Generational Risk–Is It a Big Deal? Simulating an 80-Period OLG Model with Aggregate Shocks*

Jasmina Hasanhodzic† and Laurence J. Kotlikoff‡

June 2013, revised December 2015

Abstract

The theoretical literature presumes generational risk is large enough to merit study and that such risk can be meaningfully shared via appropriate government policies. This paper assesses these propositions. It develops a computational technique to overcome the curse of dimensionality and measure generational risk in an 80-period OLG model. The model features isoelastic preferences, moderate risk aversion, Cobb-Douglas technology, and shocks to both TFP and capital depreciation.

When shocks are calibrated using empirically plausible parameter values (our baseline model), the model reproduces the U.S. variability of output and wages. But it overstates the variability of the economy's overall return to aggregate wealth. Even so, the model understates the economy's risk premium (the mean return on aggregate wealth less the mean return on safe assets). But, as shown by Constantinides, Donaldson, and Mehra (2002) and Hasanhodzic (2014), adding hard or increasing borrowing costs can readily reduce equilibrium risk-free rates to essentially any desired level with little or no impact on the economy's key macro variables. Based on this calibration, we find that generational risk is small and that Social Security can exacerbate it.

When depreciation shocks are accentuated to reproduce the observed variability of equity returns, as in Krueger and Kubler (2006), the model's risk premium is close to the equity premium observed in the stock market. Yet, with this calibration, the model greatly overstates the variability of output and wages. It also produces substantial generational risk—risk that Social Security can materially reduce.

Under both calibrations, we find that even a one-period bond market is very effective, indeed far more effective than Social Security, in sharing risks among contemporaneous generations. We also find that policy-induced intergenerational redistribution can produce far larger generational risk than macroeconomic fluctuations.

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

JEL Classification: E21, E24, E62, H55, H31, D91, D58, C63, C68

*We thank Ken Judd, Thomas Sargent, Kent Smetters, Kerk Phillips, Rick Evans, Karl Schmedders, Simon Gilchrist, and Alisdair McKay for helpful comments.

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

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

1 Introduction

Economists have examined generational risk and its mitigation via government policy in a number of models (e.g., Diamond, 1977; Merton, 1983; Bohn, 1998, 2005, and 2009; Shiller, 1999; Rangel and Zeckhauser, 2001; Smetters, 2003; Krueger and Kubler, 2006; Ball and Mankiw, 2007; and Bovenberg and Uhlig, 2008). These papers don't directly measure the size of generational risk. Instead, they presume generational risk is large enough to merit study and show that properly structured pay-as-you-go Social Security and other government policies can share it.

This paper directly measures generational risk while also assessing Social Security's ability to mitigate it. It does so by simulating an 80-period overlapping generations life-cycle model with isoelastic preferences, moderate risk aversion, Cobb-Douglas production, and aggregate shocks to both total factor productivity (TFP) and the capital depreciation rate, calibrated in two main ways suggested by the literature.

In our baseline model, shocks are calibrated using standard parameter values, as in Hansen (1985), Prescott (1986), and Ambler and Paquet (1994). This model reproduces the observed variability of output, as measured using the Net National Product data from 1929 to 2014, and, consequently, the wages, remarkably well. But it overstates somewhat the variability of the economy's overall return to aggregate wealth.

In this model, we find that generational risk is small, that even a single-period bond market can significantly enhance generational risk sharing, and that pay-as-you-go Social Security, in the presence of a bond market, either worsens or provides little additional generational risk sharing. This is true even though we've modeled Social Security on a defined contribution, balanced-budget basis (i.e., with a fixed tax rate and variable benefits) to help share risks across contemporaneous young and old generations.

Our base model exhibits a small risk premium, but this makes sense. Since there is little uncorrelated risk to share among the living, there is little demand for bonds relative to supply, which explains their low price. Low prices for safe bonds means, of course, high safe rates of return and a small risk premium. Indeed, in this model, the average rate of return on the risk-free bond, 3.797 percent, is almost as high as the 3.866 percent average return on capital. But, as shown by Hasanhodzic (2014), building on Constantinides, Donaldson, and Mehra (2002), adding borrowing costs to a model very similar to ours can readily reduce equilibrium risk-free rates to essentially any desired level, with little impact on the economy's key macro variables.

When depreciation shocks are accentuated to reproduce the observed variability of equity returns as in Krueger and Kubler (2006), our model's risk premium is close to the equity premium observed in the U.S. stock market and, consequently, also delivers the stock market's Sharpe ratio. But this calibration greatly overstates the variability of output and wages.

