A Study of Generational Risk in Life-Cycle Models*

Jasmina Hasanhodzic† **and Laurence J. Kotlikoff**‡

June 19, 2019

Abstract

We study an 80-period Overlapping Generations model with aggregate shocks to assess the impact of Social Security and financial markets on the size of generational risk. Generational risk is measured as differences across age cohorts in realized lifetime utility, and as the scope for contemporaneous generations to share risk. We find that Social Security is effective in reducing generational risk. The one-period bond market can be more effective than Social Security in sharing risks among contemporaneous generations, but accentuates generational risk across age cohorts. To solve the model, we develop a new algorithm, building on Judd, Maliar, and Maliar (2011).

Keywords: Intergenerational Risk Sharing; Government Transfer Policies; Aggregate Shocks; Incomplete Markets; Stochastic Simulation; Equity Premium.

*This work builds on but is different from the previous unpublished work by the authors (Hasanhodzic and Kotlikoff 2013, 2017). Unlike the latter work, this paper does not attempt to quantify generational risk in the economy, but instead focuses on the effect of policies and financial markets on that risk. We thank Alan Auerbach, Rick Evans, Simon Gilchrist, Kenneth Judd, Dirk Krueger, Felix Kubler, Alisdair McKay, Kerk Phillips, Michael Reiter, Thomas Sargent, Kent Smetters, and Karl Schmedders for very helpful comments.

†Babson College Finance Division, `jhasanhodzic1@babson.edu`.

‡Boston University Department of Economics, `kotlikoff@gmail.com`.

1 Introduction

Economists have long examined generational risk and its mitigation via government policy (e.g., Diamond, 1977; Merton, 1983; Bohn, 1998, 1999, 2001, 2002, 2006, 2009; Shiller, 1999; Rangel and Zeckhauser, 2001; Smetters, 2003; Krueger and Kubler, 2006; Ball and Mankiw, 2007; Campbell and Nosbusch, 2007; Storesletten, Telmer, and Yaron, 2007; Bovenberg and Uhlig, 2008; and Miyazaki, Saito, and Yamada, 2010). This paper adds to this literature by studying the impact of financial markets and Social Security on the size of generational risk. Our framework is a large-scale (80-period) life-cycle model with macro shocks. Our model features isoelastic preferences, moderate risk aversion, Cobb-Douglas production, a one-period bond market, and aggregate shocks to both total factor productivity (TFP) and the rate of capital depreciation. The TFP process is trend stationary with normal innovations in line with many macro growth models.[1] Policy is limited to Social Security.

Generational risk is measured in two ways. First, we examine the dispersion across generations born at different dates (including those whose lifespans have limited or no overlap) in their realized levels of lifetime utility. Second, we consider the scope for contemporaneous generations to share risk.

We consider two calibrations commonly considered in the literature, here labeled "base" and "alternative". The base calibration sets the variability of depreciation shocks to reproduce the variability of the economy's return to wealth. Specifically, we use data from the Bureau of Economic Analysis' National Accounts and the Federal Reserve's Financial Accounts to back out the return to the economy's wealth based on

$$r_t = \frac{W_{t+1} - W_t + C_t - E_t}{W_t}, \tag{1}$$

where W_t, E_t, C_t, and r_t stand for time-t wealth, labor income, economy-wide consumption,

[1] See, e.g., Hansen (1985), Prescott (1986), Cooley and Prescott (1995), Ríos-Rull and Santaeulalia-Llopis (2010), Gomme, Rogerson, Rupert, and Wright (2005), and Judd, Maliar, and Maliar (2011).

and the rate of return on economy-wide wealth. The wealth data come from the Federal Reserve. The labor income and consumption data are from the National Accounts.[2] Some components of the Federal Reserve wealth series are carried at book. Consequently, our r_t series may exhibit too little variability. This calibration is in line with, for example, Mehra and Prescott (1985), where the risky asset traded is the claim to the entire capital in the economy, rather than just to corporate dividends. As an alternative, we calibrate our depreciation shocks such that our model's return volatility matches that of the equity market, as, for example, in the influential paper by Krueger and Kubler (2006) (KK).[3]

Under both calibrations, stochastic depreciation breaks the positive correlation between stock returns and wages. In particular, the correlation coefficient between the wage and the rate of return on capital is -0.089 in the base model and -0.098 in the alternative model, both of which are within Davis and Willen's (2000) empirical estimates. Such different risk exposures produce significant risk-sharing opportunities between workers and retirees and lead them to share risk via the model's bond market, by lending to and borrowing from each other.[4]

Our model generates modest or significant generational risk depending on the calibration. Under the base calibration, generational risk is modest according to both measures. The average annual compensating differential (calculated as the generation-specific consumption adjustment) needed to achieve the average, across all newborns, in realized lifetime utility is less than 5 percent. The average annual compensating differential needed to achieve full risk sharing among contemporaneous generations is less than 1 percent. Under our alternative calibration, generational risk is significant. It more than quadruples compared to the base, averaging 21 percent and 4 percent under the two measures, respectively.

[2] We measure total labor income assuming labor's share of proprietorship and partnership income is the same as for national income.

[3] Other examples include Smetters (2003) and Harenberg and Ludwig (2015, 2016).

[4] Younger generations are hurt more by adverse TFP shocks than older generations whose relatively large holdings of capital are not impacted. And older generations are hurt more by adverse capital depreciation shocks than are younger generations who hold limited amounts of capital.

Our main finding is that regardless of the initial size of generational risk, Social Security is effective in reducing it. This is the case in all models when generational risk is measured as dispersion in realized lifetime utility across newborns born at different points in time, although the impact is modest in the base model and significant in the alternative model. Social Security is also effective in sharing risk across contemporaneous generations, except when this risk is already very small to start with, in which case it can slightly exacerbate it.

We also find that an unfettered bond market is highly effective in sharing risk among overlapping generations, indeed more effective than Social Security. However, the absolute sizes of the generation-specific short and long positions generated by the model are unrealistically large, rendering such investments an unfeasible alternative to Social Security. Interestingly, although the bond market helps contemporaneous generations share risk, its presence accentuates differences in realized lifetime utility across newborn.

Our base calibration features a very small risk premium. However, adding one ingredient to the model, namely costs of borrowing that rise with the amount borrowed, leads to a realistic risk premium and Sharpe ratio. Such borrowing constraints are called "soft" because they place no absolute limit on the extent of borrowing. Constantinides, Donaldson, and Mehra (2002) were the first to use borrowing constraints, albeit, hard ones, in a life-cycle model to produce reasonably-sized risk premiums.[5] Adding soft borrowing constraints has little impact on the economy's key macro variables and does not change our conclusions with respect to the base model without such costs. Our alternative calibration produces a realistic equity premium—close to 5 percent—with no need for borrowing costs.

The above-mentioned landmark paper by Krueger and Kubler (2006) (KK) is the first to study intergenerational risk sharing role of Social Security in a quantitative OLG model. We follow them in a number of modeling choices. First, like them, we specify aggregate

[5] Our implementation borrows from Hasanhodzic (2014) who implements Constantinides et al.'s idea via soft borrowing constraints in a model like ours. Soft borrowing constraints are also featured in the theoretical models of Altonji and Siow (1987), Milde and Riley (1988), Wirjanto (1995), Fernandez-Corugedo (2002), and Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007).

risk as trend stationary, with shocks to both total factor productivity and rate of depreciation. Second, like their, our production function is Cobb-Douglas in capital and exogenously supplied labor. Finally, like them, we abstract from richer economic structures (such as, for example, capital adjustment costs, intragenerational heterogeneity, housing, and cohort-specific or demographic shocks) to isolate the effect of aggregate risk on intergenerational risk sharing.[6] However, while KK posit 9 periods we have 80. The extra number of periods appears important for the assessment of the relative roles of financial markets and Social Security in sharing generational risk. One reason is that more periods permits better averaging of shocks over the life cycle. Another reason is that additional periods also provide agents with far more opportunities to adjust their behavior through time and to self insure using financial markets.

KK use their model to ask whether Social Security's ability to mitigate generational risk (first pointed out by Merton, 1983) can effect a Pareto improvement even if the system is unfunded and crowds out capital. Their general conclusion, which is no, tells us that Social Security's risk mitigation is too small to overcome its crowding out effects in general equilibrium. But it leaves open the question of just how effective Social Security is in mitigating generational risk. Answering that question requires an analysis that considers direct measurement of generational risk. While we find that Social Security is an effective risk-mitigating tool, it being ineffective would also be consistent with KK's finding. Moreover, their finding might be interpreted as saying that generational risk is small to start with, leaving little scope for Social Security to enact an improvement. We find the opposite of that interpretation, as under the calibration of depreciation shocks corresponding to KK (our alternative calibration), generational risk is significant. This finding is not obvious, as it depends on the calibration.

Ríos-Rull (1994, 1996) studies large-scale OLG models subject to aggregate productivity

[6]In a related paper (Hasanhodzic and Kotlikoff, forthcoming at the *Journal of Money, Credit and Banking*), we use the same model to to price safe and risky government obligations using consumption-asset pricing.

shocks calibrated to U.S. aggregate wealth as in our base model. His papers ask, in part, whether the degree of completeness in risk-sharing arrangements materially affects aggregate variables in the economy. His answer is no. He also finds a very small risk premium, which he interprets as evidence of little macro risk. Our findings are similar in some ways and different in others. First, we find a small risk premium in our base simulation (absent borrowing costs), but not under our alternative calibration, which produces a realistic premium. Second, like Ríos-Rull, we find that fluctuations in macro variables are similar regardless of the presence of a bond market or Social Security. But, third, under both of our calibrations, the bond market materially reduces an aggregate risk, namely generational risk. Social Security also materially reduces aggregate risk. Hence, we find that risk-sharing arrangements matter, at least for risk sharing.

