\frac{1-\theta}{\gamma}} \right]^{\frac{\gamma}{1-\theta}-1} \right)^{-1}$. Using the FOC for any two ages j and $j+1$

gives the standard Euler equation

$$\frac{\tilde{c}_{j+1}}{\tilde{c}_j} = (\tilde{\beta}(1+r))^{\frac{\gamma}{\theta+\gamma-1}} = (\tilde{\beta}(1+r))^\psi.$$

where ψ is the *IES*. We consequently have

$$\frac{\tilde{c}_j}{\tilde{c}_1} = (\beta(1+r))^{\psi(j-1)}.$$

Using this in the resource constraint, which holds with equality in the optimum, and defining by \widetilde{pvi}_1 total (human) wealth of the household, we get

$$\begin{aligned} \widetilde{pvi}_1 &\equiv \sum_{j=1}^J \tilde{y}_j \left(\frac{1}{1+r}\right)^{j-1} = \tilde{c}_1 \sum_{j=1}^J \frac{\tilde{c}_j}{\tilde{c}_1} \left(\frac{1}{1+r}\right)^{j-1} \\ \Leftrightarrow \quad \widetilde{pvi}_1 &= \tilde{c}_1 \sum_{j=1}^J \left((\tilde{\beta}(1+r))^\psi \left(\frac{1}{1+r}\right) \right)^{j-1} = \tilde{c}_1 \sum_{j=1}^J b^{j-1} = \frac{1}{m_1} \tilde{c}_1 \end{aligned}$$

where $b \equiv (\tilde{\beta}(1+r))^\psi \left(\frac{1}{1+r}\right)$ and $m_1 \equiv (\sum_{j=1}^J b^{j-1})^{-1}$ is the marginal propensity to consume out of initial wealth in period 1. We accordingly get, for any age j , that

$$\tilde{c}_j = m_j \widetilde{pvi}_1, \text{ where } m_j \equiv \tilde{\beta}(1+r)^{\psi(j-1)} m_1.$$

Using this in the utility function we get

$$\tilde{u}_1 = \left[\sum_{j=1}^J \beta^{j-1} \left(m_j \widetilde{pvi}_1 \right)^{\frac{1-\theta}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} = \left[\sum_{j=1}^J \beta^{j-1} (m_j)^{\frac{1-\theta}{\gamma}} \right]^{\frac{\gamma}{1-\theta}} \widetilde{pvi}_1,$$

establishing linearity of the utility function in initial wealth.

Consequently, the CEV in partial equilibrium—where m_j does not change between any two policies A and B because it is only a function of the constant parameters r, β, ψ —is equal to the percentage change in wealth and given by

$$g_c = \frac{\tilde{u}_1^A}{\tilde{u}_1^B} - 1 = \frac{\widetilde{pvi}_1^A}{\widetilde{pvi}_1^B} - 1.$$

□

Derivation of equations (18a-18b). Here, we derive the stock and bond return in the

quantitative model. Recall that $\bar{\kappa}_f$ is the exogenous and constant debt-equity ratio. First, we restate eq. (17),

$$K(z^t) = S(z^t) + B(z^t) = S(z^t)(1 + \bar{\kappa}_f), \quad (33)$$

S and B denote the quantities of stock and bond, respectively. The return on capital then satisfies

$$r(z^t)K(z^t) = r(z^t)S(z^t)(1 + \bar{\kappa}_f).$$

The return on capital equals the standard first-order condition of the firm, as shown in equation (16b). Out of this total return on capital, bondholders receive

$$r_b(z^{t-1})B(z^t) = r_b(z^{t-1})\bar{\kappa}_f S(z^t),$$

where the bond return is determined one period ahead, since it is one-period risk-free. Stock holders receive the remainder, which is

$$r_s(z^t)S(z^t) = r(z^t)S(z^t)(1 + \bar{\kappa}_f) - r_b(z^{t-1})\bar{\kappa}_f S(z^t).$$

From the last equation, we immediately get (18b). Taking expectations and rearranging yields (18a). \square