We call this alternative calibration the "extra-shocks model". Using it generates very different results. Now generational risk is large and Social Security is quite effective in sharing it across cohorts born at different points in time. Indeed, the higher is the variability of output and wages, the larger is this form of generational risk and the more scope there is for Social Security to share it. However, even in the extra-shocks model, the bond market is still more effective than Social Security in sharing risks across contemporaneous generations. We've also examined other calibrations, which feature highly persistent or volatile TFP shocks. These calibrations also generate substantial generational risk but also at the price of unrealistically high output and wage variability.

Our concern with the extra shocks calibration stems not simply from its overstatement of the volatility of real variables. Its use to reproduce the equity premium can also be questioned given the Modigliani-Miller Theorem, which states that corporate leverage can change with no economic consequences. But different degrees of leverage will produce different equity premiums. Nor can one rely on reported debt-equity ratios to pin down the equity premium. As illustrated in Hasanhodzic (2014), the leverage ratio suffers from the same labeling problem discussed in Kotlikoff and Green (2008) with respect to officially reported government debt.

In our base model, generational risk is small for three reasons. First, long-lived agents can self insure by saving more in good times and less in bad times. This suggests the importance of modeling agents that live for a realistic number of years. Second, TFP and capital depreciation shocks harm both the young and the old. Consequently, generations, at least those that overlap, experience common shocks to a large extent. Third, shocks that differentially impact contemporaneous agents can be partially insured via financial exchange.

The second and third reasons also help explain our finding of small unshared generational risk across contemporaneous generations in the extra-shocks model in the presence of bonds. But the very high variability of wages that this model generates, together with its persistence (due to the TFP process), means that when a cohort is born matters greatly to it realized lifetime utility. This is true despite the ability of cohorts alive at a point in time to pool risk among themselves via financial trade. In short, financial trade among the living matters, but so does the absence of financial trade between the living and the unborn.

We measure generational risk in three ways. First, we place all generations into Rawl's (1971) original position and ask whether being born at one date is materially worse in terms of expected lifetime utility than being born at another date. Second, we examine the dispersion across generations born at different dates in their realized levels of lifetime utility.

Third, we consider the scope for contemporaneous generations to share risk. Each of these three measures suggest, in our base case and, we believe, most realistic calibration, that generational risk is very small.

Literature Review

Solving OLG models with both macro shocks and long-lived agents is challenging due to the curse of dimensionality. Krussel and Smith (1998) banish the curse for certain single-agent models. They show that the state space in such models can be adequately approximated by sufficient statistics, such as the size of the economy's capital stock. Gourinchas (2000) and Storesletten, Telmer, and Yaron (2007) apply the Krusell and Smith approach in 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. This possibility was envisioned by Krusell and Smith (1998), themselves.

Krueger and Kubler (2004, 2006) represents another major milestone in battling the dimensionality curse. They calculate solutions for life-cycle models experiencing macro shocks. They do so using a technique, based on Smolyak's (1963) algorithm, that efficiently chooses grid values.[1] This algorithm guarantees uniform approximation over a small (sparse) set of points in the multidimensional hypercube. Their technique considers economic behavior over the entire state space. Unfortunately, it cannot fully overcome the curse of dimensionality at least for computing models with realistic lifespans measured in years. Indeed, Krueger and Kubler (2006) limit their model to 9 periods for computational feasibility.

Ríos-Rull (1994, 1996) uses local perturbation methods to solve large-scale (55-period) OLG models subject to aggregate productivity shocks. His papers ask, in part, whether the degree of completeness in risk-sharing arrangements materially affects the economy. His answer, like ours, is no. However, we consider larger shocks and use a global rather than a local solution method.

Our means of dispelling the curse differs markedly from the aforementioned approaches. It builds on Judd, Maliar, and Maliar (2009, 2011). These papers operationalize Marcet's (1988) 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 reach with non-trivial probability. But Judd et al.'s (2009, 2011) model features infinitely-lived agents with no access to a bond market. Hence, it is quite different from our framework. Moreover, their model doesn't admit generational risk since altruistic dynasties will share risk among their current and future members as a matter of course.

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

Our solution method, detailed in the Appendix, is simple and can be easily extended and modified to handle complex OLG models. First, we draw 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 from its initial condition. 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.

Unlike the initial 2013 version of this paper, which followed Judd et al.'s (2009, 2011) use of assets as state variables, the state vector here consists of cash-on-hand and aggregate shocks. This change does not affect the results but renders our simulation method more robust.

Reiter (2015) develops a method for solving multi-period OLG models. He independently characterizes the state 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. His numerical results for six- and larger-period models include relatively large risk premiums. They also support our finding that generational risk is small.

Krueger and Kubler (2006) breakthrough study focuses not on measuring the size of generational risk, per se, but on the potential for pay-as-you-go social security to mitigate such risk. They conclude that, for purposes of effecting a Pareto improvement, generational risk is not large enough to offset the intergenerational redistribution arising under the policy.