We proceed below by detailing our model, specifying its calibration, presenting results, and concluding. For this study we develop a new computation method, building on Judd, Maliar, and Maliar (2011), which we detail in the Appendix.

2 The Model

Our model is intentionally bare bones to maximize the potential for generational risk. Thus, we omit variable labor supply, which helps cohorts self insure against macro shocks. We also omit progressive income taxation, which redistributes, in part, from winning to losing generations. And we ignore all social insurance programs other than Social Security.

Each agent works through retirement age R, dies at age G, and maximizes expected lifetime utility. There are no short sale constraints, but in some versions of the model agents face borrowing costs which are increasing at the margin in the amount borrowed. There are also no adjustment costs, so firms maximize static profits.

2.1 Endowments and Preferences

The economy is populated by G overlapping generations that live from age 1 to age G. All agents within a generation are identical and are referenced by their age g at time t. Each cohort supplies ℓ_g units of labor per period, which equals 1 before and 0 after retirement. Hence, total labor supply equals R. Utility in a given year is time-separable and isoelastic, with risk aversion coefficient γ. Thus,

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}. \tag{2}$$

2.2 Technology

Production is Cobb-Douglas with output Y_t given by

$$Y_t = z_t K_t^{\alpha} L_t^{1-\alpha}, \tag{3}$$

where z is total factor productivity, α is capital's share, K_t is capital, and L_t is labor demand. Note that we abstract from capital adjustment costs. As shown in Hasanhodzic (2015), adding such costs make very little difference to the model's outcomes. Equilibrium factor prices satisfy

$$w_t = z(1-\alpha)\left(\frac{K_t}{R}\right)^{\alpha}, \tag{4}$$

$$r_t = z\alpha\left(\frac{K_t}{R}\right)^{\alpha-1} - \delta_t, \tag{5}$$

with depreciation shock, $\delta_t \sim \mathcal{N}(0, \psi^2)$.

Total factor productivity, z, obeys

$$\ln(z_{t+1}) = \rho \ln(z_t) + \epsilon_{t+1}, \tag{6}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$.

2.3 Financial Markets

Households save and invest in either risky capital or one-period safe bonds. Investing 1 unit of consumption in bonds at time t yields $1 + \bar{r}_t$ units of the model's single good in period $t+1$. The return, $\bar{r}_t$, is indexed by t because it is determined at t. The asset demand of a household age g at time t is given by $\theta_{g,t}$ and its share of assets invested in bonds is given by $\alpha_{g,t}$. The supply of capital in period t, K_t, satisfies

$$K_t = \sum_{g=1}^{G} \theta_{g,t-1}. \tag{7}$$

Bonds are in zero net supply, hence for all t,

$$\sum_{g=1}^{G} \alpha_{g,t}\theta_{g,t} = 0. \tag{8}$$

As shown in Green and Kotlikoff (2008), fiscal policy can be labeled in an infinite number of ways to produce whatever time path of explicit and implicit debts the government wishes to report. Such relabeling makes no difference to this or any other neoclassical model.[7] Hence, our model can be viewed as including government debt or not depending on the reader's preferences. With government debt included in the policy's labeling, the left-hand-side of (8) would be larger by the amount of debt. But the right-hand-side would also be larger by exactly the same amount, leaving the capital stock unchanged.[8]

How can a one-period bond market among the living impact generational risk? It obviously can't be used to share risk between the living and the unborn. But it can help contemporaneous generations share risks. For example, workers can hedge negative TFP

[7] I.e., all relabeled models are isomorphisms.

[8] The ability to relabel a given model with a given generational redistribution policy so that it has whatever time-path of government debt one wishes to report does not imply that changes in generational policy have no impact. It simply implies that the size of government debt does not help measure this or any other policy.

shocks by buying bonds from retirees who can use the proceeds to buy stock. Retirees are in a position to sell bonds to workers because part of their resources, namely the principal of their assets, is insulated from TFP shocks. This is particularly the case for the oldest elderly who have the fewest years left to live and whose consumption is disproportionately determined by their stock of assets as opposed to the return on their assets.

Depreciation shocks reverse this logic, but to a lesser degree than one might first think. A large negative depreciation shock directly hurts retirees, who are the primary owners of capital. But, thanks to the induced shortage of capital, it also helps them by raising the rate of return over the short term. As for workers, the reduction in the stock of capital leads to a short-term reduction in their wage.

2.4 Borrowing Costs

Following Hasanhodzic (2014), we model borrowing costs via a function proposed by Chen and Mangasarian (1996). The function is smooth and rising for negative bond holdings and essentially zero when bond holdings are close to zero or positive. Specifically, to borrow the amount of $\alpha\theta$ households have to pay the borrowing cost of $f(\alpha)\theta$, where

$$f(\alpha) = 0.2\left(-b\alpha - 1 + \frac{1}{5}\ln(1 + \mathrm{e}^{5b\alpha+5})\right) \tag{9}$$

and b is a parameter described in Section 3 governing the slope of f. Since f is increasing in bond shares (α), for a given level of assets (θ) the marginal borrowing cost is increasing in total amount borrowed ($\alpha\theta$).[9] For more details about the rationale for our bond-holding penalty function and its comparison to standard hard borrowing constraints used in Con-

[9] As explained in Hasanhodzic (2014), specifying the borrowing costs via $f(\alpha)\theta$ rather than $f(\alpha\theta)$ insures that the model remains scalable. This specification also makes economic sense. With 2θ in assets and some α, the marginal costs would be the same as with θ in assets and that same alpha, since the extra assets could be used as collateral. This is in line with Goodfriend (2005) and Goodfriend and McCallum (2007), where collateral is as a valuable input in loan production because it enables a bank to enforce the repayment of loans with less monitoring (i.e., the greater is the borrower's collateral, the more productive is the intermediary's monitoring effort).

stantinides, Donaldson, and Mehra (2002), see Hasanhodzic (2014).

2.5 Government

Social Security benefits are financed by a wage tax, τ, and provided to all retirees on a per-capita basis. Let $H_{g,t}$ denote the tax levied on the age-g household at time t and let $B_{g,t}$ denote the benefit paid to the age-g household at time t. Then

$$H_{g,t} = \tau w_t \ell_g \tag{10}$$

and

$$B_{g,t} = (1 - \ell_g) \frac{\sum_{g=1}^{G} H_{g,t}}{80 - R}. \tag{11}$$

Since tax rate (τ) is fixed, equation 11 implies that the current young are paying for the Social Security benefits of the current old.[10]

2.6 Household Problem

At time t, the economy's state is $(s_t,\ z_t)$, with $s_t = (x_{1,t}, \ldots, x_{G-1,t})$ denoting the set of age-specific holdings of cash-on-hand. [11]

$$V_g(s, z) = \max_{c,\theta,\alpha} \left\{ u(c) + \beta \mathrm{E} \left[V_{g+1}(s', z') \right] \right\} \quad \text{for} \quad g < G, \text{ and} \tag{12}$$

$$V_G(s, z) = u(c) \tag{13}$$

[10] In the model without Social Security, both $H_{g,t}$ and $B_{g,t}$ are zero for all g and all t.

[11] Note that $x_{G,t}$, the cash on hand of the oldest generation is not included in the state vector. When the depreciation shock, δ, is zero, the value of $x_{G,t}$ can be inferred from the other state variables. When δ is random, this is no longer the case. Now the initial value of $x_{G,t}$ (or equivalently the initial value of δ) is needed to fully characterize the economy's initial-period consumption vector. But we still exclude $x_{G,t}$ from the state vector because it provides no addition information about the economy's future evolution. Also, we can directly calculate $x_{G,t}$ and, thus, the consumption of the old for periods beyond the first.

subject to

$$c_{1,t} = \ell_1 w_t - \theta_{1,t} - H_{1,t} + B_{1,t}, \tag{14}$$

$$\begin{aligned} c_{g,t} &= \ell_g w_t + \left[\alpha_{g-1,t-1}(1+\bar{r}_{t-1}) + (1-\alpha_{g-1,t-1})(1+r_t)\right]\theta_{g-1,t-1} - \theta_{g,t} \\ &\quad - If(\alpha_{g-1,t-1}\theta_{g-1,t-1}) - H_{g,t} + B_{g,t}, \\ &\quad \text{for } \ 1 < g < G, \ \text{ and} \end{aligned} \tag{15}$$

$$\begin{aligned} c_{G,t} &= \ell_G w_t + \left[\alpha_{G-1,t-1}(1+\bar{r}_{t-1}) + (1-\alpha_{G-1,t-1})(1+r_t)\right]\theta_{G-1,t-1} \\ &\quad - If(\alpha_{G-1,t-1}\theta_{G-1,t-1}) - H_{G,t} + B_{G,t}, \end{aligned} \tag{16}$$

where $c_{g,t}$ is the consumption of a g-year old at time t, I is an indicator variable that equals 1 if there are borrowing costs and equals 0 otherwise, f is the borrowing cost function described above, and (14)–(16) are budget constraints for age group 1, those between 1 and G, and that for age group G.

2.7 Equilibrium

Given the initial state of the economy $(x_{1,0}, \ldots, x_{G-1,0}, z_0)$, the recursive competitive equilibrium is defined as follows.

Definition. The recursive competitive equilibrium is governed by the consumption functions, $c_g(s,z)$, the share of savings invested in bonds, $\alpha_g(s,z)$, factor demands of the representative firm, $K(s,z)$ and $L(s,z)$, government policy, $H(s,z)$ and $B(s,z)$, as well as the pricing functions $r(s,z)$, $w(s,z)$, and $\bar{r}(s,z)$ such that:

1. Given the pricing functions, the value functions (12) and (13) solve the recursive problem of the households subject to the budget constraints (14)–(16), and θ_g, α_g, and c_g are the associated policy functions for all g and for all dates and states.