Of course, asking pay-as-you-go Social Security to produce a Pareto improvement is a tall order given the extent to which the policy redistributes across generations, both directly and via factor-price changes. In partial equilibrium, Krueger and Kubler's model achieves this goal. But studying Social Security's risk sharing necessitates general equilibrium analysis. And as Bovenberg and Uhlig (2008) make clear, doing so requires controlling for policy-generated intergenerational redistribution. Their study shows that fully funded defined benefit pensions combined with defined contribution pensions can improve the allocation of risk. Their work compliments Bohn's (1998, 2005, 2008) extensive analysis of the risk-sharing and risk-making properties of particular government policies.

However, Bovenberg and Uhlig's 2-period framework, like Krueger and Kubler's 9-period model, appears to have too few periods to adequately measure an issue these authors don't directly address – the size of the risk needing to being shared. Compressing 80 or so years of annual shocks into 2 or even 9 periods leaves agents with far fewer opportunities to adjust their behavior (i.e., to self insure) in response to the shocks they experience through time.

Summary of Our Findings

To be clear, we fully recognize that OLG models have missing markets, i.e., that generational risk-sharing is inherently incomplete. What we ask is whether this risk is large enough for anyone, be they applied theorists or social planners, to worry about. This question is particularly relevant in the presence of bond markets, which contemporaneous generations can use to share risks. Stated differently, if generational risk is small and bond markets do a good job sharing what limited risk prevails, there may be no need to consider pay-as-you-go Social Security or other government mechanisms to effect generational risk sharing. Such a conclusion would mirror Lucas' (1987) finding that the welfare cost of macroeconomic risk is small and doesn't justify stabilization policy.

This is, indeed, what we find in our base model, where macroeconomic risk is calibrated as is standard in the RBC literature. The absolute value of the compensating differential factor—the percentage by which we need to expand or contract a cohort's annual consumption across all states it might experience to equate its expected lifetime utility to the average expected lifetime utility of all newborns—is 0.089 percent, on average. And this absolute compensating differential factor has a very small standard deviation—just 0.067 percent. Hence, Rawlsian justice prevails from an ex-ante, original position, and there is little risk associated with one's birth date.

Second, the compensating consumption differential factor (again, measured in absolute value and applied from birth to death) needed to equate realized lifetime utility of each cohort through time to the average across newborns of their realized lifetimes utilities is 1.978 percent, on average, and has a standard deviation of 1.113 percent.

Third, there is limited need to share generational risks among contemporaneous generations. As Abel and Kotlikoff (1988) show, full risk sharing among contemporaneous agents with isoelastic preferences requires equal percentage changes in consumption from one period to the next. Hence, one can measure the potential for generational risk sharing by calculating the absolute percentage adjustments in concomitant agents' consumption levels needed to achieve the same percentage change. In our base model with 830 years of observations, the largest such absolute adjustment of any agent's consumption needed to produce uniform percentage changes is 1.011 percent.

When macroeconomic risk is amplified so that the rate of return on capital matches the observed variability of equity returns as is the case in our extra-shocks model, the compensating differential factor needed to equate the realized utility of newborns through time to the average realized lifetime utility is at least 8 times larger than in our base model. This is to be expected given that the variability of output and wages that the extra-shocks model entails is very large—at least 4.5 times larger than that observed in the data and in our base model—and given the persistence of the TFP process governing the economy. However, even in this case, the absolute percent adjustments needed to achieve perfect risk sharing among the co-living are at most 2.828 percent.

Finally, precise calibration aside, we find that policy-induced intergenerational redistribution produces far larger generational risk than macroeconomic fluctuations.

We proceed below by detailing the model, specifying its calibration, presenting results, and concluding. The details of our algorithm are provided in the Appendix.

2 The Model

Our model is the canonical life-cycle model in its barest bones. We have purposely kept it simple to maximize the potential for generational risk. For example, we omit variable labor supply, which would help cohorts self insure against the model's macro shocks. Each agent works through retirement age R, dies at age G, and maximizes expected lifetime utility. There are neither borrowing nor short sale constraints on stock, although the latter would be irrelevant as no one chooses to short stock in neither in our base model nor in our extra-shocks model. 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 of workers supplies 1 unit of labor per period. More precisely, letting ℓ_g denote the labor supply of generation g, ℓ_g equals 1 for workers and 0 for retirees. Hence, total labor supply equals R. As indicated in equation (1), utility in a given year is time-separable and isoelastic, with risk aversion coefficient γ.