2. Wages and rates of return on capital satisfy (4) and (5).

3. The government budget constraint (11) is satisfied.

4. All markets clear.

5. Finally, for all age groups $g = 1, \ldots, G-1$, optimal intertemporal consumption and investment choice satisfies

$$1 = \beta E_z\big[\big(1 + r(s', z')\big)\frac{u'(c_{g+1}(s', z'))}{u'(c_g(s, z))}\big] \qquad (17)$$
$$+ I\beta E_z\big[\big(\alpha_g(s,z)(\bar{r}(s,z) - r(s', z')) - f(\alpha_g(s,z))\big)\frac{u'(c_{g+1}(s', z'))}{u'(c_g(s, z))}\big]$$

and

$$0 = E_z\left[u'(c_{g+1}(s', z'))\big(\bar{r}(s,z) - r(s', z') - If'(\alpha_g(s,z))\big)\right], \qquad (18)$$

where E_z is the conditional expectation of z' given z.

3 Calibration

As indicated below, our calibration is standard. We assume neither demographic nor systematic technological change since our focus is on risk, not long-run trends.

3.1 Endowments and Preferences

Agents work for 45 periods and live for 80. We set the quarterly subjective discount factor, β, to 0.99. This implies an annual value of 0.96 for β. In the base model as well as the model with risk-inducing policy, risk aversion γ equals 2. In the alternative model (i.e., the model with depreciation shocks calibrated as in Krueger and Kubler, 2006) it equals 5. With this level of risk aversion, the alternative model delivers a realistic risk premium absent borrowing costs. But, as indicated, it comes at cost, namely an overly volatile output series.

3.2 Technology

Base Model

For the base model, we calibrate the TFP process, z, based on Hansen (1985) and Prescott (1986).[12] Hansen estimates a quarterly value for the autocorrelation coefficient, ρ, of 0.95 and a standard deviation, σ, of the innovation ϵ ranging from 0.007 to 0.01. Prescott's (1986) estimates are 0.9 for ρ and 0.00763 for σ.

Our assumed quarterly values for ρ and σ are 0.95 and 0.01, respectively. On an annual basis they are 0.814 and 0.019, respectively, generating a mean TFP value of 0.997 with a standard deviation of 0.033 and a coefficient of variation of 3.293 percent.

In our base model, we set the quarterly value of the standard deviation, ψ, of the depreciation shock, δ, to 0.011 (implying an annual value of 0.045).[13] This is higher than the 0.0052 quarterly estimate of Ambler and Paquet (1994). With this calibration of the shocks, the wage displays a standard deviation of 0.117 around a mean of 1.934, for a coefficient of variation of 6.047 percent.

The Alternative Model

In the alternative model, the TFP process is calibrated as above, but the quarterly standard deviation of the depreciation shock is increased to 0.034 (implying an annual value of 0.137) to reproduce the Sharpe ratio of the stock market.[14] This accords with the size of the standard deviation of the depreciation shock assumed by Krueger and Kubler (2006).

Empirical estimates of the historic equity premium and the standard deviation of the return on stocks—and therefore the Sharpe ratio—vary greatly depending on the time period used and the security chosen to proxy for the safe asset.

[12] This TFP formuation is standard. See, e.g., Cooley and Prescott (1995), Ríos-Rull and Santaeulalia-Llopis (2010), Gomme, Rogerson, Rupert, and Wright (2005), and Judd, Maliar, and Maliar (2011).

[13] We interpret Y (equation 3) as the net production function, and hence set the mean value of depreciation at zero.

[14] The model's Sharpe ratio is defined as the difference in mean real returns to capital and safe bonds divided by the standard deviation of the real return to capital.

The equity premium often targeted in academic studies is 4 percent (see, e.g., Jagannathan, McGrattan, and Scherbina, 2001). This accords with Siegal's (1998) estimate based on data for the last two centuries. Mehra (2008) suggests that the historic equity premium ranges from 2 to 4 percent if an inflation-indexed, default-free bond portfolio is used as a proxy for the risk-free rate. Jagannathan, McGrattan, and Scherbina (2001) find that the equity premium has declined over time, averaging just 0.7 percent after 1970. As for the standard deviation of stock returns, Constantinides, Donaldson, and Mehra (2002) report a range of 13.9 to 15.8 percent.[15] Combining a 4 percent equity premium and a 14 percent standard deviation of the real equity return implies an empirical Sharpe ratio of 0.286.

Our model's alternative calibration generates a risk premium of 4.634 percent and a standard deviation of the return to capital of 13.922 percent, producing a Sharpe ratio of 0.333. These values are in line with the historical record of measured equity returns. They also accord with the practitioner literature (see, e.g., Hasanhodzic, Lo, and Patel, 2009).

3.3 Borrowing Costs

The borrowing costs function is calibrated so that the ratio of the marginal borrowing cost to the risk-free rate ranges from 15 to 20. Many credit card companies now charge interest rates between 15 and 25 percent. Given that the yield on 30-year TIPS is 1 percent and even lower on shorter-term TIPS, our borrowing cost ratio seems reasonable. Moreover, even those borrowers with high incomes, excellent credit, and considerable home equity, face home equity rates that can be 10 or more times higher than the one-month T-Bill rate. To achieve this ratio, the borrowing cost parameter b is set to 25 in the model without Social Security and to 33 in the model with it.

[15] This is in line with the standard deviation of the annualized returns of the S&P500 Total Return Index over the last 22 years.

3.4 Government

Social Security benefits are financed via a payroll tax, τ, of 15 percent.

3.5 Return to U.S. Wealth and to Safe Assets

As indicated in the Introduction, to measure the empirical equivalent to the model's return on capital we use the national income accounting identity that $W_{t+1} = W_t + r_t W_t + E_t - C_t - G_t$, where W_t stands for national wealth at time t , E_t stands for labor income at time t, C_t standards for household consumption at time t, and G_t stands for government consumption at time t. We solve this identity for annual values of r_t by plugging in values of W_t, reported in the Federal Reserve's Financial Accounts data, and E_t, C_t, and G_t , reported by the Bureau of Economic Analysis in the National Income Accounts. Our data for this calculation cover 1947-2015. All data were converted into real dollars using the PCE index and measured at producer prices. The share of labor earnings in proprietorship and partnership income was assumed to equal the overall share of labor income to national income on a year-by-year basis. The empirical counterpart to the model's safe rate of return is calculated as the annualized real return on one-month Treasuries from 1947-2015.

As reported in Table 1, the mean return to capital (actually, total U.S. assets) is 6.512 percent with a standard deviation of 4.886. The mean safe return is 1.083. These figures imply a Sharpe ratio of 1.111, which, as indicated below, is close to our base model's Sharpe ratio.

4 Results

Turning to the results, for each calibration (base and alternative), we consider three models: the model with the bond market, the model with the bond market, but with borrowing costs, and the model without the bond market. We first compare the models' asset returns with those in the data. We then turn to assessing the effect of Social Security and the bond market on generational risk. To quantify the ability of Social Security to mitigate generational risk, for each model we present results both with and without Social Security. To quantify the impact of the bond market on generational risk, we compare our measures of generational risk in models with the bond market to those in models without the bond market. In each case, our generational risk measures focus on risk in an environment where policy has been in place for at least 75 years, i.e. they are calculated post 75 years of transition.

4.1 Summary Statistics of Asset Returns

Table 1 compares returns to capital and the safe bond in our model and in the data. The base model, regardless of whether it includes borrowing costs or Social Security, reproduces the variability of the real return to U.S. wealth, as intended by the calibration. Specifically, the standard deviation of the return to capital ranges from 4.587 percent to 4.606 percent in the model compared to 4.886 percent in the data.

Table 1 also shows that, absent borrowing costs, our base model also produces a risk-premium puzzle. The rate of return to capital averages 3.916 percent without Social Security and 4.892 percent with it, while the corresponding safe rate of return averages 3.632 percent and 4.622 percent in the two cases, respectively. The resulting risk premia are very small—0.285 percent without Social Security and 0.270 percent with it.

However, adding increasing costs to borrowing to the base model reduces the risk-free rate to roughly two-tenths of one percent on average. This implies a risk premium of 3.762 percent in the model without Social Security and 4.772 percent with it. The latter is close to the 5.429 percent measured in the data.

Table 1 also reveals that the alternative model overstates the standard deviation of the returns to capital by almost a factor of three (13.922 percent in the model vs. 4.886 percent in the data). However, the calibration of this model was intended to be in line with the historical record of measured equity returns rather than the return to U.S. wealth. Indeed, as specified in Section 3.2 above, this model's risk premium and Sharpe ratio (4.634 percent and 0.333, respectively) accord with those of the stock market.

Summary Statistics of Rates of Return

	Mean Return to Capital (%)	Mean Safe Rate of Return (%)	S.D. of Return to Capital (%)	Risk Premium (%)	Sharpe Ratio
	Base Model				
No Social Security	3.916	3.632	4.592	0.285	0.062
Social Security	4.892	4.622	4.606	0.270	0.059
	Base Model With Borrowing Costs				
No Social Security	3.999	0.237	4.587	3.762	0.820
Social Security	4.994	0.222	4.599	4.772	1.037
	Alternative Model				
No Social Security	2.640	-1.994	13.922	4.634	0.333
Social Security	4.414	0.241	14.032	4.173	0.297
	Data Based On Return to U.S. Wealth				
	6.512	1.083	4.886	5.429	1.111

Table 1: Summary statistics of rates of return in the model and in the data.

Our models' mean return to capital ranges from 2.640 percent to 4.994, depending, crucially on the presence of Social Security. This falls short of the 6.512 percent mean return to U.S. wealth (the benchmark for the base model) or the 10.900 percent mean return on a value-weighted market portfolio (the benchmark for the alternative model).[16] But we could easily introduce additional government policies, such as government consumption financed by an income tax, which would crowd out capital and raise capital's return to match that of the data.[17]

Table 2 compares the variability of output for different variants of our model to that in the data. Following Prescott (1986), we detrend real net national product for the years 1929 through 2015 and form the standard deviation of the percent deviations from trend. Our model abstracts from growth, so we simply form the standard deviation of our model's percentage deviation of annual output from its mean.[18]

[16]The value of 10.900 percent is the annualized monthly average of the value-weighted market portfolio from the WRDS CRSP database.

[17]We chose not to do so to isolate the impact of Social Security on generational risk.

[18]The standard deviation of detrended per capita output is virtually identical.

Standard Deviation of Percent Deviations of Output from Trend

Model/Data	S.D. (%)
Base Model	
No Social Security	5.782
Social Security	5.685
Base Model With Borrowing Costs	
No Social Security	5.275
Social Security	5.237
Alternative Model	
No Social Security	22.262
Social Security	21.516
Data	
Real Net National Product, 1929-2015	3.396

Table 2: Standard deviation of percent deviations from trend of U.S. real Net National Product, 1929–2015 and standard deviation of percent deviations of output from its mean in the base and alternative models.

As the table shows, regardless of the calibration, our models overstate actual output variability. For example, in the base model with Social Security and borrowing costs, the standard deviation of percentage output deviations is 5.237 percent in our model compared with 3.396 percent in the data. The alternative model produces a standard deviation of 21.516 percent with Social Security and 22.262 percent without it.

4.2 Generational Risk in the Base Model

In what follows we describe each generational risk measure and present corresponding results for the base model.

4.2.1 Differences Across Age Cohorts in Realized Lifetime Utility

In Table 3 we report the realized lifetime utility generational risk measure. The realized utility measure is based on the particular state to which the generation is born and the particular sequence of shocks it experiences over its lifetime. We first calculate each generation's particular realized lifetime utility and form the average of these realized values across all generations born between years 75 and 751. Next, we calculate for each generation the factor by which we need to multiply each year's realized consumption to produce the same realized lifetime utility as the first 677 generations experience on average. Finally, we compute the absolute value of this factor's deviation from 1. The closer are the percent adjustments to 0, and the less variable they are through time, the less difference does the date of birth make for the household's expected lifetime utility, i.e., the smaller is the generational risk.

Absolute Percentage Adjustments to Achieve Average Realized Utility of Newborns

	Min	Mean	Max	S.D.
Bond Market				
No Social Security	0.014	4.675	19.961	3.800
Social Security	0.000	4.107	18.097	3.388
Borrowing Costs				
No Social Security	0.006	2.663	10.333	2.082
Social Security	0.001	2.372	10.639	1.996
No Bond Market				
No Social Security	0.000	2.686	10.329	2.096
Social Security	0.007	2.394	10.639	2.006

Table 3: Means, standard deviations, minimums, and maximums of absolute percent adjustments in each cohort's annual consumption needed to achieve post-transition average realized lifetime utility of newborns in the base model.

This table shows that the compensating consumption differential needed to equate realized lifetime utility of each cohort through time to the average across newborns of their realized lifetimes utilities averages 4.675 percent with a standard deviation of 3.800 percent in the presence of a fully functional bond market. These values are modest. The maximum value of the differential is 19.961 percent, which is large, but the bond market is not designed to share risk among the living and the unborn. Rather, it may exacerbate differences in realized lifetime utility across non-overlapping generations or generations with limited overlap. Indeed, eliminating the bond market lowers the average value of the differential to 2.686, its standard deviation to 2.096, and its maximum value to 10.329.

Adding Social Security, whether in the presence or absence of bonds, helps, despite the fact that the residual risk needing to be shared is small already. For example, with no bond market, Social Security reduces this risk measure by 10.9 percent—from 2.686 percent to 2.394 percent.

Why might shocks hitting earlier generations not greatly impact later generations? All generations live for 80 periods meaning that initial good or bad shocks will have limited impact on the shocks that a generation experiences later in life. Intuitively, new shocks, even if their TFP impact persists, arrive each period and average out over time. Moreover, our economy is ergodic and naturally rebounds from bad or good states. Hence, a generation that is born into bad (good) times with a small (large) stock of capital will, on average, experience better (worse) times in the future.

4.2.2 Generational Risk Among Contemporaneous Generations

We next ask whether contemporaneous generations are experiencing materially different shocks as measured by differences in their annual percentage consumption changes. If so, such changes could be pooled either via private arrangements or government policy. Recall that full risk sharing among contemporaneous generations, indeed all agents, with the homothetic preferences considered here, requires equal percentage changes in the consumption from one period to the next (see Abel and Kotlikoff, 1988). Hence, one can measure the extent of generational risk by considering the co-movement of consumption across age groups as well as the extent of consumption adjustments that would be needed to achieve perfect consumption co-movement.

Table 4 summarizes the agent- and year-specific absolute percentage consumption adjustment needed to achieve perfect risk sharing, i.e., to ensure that all agents experience the same percentage change in the year in question.[19] It shows that the average (across agents) absolute percentage adjustment needed to achieve full risk sharing is 0.224 percent

[19] Values of 0.000 and 1.000 reflect rounding.

Absolute Percentage Adjustments to Achieve Perfect Risk Sharing Among Contemporaneous Generations

	Min	Mean	Max	S.D.
No Borrowing Costs				
No Social Security	0.000	0.224	1.098	0.163
Social Security	0.000	0.317	1.108	0.152
Borrowing Costs				
No Social Security	0.000	0.933	5.904	0.699
Social Security	0.000	0.713	3.811	0.535
No Bond Market				
No Social Security	0.000	0.929	5.853	0.696
Social Security	0.000	0.721	3.877	0.539

Table 4: Absolute percent adjustments to achieve perfect risk sharing among contemporaneous generations in the base model. Minimum, mean, and maximum are taken across all cohorts and all time. The standard deviation value is the mean across cohorts of cohort-specific standard deviations.

without Social Security, but with an unfettered bond market.[20] The standard deviation is only 0.163 percent. Hence, generational risk among contemporary cohorts is quite small even absent any government risk-sharing policy. Indeed, the maximum absolute adjustment is only 1.098 percent. Adding Social Security to the risk-sharing mix slightly exacerbates generational risk, raising the average adjustment needed to achieve perfect risk sharing to 0.317 percent.[21]

With borrowing costs and no Social Security, the absolute adjustments are larger, but still quite small—0.933 percent on average. The maximum adjustment in this case—5.904 percent—is almost six times larger than arises without borrowing costs. In this case, adding Social Security improves risk sharing. It reduces the average absolute adjustment needed for full risk sharing by one quarter. It also lowers the standard deviation of the adjustment

[20] Again, if all contemporaneous generations experienced the same shocks, appropriately scaled, everyone's consumption would, on a percentage basis, rise or fall in unison. Table 4 indicates that very little redistribution among contemporaneous generations is needed to assure that outcome.

[21] Recall that we are simulating pay-go Social Security with a fixed tax rate. Clearly, if the share of elderly were very small, TFP shocks could produce major changes in tax revenues, which would then be spread over a small number of beneficiaries. In this case, Social Security could produce more risk across contemporaneous generations. Hence, there is no reason to expect pay-go Social Security to always share risk appropriately among those alive at a point in time.

and its maximum value. Social Security has a similar effect in the model without the bond market.

Note that the unfettered bond market is more than three times as efficient as Social Security in reducing risk among contemporaneous generations. Compare, in this regard, the value of 0.224 percent with the value of 0.721 percent. The former number is the mean absolute percent adjustment with the bond market, but no Social Security. The later value is the mean adjustment with Social Security, but no bond market. However, as illustrated below, the simulated size of short and long bond positions associated with unrestricted use of this market appears unrealistically large.

4.2.3 Illustrating Consumption Co-Movement Among Contemporaneous Generations

Illustrating the Comovement in Consumption

Age	TFP	Depreciation Rate (%)	C	% Change in C	Wage	Stock Holdings	Bond Holdings	Rate of Return on Bonds (%)	Rate of Return to Capital (%)	% Change Output	% Change in Agg C
Worker											
29	0.934	-5.960	1.659	2.396	1.805	11.029	-10.160	3.446	9.544	-1.460	2.772
30	0.946	5.160	1.605	-3.276	1.859	11.916	-10.198	3.284	-1.644	2.959	-3.169
31	0.971	6.285	1.555	-3.129	1.873	11.263	-9.823	3.462	-2.537	0.754	-3.320
32	0.984	7.453	1.494	-3.907	1.858	10.563	-9.430	3.709	-3.492	-0.799	-3.907
33	0.974	3.545	1.460	-2.258	1.796	9.784	-9.007	3.951	0.571	-3.317	-1.957
34	0.950	2.733	1.428	-2.227	1.732	9.521	-8.708	4.009	1.367	-3.544	-1.765
Retiree											
69	0.934	-5.960	1.700	3.030	0.000	20.772	9.111	3.446	9.544	-1.460	2.772
70	0.946	5.160	1.646	-3.228	0.000	21.274	9.205	3.284	-1.644	2.959	-3.169
71	0.971	6.285	1.585	-3.670	0.000	19.847	8.939	3.462	-2.537	0.754	-3.320
72	0.984	7.453	1.519	-4.162	0.000	18.362	8.645	3.709	-3.492	-0.799	-3.907
73	0.974	3.545	1.490	-1.948	0.000	16.874	8.293	3.951	0.571	-3.317	-1.957
74	0.950	2.733	1.466	-1.552	0.000	16.052	8.050	4.009	1.367	-3.544	-1.765

Table 5: An example of consumption co-movement of a worker and a retiree through time.

In Table 5, we follow a worker age 29 and a retiree age 69 for six years to illustrate the co-movement in their consumption in the base model. While we consider just six years, starting in the year 745, they are very typical of our simulation results.

As columns two and three show, the model's shocks are sizable. There is, for example, close to a 5 percent difference between the largest and smallest TFP values over the period. And the depreciation shocks range from -6.0 percent to 7.5 percent. (Note, negative depreciation is a good thing, i.e., a plus, when it comes to expanding the size of the capital stock and the return to capital.)

These shocks as well as their associated impact on capital accumulation produce significant changes in the wage. In year 750, for example, the wage is 6.78 percent smaller than two years earlier. And the rate of return to capital ranges from 9.54 percent in 745 to -3.49 percent in 746.

Interestingly, aggregate consumption and output can move in different directions over short periods in our model. For example, output rises, but consumption falls between the first and second years of our table. This reflects the higher stock of capital (thanks to negative depreciation) in 745 and the higher level of TFP. But the year 746 return to capital is negative reflecting the significant depreciation rate that year. This depresses the consumption of the old as well as the young who also hold stock (capital). The size of these shocks, the size of changes to factor returns, and the size of output changes notwithstanding, there is very high co-movement of the worker's, the retiree's consumption as well as aggregate consumption.[22]

Table 5 suggests that borrowing by the young and lending by the old are instrumental to risk sharing among the living. This raises the real world question of whether the workers borrow from the old to purchase stocks. One subtle mechanism by which they do so is via their company's borrowing on their behalf and paying them more when the company does well. These payments may be contemporaneous with company's performance or be made through time. This effectively lets workers borrow to invest in stocks.[23] However, the simulated sizes of short and long bond positions associated with unrestricted borrowing and lending appear unrealistically large. For instance, 29-year-old workers optimize their

[22] Also per capital consumption since the model features zero population growth.

[23] The fact that young borrow to purchase equity is a central feature of Constantinides, Donaldson, and Mehra (2002).

exposure to shocks by holding 11.029 units of stocks financed by shorting 10.160 units of bonds. Younger workers take even more highly leveraged equity positions. This indicates that although the bond market with unlimited leverage appears highly efficient in reducing risk among contemporaneous generations, it may not be a feasible alternative to Social Security.

4.3 Generational Risk in the Alternative Model

We next present results under each generational risk measure for the alternative model.

4.3.1 Differences Across Age Cohorts in Realized Lifetime Utility

Table 6 reports our realized lifetime utility measures in the alternative model. The measures are much larger that their Table 3 (base model) counterparts. For example, under the alternative calibration with neither bonds nor Social Security, one would need to expand or contract a newborn's consumption in all years of her life by 15.141 percent, on average, in order to achieve average realized utility across newborns. The maximum adjustment is 83.694 percent. The corresponding figures in the base model are 2.686 percent and 10.329 percent. Thus, the alternative calibration is highly effective in generating risk across newborns through time. And, as Table 6 shows, Social Security is highly effective in mitigating that risk. The same table also confirms that the presence of the bond market increases generational risk according to this measure. However, in the alternative model, the increase is much more pronounced.

4.3.2 Generational Risk Among Contemporaneous Generations

Table 7 presents our measures of generational risk among contemporaneous generations for the alternative model. As indicated, the average absolute percentage adjustments needed to achieve full risk sharing are less than half a percent without Social Security, but with bonds. Hence, in the presence of a bond market, generational risk among contemporary cohorts is

Absolute Percentage Adjustments to Achieve Average Realized Utility of Newborns
Alternative Model

	Min	Mean	Max	S.D.
Bond Market				
No Social Security	0.096	20.078	78.201	16.675
Social Security	0.008	16.895	66.167	14.229
No Bond Market				
No Social Security	0.105	15.141	83.694	13.758
Social Security	0.004	9.248	32.863	8.435

Table 6: Means, standard deviations, minimums, and maximums of absolute percent adjustments in each cohort's annual consumption needed to achieve post-transition average realized lifetime utility of newborns in the alternative model.

quite small even absent any government risk-sharing policy. Indeed, the maximum value is only 7.014 percent. Adding Social Security reduces these small values even further. Now the average absolute adjustment is 0.262 percent and the maximum adjustment is 4.172 percent.

Absolute Percentage Adjustments to Achieve Perfect Risk Sharing Among Contemporaneous Generations
Alternative Model

	Min	Mean	Max	S.D.
Bond Market				
No Social Security	0.000	0.437	4.557	0.341
Social Security	0.000	0.296	4.179	0.289
No Bond Market				
No Social Security	0.000	3.473	24.012	2.627
Social Security	0.000	2.307	17.919	1.789

Table 7: Absolute percent adjustments to achieve perfect risk sharing among contemporaneous generations in the alternative model. Minimum, mean, and maximum are taken across all cohorts and all time. The standard deviation value is the mean across cohorts of cohort-specific standard deviations.

Without bonds or Social Security, the absolute adjustments are larger—2.943 percent on average, with a much larger maximum value of 23.983 percent. Social Security reduces the average and the maximum values of the adjustment to 2.004 percent and 15.335 percent, respectively.

But, as was the case in the base model, if the bond market could be used with unrestricted leverage, it would be more effective than Social Security in reducing generational risk among the living. To see this, compare the value of 0.486 percent with the value of 2.004 percent. The former value is the mean absolute percent adjustment with the bond market, but no Social Security. The later value is the adjustment with Social Security, but no bond market.

5 Conclusion

The mitigation of generational risk via financial markets or Social Security and other government policies has long intrigued economists. This paper adds to the literature on this issue. Our 80-period OLG model with aggregate risk generates modest or large generational risk depending on the calibration. Our base calibration sets depreciation shocks to replicate the variability of returns to total U.S. wealth. Our alternative calibration reproduces the variability of returns to U.S. equities.

Our main finding is that Social Security is a significant generational risk-mitigating institution regardless of the calibration. But it's not as powerful in sharing risk across contemporaneous generations as the one-period bond market. One should, however, view our bond-market results as illustrative of the potential power of financial exchange to share risk, not as evidence that such risk sharing necessarily arises. The reason is that full financial risk sharing entails, in our model, unrealistically large short and long bond positions for particular cohorts. Interestingly, though, although the bond market helps contemporaneous generations share risk, its presence accentuates differences in realized lifetime utility across generations.

An interesting future direction would be to move away from our trend stationary environment and consider the effect on generational risk of a random walk in TFP in this model. This is likely more important for studying generational risks in countries whose TFP processes exhibit either much higher serial correlation than observed for the U.S. or breaks in their TFP growth rates.[24] Other examples include the idiosyncratic heterogeneity included in Harenberg and Ludwig (2015), the assumption of less elastic capital supply considered in Glover et al. (2016), and uncertain intergenerational policy and central bank policy, which appears to have limited the incomes of many elderly during the Great Recession. There are also other mechanisms we can use to achieve the equity premium, including Croce's (2014) combination of long-run uncertainty about the productivity growth rate, convex adjustment costs, and Epstein-Zin preferences, or Jermann's (1998) combination of habit preferences and adjustment costs.[25] While, as shown in this paper, the way the economy's risks are modeled may affect the size of generational risk, we suspect that our results on how different mechanisms (bond market and Social Security) allocate risk conditional on its size will remain intact. Ultimately, models such as ours need to be measured against generation-specific consumption data.[26]

Beyond these points, the paper demonstrates the feasibility of simulating realistic, large-scale OLG models with aggregate shocks in which generational policy matters as appears so evident in real economies.

[24]Prescott(1986) indicates that modeling the growth rate in U.S. TFP as a random walk "result in essentially the same fluctuations" as the autoregressive process, which he used and we employ. This said, generational risk, even in the U.S., might be different based on a random walk in TFP growth.

[25]Jermann (1998) found that introducing adjustment costs in the standard real business cycle model with CRRA preferences does not yield a sizeable equity premium. Hasanhodzic (2015) confirmed his finding in a model like ours.

[26]Glover et al. (2016) provide empirical evidence suggesting major intergenerational redistribution between the young and the old associated with changes in asset prices relative to labor income. Unfortunately, they don't provide measures of changes in relative consumption by age group. Moreover, half of their asset price changes appears to be due to housing price changes. But since they don't model housing services as a commodity, their empirical analysis misses the wealth effects stemming from the decline in the relative price of housing. Intuitively, households that retained their homes during the Great Recession received the same stream of housing services notwithstanding the housing price decline.

References

Abel, A. and L. J. Kotlikoff, 1988, “Does the Consumption of Different Age Groups Move Together? A New Nonparametric Test of Intergenerational Altruism”, NBER Working Papers 2490, National Bureau of Economic Research.

Altonji, J. G. and A. Siow, 1987, “Testing the Response of Consumption to Income Changes with (Noisy) Panel Data”, *Quarterly Journal of Economics* 102, 293–328.

Ambler, S. and A. Paquet, 1994, “Stochastic Depreciation and the Business Cycle”, *International Economic Review* 35, 101–116.

Auerbach, A. J., J. Gokhale, and L. J. Kotlikoff, 1994, “Generational Accounting: A Meaningful Way to Evaluate Fiscal Policy”, *Journal of Economic Perspectives* 8, 73–94.

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

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

Barro, R. J., 2006, “Rare Disasters and Asset Markets in the Twentieth Century”, *The Quarterly Journal of Economics* 121, 823–866.

Bodie, Z., R. C. Merton, and W. F. Samuelson, 1992, “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model”, *Journal of Economic Dynamics and Control* 16, 427–449.

Bohn, H., 1998, “Risk Sharing in a Stochastic Overlapping Generations Economy”, University of California at Santa Barbara, Economics Working Paper Series, Department of Economics, UC Santa Barbara.

Bohn, H., 1999, "Should the Social Security Trust Fund Hold Equities?: An Intergenerational Welfare Analysis", *Review of Economic Dynamics* 2, 666–697.

Bohn, H., 2001, "Social Security and Demographic Uncertainty: The Risk-Sharing Properties of Alternative Policies", in J. Campbell and M. Feldstein, editors, "Risk Aspects of Investment Based Social Security Reform", Chicago: University of Chicago Press, 203–241.

Bohn, H., 2002, "Retirement Savings in an Aging Society: A Case for Innovative Government Debt Management", in A. Auerbach and H. Herrman, editors, "Ageing, Financial Markets and Monetary Policy", Berlin: Springer, 139–181.

Bohn, H., 2006, "Who Bears What Risk? An Intergenerational Perspective", in D. Blitzstein, O. S. Mitchell, and S. P. Utkus, editors, "Restructuring Retirement Risks", Oxford, UK: Oxford University Press, 10–36.

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

Bovenberg, L. and H. Uhlig, 2008, "Pension Systems and the Allocation of Macroeconomic Risk", in "NBER International Seminar on Macroeconomics 2006", National Bureau of Economic Research, 241–344.

Campbell, J. Y. and Y. Nosbusch, 2007, "Intergenerational Risk Sharing and Equilibrium Asset Prices", *Journal of Monetary Economics* 54, 2251–2268.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Rios-Rull, 2007, "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default", *Econometrica* 75, 1525–1589.

Chen, C. and O. Mangasarian, 1996, "A Class of Smoothing Functions for Nonlinear and Mixed Complementarity Problems", *Computational Optimization and Applications* 5, 97–138.

Constantinides, G. M., J. B. Donaldson, and R. Mehra, 2002, "Junior Can't Borrow: A New Perspective On the Equity Premium Puzzle", *The Quarterly Journal of Economics* 117, 269–296.

Cooley, T. and E. Prescott, 1995, "Economic Growth and Business Cycles", in "Frontiers of Business Cycle Research", Princeton University Press.

Croce, M. M., 2014, "Long-run Productivity Risk: A Hope for Production-based Asset Pricing?", *Journal of Monetary Economics* 66, 13–31.

Davis, S. J. and P. Willen, 2000, "Using Financial Assets to Hedge Labor Income Risks: Estimating the Benefits", Working Paper.

Den Haan, W. J. and A. Marcet, 1990, "Solving the Stochastic Growth Model by Parameterizing Expectations", *Journal of Business and Economic Statistics* 8, 31–34.

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

Fernandez-Corugedo, E., 2002, "Soft Liquidity Constraints and Precautionary Saving", Working paper, Bank of England.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull, 2011, "Intergenerational Redistribution in the Great Recession", Technical report, National Bureau of Economic Research.

Gomme, P., R. Rogerson, P. Rupert, and R. Wright, 2005, "The Business Cycle and the Life Cycle", *NBER Macroeconomics Annual* 19.

Goodfriend, M., 2005, *Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley*, chapter Narrow Money, Broad Money, and the Transmission of Monetary Policy, Board of Governors of the Federal Reserve System.

Goodfriend, M. and B. T. McCallum, 2007, "Banking and Interest Rates in Monetary Policy Analysis: A quantitative exploration", *Journal of Monetary Economics* 54, 1480–1507.

Gourinchas, P.-O., 2000, "Precautionary Savings, Life Cycle and Macroeconomics", Econometric Society World Congress 2000 Contributed Papers 0793, Econometric Society.

Green, J. and L. J. Kotikoff, 2008, "On the General Relativity of Fiscal Language", in A. J. Auerbach and D. N. Shaviro, editors, "Institutional Foundations of Public Finance", Cambridge: Harvard University Press, 241–256.

Hansen, G. D., 1985, "Indivisible Labor and the Business Cycle", *Journal of Monetary Economics* 16, 309–327.

Harenberg, D. and A. Ludwig, 2015, "Social Security in an Analytically Tractable Overlapping Generations Model with Aggregate and Idiosyncratic Risks", *International Tax and Public Finance* 22, 579–603.

Harenberg, D. and A. Ludwig, 2016, "Idiosyncratic Risk, Aggregate Risk, and the Welfare Effects of Social Security", SAFE working paper no. 59. Available at SSRN: https://ssrn.com/abstract=2464170.

Hasanhodzic, J., 2014, "Increasing Borrowing Costs and the Equity Premium", SSRN working paper. Available at SSRN: https://ssrn.com/abstract=2425330.

Hasanhodzic, J. and L. Kotlikoff, 2013, "Generational Risk–Is It a Big Deal?: Simulating an 80-Period OLG Model with Aggregate Shocks", NBER Working Papers 19179, National Bureau of Economic Research.

Hasanhodzic, J. and L. Kotlikoff, 2017, "Generational Risk–Is It a Big Deal?: Simulating an 80-Period OLG Model with Aggregate Shocks", SSRN working paper. Available at SSRN: https://ssrn.com/abstract=2994899.

Hasanhodzic, J. and L. Kotlikoff, forthcoming, "Valuing Government Obligations When Markets are Incomplete", *Journal of Money, Credit and Banking* .

Hasanhodzic, J., A. W. Lo, and P. Patel, 2009, "The CS 130/30 Index: A Summary and Performance Comparison", White paper, Credit Suisse.

Hayashi, F., 1987, "Tests for Liquidity Constraints: A Critical Survey and Some New Observations", in F. T. Bewley, editor, "Advances in Econometrics, Fifth World Congress", Cambridge University Press, Cambridge, England.

Jagannathan, R., E. McGrattan, and A. Scherbina, 2000, "The Declining U.S. Equity Premium", *Federal Reserve Bank of Minneapolis Quarterly Review* 24, 3–19.

Jerrman, U., 1998, "Asset Pricing in Production Economies", *Journal of Monetary Economics* 41, 257–275.

Judd, K., L. Maliar, and S. Maliar, 2009, "Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models", NBER Working Papers 15296, National Bureau of Economic Research.

Judd, K., L. Maliar, and S. Maliar, 2011, "Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models", *Quantitative Economics* 2, 173–210.

Krueger, D. and F. Kubler, 2004, "Computing Equilibrium in OLG Models with Stochastic Production", *Journal of Economic Dynamics and Control* 28, 1411–1436.

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

Krusell, P. and A. A. Smith, 1997, "Income And Wealth Heterogeneity, Portfolio Choice, And Equilibrium Asset Returns", *Macroeconomic Dynamics* 1, 387–422.

Krusell, P. and A. A. Smith, 1998, "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy* 106, 867–896.

Krusell, P. and A. A. Smith, 1999, "On the Welfare Effects of Eliminating Business Cycles", *Review of Economic Dynamics* 2, 245–272.

Lucas, R., 1987, *Models of Business Cycles*, Oxford, UK: Basil Blackwell.

Malin, B. A., D. Krueger, and F. Kubler, 2011, "Solving the Multi-Country Real Business Cycle Model Using a Smolyak-Collocation Method", *Journal of Economic Dynamics and Control* 35, 229–239.

Marcet, A., 1988, "Solution of Nonlinear Models by Parameterizing Expectations", Technical report, Carnegie Mellon University.

Marcet, A. and G. Lorenzoni, 2001, "The Parameterized Expectations Approach: Some Practical Issues", in R. Marimon and A. Scott, editors, "Computational Methods for the Study of Dynamic Economies", Chicago: Oxford University Press.

Marcet, A. and D. A. Marshall, 1994, "Solving Nonlinear Rational Expectations Models By Parametrized Expectations: Convergence to Stationary Solutions", Discussion paper 91, Federal Reserve Bank of Minneapolis.

Mehra, R., 2008, "The Equity Premium Puzzle: A Review", *Foundations and Trends in Finance* 2, 1–81.

Mehra, R. and E. C. Prescott, 1985, "The Equity Premium: A Puzzle", *Journal of Monetary Economics* 15, 145–161.

Merton, R. C., 1983, "On the Role of Social Security as a Means for Efficient Risk-Bearing in an Economy Where Human Capital Is Not Tradeable", in Z. Bodie and J. Shoven, editors, "Financial Aspects of the United States Pension System", Chicago: University of Chicago Press.

Milde, H. and J. G. Riley, 1988, "Signaling in Credit Markets", *Quarterly Journal of Economics* 103, 101–129.

Miyazaki, K., M. Saito, and T. Yamada, 2010, "On the Intergenerational Sharing of Cohort-Specific Shocks on Permanent Income", *Macroeconomic Dynamics* 14, 93–118.

Prescott, E., 1986, "Theory Ahead of Business-Cycle Measurement", *Carnegie-Rochester Conference Series on Public Policy* 25, 11–44.

Rangel, A. and R. Zeckhauser, 2001, "Can Market and Voting Institutions Generate Optimal Intergenerational Risk Sharing?", in "Risk Aspects of Investment-Based Social Security Reform", NBER Chapters, National Bureau of Economic Research, 113–152.

Rawls, J., 1971, *A Theory of Justice*, Cambridge, MA: Belknap Press of Harvard University Press.

Reiter, M., 2015, "Solving OLG Models with Many Cohorts, Asset Choice and Large Shocks", SSRN working paper. Available at SSRN: https://ssrn.com/abstract=2709404.

Rietz, T. A., 1988, "The Equity Risk Premium: A Solution", *Journal of Monetary Economics* 22, 117–131.

Rios-Rull, J. and R. Santaeulalia-Llopis, 2010, "Redistributive Shocks and Productivity Shocks", *Journal of Monetary Economics* 57, 931–948.

Rios-Rull, J.-V., 1994, "On the Quantitative Importance of Market Completeness", *Journal of Monetary Economics* 34, 463–496.

Rios-Rull, J.-V., 1996, "Life-Cycle Economies and Aggregate Fluctuations", *Review of Economic Studies* 63, 465–489.

Shiller, R., 1999, "Social Security and Intergenerational, Intragenerational, and International Risk Sharing", *Carnegie-Rochester Conference Series on Public Policy* 50, 165–204.

Siegel, J. J., 1998, *Stocks for the Long Run: The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies*, McGraw-Hill.

Smetters, K. A., 2003, "Trading with the Unborn: A New Perspective on Capital Income Taxation", NBER Working Papers 9412, National Bureau of Economic Research.

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

Weitzman, M. L., 2007, "Subjective Expectations and Asset-Return Puzzles", *American Economic Review* 97, 1102–1130.

Wirjanto, T. S., 1995, "Aggregate Consumption Behaviour and Liquidity Constraints: The Canadian Evidence", *Canadian Journal of Economics* 28, 1135–1152.

A Computational Appendix

For Online Publication

OLG models with macro shocks and long-lived agents encounter the curse of dimensionality. Ríos-Rull (1994, 1996) uses local perturbation methods to solve such models. However, when shocks are large, as in our alternative calibration, a global method is preferred. Krussel and Smith's (1998) landmark work banishes the curse for certain single-agent models. They show that the state space in such models may be adequately approximated by sufficient statistics, such as the size of the economy's capital stock. Gourinchas (2000), Storesletten, Telmer, and Yaron (2007) and Harenberg and Ludwig (2016) apply the Krusell and Smith approach to OLG models. They find it works well for their purposes. But Krueger and Kubler (2004) argue that Krusell and Smith's low-dimensional approximation approach cannot, as a general matter, adequately handle OLG economies—a robustness concern raised by Krusell and Smith (1998) themselves.

Krueger and Kubler (2004, 2006) represents a major milestone in battling the curse in OLG models. They calculate solutions for life-cycle models experiencing macro shocks. They do so by applying Smolyak's (1963) algorithm to efficiently choose grid values.[27] However, their technique cannot overcome the curse at least in computing models with realistic lifespans measured in years. Indeed, Krueger and Kubler (2006) limit their model to 9 periods for computational feasibility.[28]

Our means of dispelling the curse builds on the Generalized Stochastic Simulation Algorithm (GSSA) developed by Judd, Maliar, and Maliar (2009, 2011). GSSA, in turn, builds on the Parametrized Expectations Approach (PEA) developed by Marcet (1988), Den Haan and Marcet (1990), Marcet and Marshall (1994), and Marcet and Lorenzoni (1998), notably by replacing the Monte Carlo integration method employed in these studies with a more accurate Gaussian quadrature method.

Marcet (1988) is seminal. It contains the fundamental insight that, for computational purposes, the state space can be limited to the economy's ergodic set. I.e., economic behavior needs to be calculated only for states the economy will actually visit with non-trivial probability.

As indicated in Marcet and Marshall (1992) and Judd et al. (2009, 2011), PEA and its enhancement, GSSA, have been used to solve a wide range of economic models featuring infinitely-lived agents. Our paper appears to be the first OLG model to implement as

[27] Malin, Krueger, and Kubler (2011) detail this method.

[28] Brumm and Scheidegger (2017) propose an adaptive sparse grid approach to efficiently solve high-dimensional dynamic models, although their applications do not include an OLG model.

well as build on PEA/GSSA. Unlike Judd et al.'s (2009, 2011) use of assets and shocks as state variables, the state vector here consists of cash-on-hand and shocks.[29] This renders the computation of the agents' bond holdings and the risk-free rate computationally challenging, as each agent's first order conditions depend on all agents' bond holdings. Our innovation is to develop an inner loop to deal with this challenge.

Our solution method is simple and can be easily extended to handle more complex OLG models. Here we first provide a high level overview and then a more detailed one. At a high level, we start by drawing a sequence of aggregate shocks. Second, we guess consumption functions for each of our 80 generations as linear polynomials of the economy's state vector. Third, we project the economy forward for 830 years from its initial conditions.[30] This involves clearing the bond market if one is assumed. Fourth, we use the model's Euler conditions to update our guessed decision functions. And fifth, we repeat steps two through four until the Euler conditions are satisfied to a high degree of precision.

More specifically, our algorithm contains outer and inner loops. The outer loop solves for consumption functions of each generation. This is GSSA. The inner loop uses a combination of techniques from the numerical analysis literature—Broyden, Gauss-Seidel, and Newton's method—to compute the agents' bond holdings and the risk-free rate that clears the bond market. This is our innovation.

Recall that the state vector consists of cash-on-hand variables, $x_{g,t}$, of generations 1 through $G-1$ and exogenous shocks. Given the information at time t, agents decide how much of their cash on hand to consume, $c_{g,t}$. They also choose the proportion $\alpha_{g,t}$ of their savings to allocate to bonds at the prevailing risk-free rate $\bar{r}_t$.

The outer loop starts by making an initial guess of generation-specific consumption functions c_g as polynomials (linear, for this paper) in the state vector and the prevailing depreciation shock.[31] Next, we take a draw of the path of shocks for T periods. We then run the model forward for T periods using the economy's initial condition (i.e., the state vector of age specific cash on hand holdings that arise in the long run with zero economic shocks), guessed consumption functions and the drawn shocks. I.e., we compute cash-on-hand variables at time $t+1$ using the information we have at time t and the exogenous shocks at time

[29] Reiter (2015) develops a method for solving multi-period OLG models. He independently chose to characterize the state vector in terms of cash-on-hand. His main focus is the computation of global higher-order approximations to medium-sized OLG models, including models with asset short-sale constraints, through an efficient implementation of quasi-Newton methods.

[30] We chose 830 years to produce roughly 10 data points per 82 consumption-function coefficients entering our polynomial approximations to the 80 generations' consumption functions. This said, all the data are used to estimate all the coefficients. Longer simulation periods produce no differences in results.

[31] Although we do not include δ as part of the theoretical state space, using it as a regressor for approximating the consumption functions proved valuable.

$t+1$. Since the α's and the $\bar{r}$ that are determined at time t are realized at time $t+1$, their knowledge is necessary to compute cash-on-hand variables in period $t+1$.

In running the model forward, at each time t, we compute the agents' choice of bond shares and the risk-free rate that clears the bond market. To solve for $\bar{r}_t$, we use Broyden's method based on the bond-market clearing condition. This condition requires that the sum of bond holdings at time t equals zero. The bond holdings at time t of each agent age g is $\alpha_{g,t}\theta_{g,t}$. The choice of the $\alpha_{g,t}$'s make them functions of $\bar{r}_t$. Hence, for given values of the $\theta_{g,t}$'s, the bond-market clearing condition is a function of $\bar{r}_t$ and can be used, via Broyden's method, to find the $\bar{r}_t$ that sustains market clearing.

For any given $\bar{r}_t$, the choice of $\alpha_{g,t}$'s is determined by Gauss-Seidel iterations to solve the system of simultaneous $G-1$ generation-specific Euler equations governing the choices of the $G-1$ α's for the new values of those α's. Specifically, for given guesses of each agent's value of α, other than that of agent i, we apply Newton's method to agent's i's Euler equation to determine the new guessed value of α for agent i.[32]

Simulating the model forward produces the data needed to update our guessed consumption functions. Specifically, for each age group g and each period t, we evaluate the Euler condition to determine what that age group's consumption should be in that period. This calculation is based on the derived period-t state variables and the current guessed consumption functions of all agents, which enter, via their impact on the state vector of cash of hand at $t+1$, into the determination of any given age-g agent's marginal utility of consumption at $t+1$. The expected value is evaluated using Gaussian quadrature, as in GSSA.

Following PEA/GSSA, we then regress these time series of age-specific consumption levels on the state variables plus the depreciation shock and use the new regression estimates to update, with dampening, the polynomial coefficients of each guessed consumption function. We iterate the updating of these functions based, always, on the same draw of the path of shocks until consumption functions converge.

We evaluate the accuracy of our solutions using two methods proposed in the literature—out-of-sample deviations from the exact satisfaction of the Euler equations and the statistic proposed by Den Haan and Marcet (1989, 1993).

A.1 Out-of-Sample Deviations from the Perfect Satisfaction of Euler Equations

A satisfactory solution requires that generation-specific Euler equations (17) hold out of sample. Hence, to test the accuracy of our solution, we draw a fresh sequence of 1660 sets of

[32] Taking other unknowns as given is Gauss-Seidel.

shocks for each simulated model. We then run the model forward for 1660 years (twice the length of the original simulation), imposing the drawn shocks, using the original consumption functions, c_g, and clearing the bond market by rerunning the model's inner loop each year as we move through time. To calculate out-of-sample, unit-free deviations from full satisfaction of the Euler equations, we form

$$\begin{aligned}\epsilon(s,z) = {} & \beta E_z\big[\big(1+r(s',z')\big)\frac{u'(c_{g+1}(s',z'))}{u'(c_g(s,z))}\big]-1 \\ & + I\beta E_z\big[\big(\alpha_g(s,z)(\bar{r}(s,z)-r(s',z'))-f(\alpha_g(s,z))\big)\frac{u'(c_{g+1}(s',z'))}{u'(c_g(s,z))}\big]-1\end{aligned} \tag{A.1}$$

for each period in the newly simulated time path and for each generation $g \in 1,\ldots,G-1$. Finally, we compute the average, across time, of the absolute value of the deviations from these Euler equations for each generation.

Table A.1 reports summary statistics, across generations, of their average absolute deviations from Euler equations for each model considered.[33] As indicated, in all cases these deviations are at most 0.004. And in most cases, they are zero to the third decimal place.

The portfolio choice equations (18) and the bond market-clearing condition (8) hold almost perfectly by construction, since the α's and $\bar{r}$ that satisfying them are calculated in the inner loop with a high degree of precision. In particular, the average absolute deviations from these equations, which theoretically should equal zero, are at most 0.0005 and 0.00001, respectively, and in most cases equal zero to the seventh decimal place.

A.2 The Den Haan-Marcet Statistic

An alternative precision test is provided by Den Haan and Marcet (1989, 1993). Taylor and Uhlig (1990) use this test to compare alternative solution methods for nonlinear stochastic growth models. We follow Taylor-Uhlig's particular implementation method.

As above, we start with a fresh draw of shocks over T periods and simulate the model forward based on these shocks, using the original consumption functions and clearing the bond market each period based on the inner loop technique (discussed above). We set T again to 1660. Then, for each generation-specific Euler equation (17), we compute the

[33] Note, these deviations are not Euler errors, which capture differences in period t's marginal utility and period $(t+1)$'s realized marginal utility (properly weighted by β and $r(s',z')$). Rather, they reference mistakes in satisfying the Euler equation, i.e., the discrepancy in period t between the marginal utility and its properly weighted time-t expectation.

Mean Absolute Euler Equation Deviations

	Min	Mean	Max
	Base Model		
No Social Security	0.0000	0.0002	0.0005
Social Security	0.0001	0.0003	0.0005
	Base Model With Borrowing Costs		
No Social Security	0.0002	0.0013	0.0031
Social Security	0.0004	0.0009	0.0018
	Base Model Without Bonds		
No Social Security	0.0000	0.0002	0.0003
Social Security	0.0001	0.0003	0.0006
	Alternative Model		
No Social Security	0.0009	0.0024	0.0039
Social Security	0.0013	0.0026	0.0039
	Alternative Model Without Bonds		
No Social Security	0.0012	0.0016	0.0018
Social Security	0.0015	0.0022	0.0027
	Risk-Inducing Policy		
Bond Market	0.0000	0.0001	0.0005
No Bond Market	0.0000	0.0000	0.0002
	Sensitivity Analysis Models		
More Volatile TFP	0.0000	0.0006	0.0011
More Persistent TFP	0.0000	0.0001	0.0001
Very High Risk Aversion	0.0000	0.0000	0.0000

Table A.1: Minimum, mean, and maximum across generations of the average, across time, of the absolute value of the generation-specific, out-of-sample deviations from the perfect satisfaction of Euler equations.

residual, η_g, where g references the generation's age at time t.

$$\eta_g(t) = \beta\big(1 + r(t+1)\big)\frac{u'(c_{g+1}(t+1))}{u'(c_g(t))}. \tag{A.2}$$

We next regress, separately for each generation, their 1600 η_g values on a matrix x_g consisting of a constant, five lags of c_g, and five lags of z. The predicted values of the regression equation, $\hat{a}_g$,

$$\hat{a}_g = \big(\Sigma x_g(t)' x_g(t)\big)^{-1}\big(\Sigma x_g(t)' \eta_g(t)\big), \tag{A.3}$$

are then used to construct the Den Haan-Marcet statistic m_g as follows:

$$m_g = \hat{a}_g' \big(\Sigma x_g(t)' x_g(t)\big)\big(\Sigma x_g(t)' x_g(t) \eta_g(t)^2\big)^{-1}\big(\Sigma x_g(t)' x_g(t)\big)\hat{a}_g. \tag{A.4}$$

If the generation-specific Euler equations (17) are satisfied, then $E_{t-1}[\eta_g(t)] = 0$ must hold. This implies that the coefficient vector, and, therefore, m_g is zero, which is the null hypothesis. Note that our solution method does not enforce this property, so as Den Haan and Marcet (1994) point out, theirs is a challenging test.

The Den Haan-Marcet Statistic

	Min	Mean	Max
		Base Model	
No Social Security	6.240	12.292	21.285
Social Security	12.315	18.310	24.828
		Alternative Model	
No Social Security	5.506	10.837	23.188
Social Security	5.598	10.184	16.851

Table A.2: The minimum, mean, and maximum values across generations of the Den Haan-Marcet statistic in the base and alternative models.

Under the null, m_g is distributed as $\chi^2(11)$ asymptotically. Based on a two-sided test at the 2.5 percent significance level, we would fail to reject the null if m_g lies outside the interval (3.82, 21.92). In Table A.2 we compute the minimum, mean, and maximum across generations of generation-specific statistics m_g in the base and alternative models with and without Social Security. The mean across generations of the statistic is well within the acceptance interval for all four models. The same holds for the minimum value. Two of the four maximum values are slightly above the top value of the acceptance range.

A.3 Algorithm Step by Step

The following is a step-by-step description of our algorithm.

Initialization:

- Set $\bar{z}$ and $\bar{\delta}$ to their average values and solve for the nonstochastic steady state cash on hand of each age group without bond, $\bar{s} = (\bar{x}_1, \ldots, \bar{x}_{G-1})$. Note that the nonstochastic steady state corresponds to the state vector of age specific cash on hand holdings that arise in the long run with zero economic shocks. Let $(s_0, z_0, \delta_0) = (\bar{s}, \bar{z}, \bar{\delta})$ be the starting point of the simulation.

- Approximate $G - 1$ consumption functions by polynomials in the state variables and the shock delta: $c_1(s, z, \delta) = \phi_1(s, z, \delta; b_1), \ldots, c_{G-1}(s, z, \delta) = \phi_{G-1}(s, z, \delta; b_{G-1})$, where $b_1, \ldots, b_{G-1}$ are polynomial coefficients. We use linear polynomials. To start the iterations, we make the following initial guess for the coefficients:
$b_1 = (0, 0.9\bar{c}_1/\bar{x}_1, 0, \ldots, 0, 0.1\bar{c}_1, 0), \ldots, b_{G-1} = (0, 0, \ldots, 0, 0.9\bar{c}_{G-1}/\bar{x}_{G-1}, 0.1\bar{c}_{G-1}, 0)$.

- Take draws of the exogenous path of shocks for T years. We set T to 830, which corresponds to roughly 10 observations per polynomial coefficient.

Outer loop:

- The first step in the outer loop is to simulate the model forward, i.e. compute the state space for $t = 1, \ldots, T$. To do so, at each time t we proceed as follows:
  - Recall that at time t, the state vector consists of the vector of cash-on-hand variables of generations 1 through $G - 1$, $s_t = (x_{1,t}, \ldots, x_{G-1,t})$ and exogenous shocks.
  - Using this state vector, for each age group g, calculate its consumption $c_g^{(p)}$ given the current guess for the coefficients $b_g^{(p)}$, where the subscript (p) denotes the current iteration of the outer loop. I.e., $c_{g,t}^{(p)}$ equals the inner product of the vector $(1, s_t, z_t, \delta_t)$ with the vector of coefficients $b_g^{(p)}$. Compute the generation-specific asset demands, $\theta_{g,t}$, as the difference between cash on hand and consumption, $\theta_{g,t} = x_{g,t} - c_{g,t}$. Note that the sum of asset demands of generations 1 through $G - 1$ is the capital stock at the beginning of period $t + 1$, k_{t+1}.
  - At this point enter the inner loop to compute the agents' choices of bond shares, $\alpha_{g,t}$, for generations 1 through $G - 1$, and the risk free rate $\bar{r}_t$. Recall that these are needed to compute the cash-on-hand variables at time $t + 1$.

- Inner loop:
    * Use Broyden's method to solve (8) for $\bar{r}_t$. To start, make an (arbitrary) initial guess for the value of $\bar{r}_t$.
    * Given $\bar{r}_t$, solve the system of $G-1$ equations given by (18) for $g = 1, \ldots, G-1$, for $G-1$ unknowns, $\alpha_{1,t}, \ldots, \alpha_{G-1,t}$. To do so, approximate the expectation by Gaussian quadrature.[34] Notice that the consumption at time $t+1$, $c_{g,t+1}$, on the right-hand-side of each equation (18) needs to be approximated by the polynomial in the state vector plus δ. Hence, each of these equations depends on the entire distribution of the cash-on-hand variables, and through them, on all of the generation-specific α's, $\alpha_{1,t}, \ldots, \alpha_{G-1,t}$. To solve a nonlinear system of $G-1$ nonlinear equations in $G-1$ unknowns we use the Gauss-Seidel algorithm, which reduces the problem of solving for $G-1$ unknowns simultaneously in $G-1$ equations to that of iteratively solving $G-1$ equations in one unknown.[35] We solve each of these nonlinear equations in one unknown α using Newton's method.
    * Use $\alpha_{g,t}$ found above for all g to calculate (8) and update $\bar{r}_t$.
- Given $\alpha_{g,t}$ for $g = 1, \ldots, G-1$, $\bar{r}_t$, k_{t+1}, and exogenous shocks, we can now compute each generation's cash on hand in period t+1 as the sum of their labor and capital income (plus or minus any government transfers) at time t+1.

- Note that for each age group g and each state (s_t, z_t), $t = 1, \ldots, T$, (17) implies

$$\begin{aligned} c_g(s,z) = \big\{ &\beta E_z\big[\big(1 + r(s',z')\big)u'(c_{g+1}(s',z'))\big] \\ &+ I\beta E_z\big[\big(\alpha_g(s,z)(\bar{r}(s,z) - r(s',z')) - f(\alpha_g(s,z))\big)u'(c_{g+1}(s',z'))\big]\big\}^{-\frac{1}{\gamma}}. \end{aligned} \tag{A.5}$$

  Denote the right-hand-side of (A.5) by y_g and evaluate the expectation using Gaussian quadrature.

- For each age group g, regress y_g on (s_t, z_t, δ_t) and a constant term using regularized least squares with Trikhonov regularization. Denote the estimated regression coefficients by $\hat{b}_g^{(p)}$.

[34] We use 4 nodes in the quadrature, using more does not change the results.
[35] As the starting point for Gauss-Seidel we use the α's computed at time t-1.

- Check for convergence: If

$$\frac{1}{G-1}\sum_{g=1}^{G-1}\frac{1}{T}\sum_{t=1}^{T}\left|\frac{x_g^{(p-1)}-x_g^{(p)}}{x_g^{(p-1)}}\right|<\epsilon,$$

end. Otherwise, for each age group g update the coefficients as $b_g^{(p+1)} = (1-\xi)b_g^{(p)} + \xi\hat{b}_g^{(p)}$ and return to the beginning of the outer loop. We use $\xi = 0.1$ and $\epsilon \in [10^{-7}, 10^{-13}]$.