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

2.2 Technology

Production is Cobb-Douglas with output Y_t given by

$$Y_t = z_t K_t^{\alpha} L_t^{1-\alpha}, \tag{2}$$

where z is total factor productivity, α is capital's share of output, K_t is capital, and L_t is labor demand, which equals labor supply, i.e., R. Equilibrium factor prices are given by

$$w_t = z(1-\alpha)\left(\frac{K_t}{R}\right)^{\alpha}, \tag{3}$$

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

where δ_t is a normal random variable referencing depreciation shocks.

Total factor productivity, z, obeys

$$\ln(z_{t+1}) = \rho \ln(z_t) + \epsilon_{t+1}, \tag{5}$$

where $\epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2)$.

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 since it is known at time t. The asset demand of a household age g at time t (the assets brought into $t+1$) 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

$$K_t = \sum_{g=1}^{G} \theta_{g,t-1}. \tag{6}$$

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

$$\sum_{g=1}^{G} \alpha_{g,t}\theta_{g,t} = 0. \tag{7}$$

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.[2]

[2] I.e., all such relabeled models are isomorphisms.

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 (7) 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.

How can a bond market 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 risk. For example, workers can hedge bad TPF shocks that lower their wages by buying bonds from retirees who can use the proceeds of their bond sales 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 think. A large negative depreciation shock certainly hurts retirees, who are the primary owners of capital. But it also helps them by raising their rate of return in the next and, other things equal, subsequent periods. Workers, on the other hand, are hurt by negative depreciation shocks because they end up with less capital with which to work and, therefore, a lower wage. This too occurs in the period after the negative depreciation shock. And wages can remain lower than would otherwise have been the case to the extent that subsequent TFP and depreciation shocks are comparatively small.

2.4 Government

We model Social Security as financed by a fixed, annual 15 percent proportional tax whose benefits are provided to all retirees on an equal per-capita basis. We also consider an unrealistic, risk-inducing policy to test one of our methods of measuring generational risk. This risk-inducing policy also taxes the young and hands the proceeds to the old. But the proportion of the wages taken from the young raises steeply with wage.

Let $H_{g,t}$ denote the government tax levied on the age-g household at time t, and let $B_{g,t}$ denote the government's transfer to the age-g household at time t. Then

$$H_{g,t} = \begin{cases} \tau w_t \ell_g, & \text{with Social Security Policy} \\ \mu(w_t)\ell_g, & \text{with Risk-Inducing Policy} \end{cases} \tag{8}$$

where

$$\mu(w_t) = a\underline{w} + \frac{b\bar{w} - a\underline{w}}{\bar{w} - \underline{w}}(w_t - \underline{w}), \tag{9}$$

and $\underline{w}$ and $\bar{w}$ are minimum and maximum values of w. Parameters a, b, τ, $\underline{w}$, and $\bar{w}$ are described in Section 3. Finally,

$$B_{g,t} = (1 - \ell_g)\frac{\sum_{g=1}^{G} H_{g,t}}{80 - R}. \tag{10}$$

Implicit in the above formulation is that government taxes equal government spending,[3]

$$\sum_{g=1}^{G} H_{g,t} = \sum_{g=1}^{G} B_{g,t}. \tag{11}$$

2.5 Household Problem

At time t, the economy's state is $(s_t,\ z_t,\ \delta_t)$, with $s_t = (x_{1,t}, \ldots, x_{G-1,t})$ denoting the set of age-specific cash-on-hand holdings.[4] Households of age g in state (s, z, δ) maximize expected remaining lifetime utility given by

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

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

[3]The more appropriate statement, given the capacity to relabel, is that the government's intertemporal budget constraint is satisfied along all its possible paths.

[4]Note that depreciation shock, δ, is part of the state even though it is i.i.d. (and thus irrelevant for computing the expectation). Without its knowledge, households are unable to compute the rate of return that determines their cash on hand. Moreover, it ptovides information as to how much of the prevailing level of cash on hand is due to noise vs. the TFP shock which is likely to persist.

subject to

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

$$c_{g,t} = \ell_g w_t + \left[\alpha_{g-1,t-1}(1+\bar{r}_{t-1}) + (1-\alpha_{g-1,t-1})(1+r_t)\right]\theta_{g-1,t-1} - \theta_{g,t} - H_{g,t} + B_{g,t}, \tag{15}$$

for $1 < g < G$, and

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

where $c_{g,t}$ is the consumption of a g-year old at time t and (14)–(16) are budget constraints for age group 1, those between 1 and G, and that for age group G.

2.6 Equilibrium

Given the initial state of the economy $s_0 = (x_{1,0}, \ldots, x_{G-1,0}, z_0, \delta_0)$, the recursive competitive equilibrium is defined as follows.

Definition. The recursive competitive equilibrium is governed by the consumption functions, $c_g(s, z, \delta)$, the share of savings invested in bonds, $\alpha_g(s, z, \delta)$, factor demands of the representative firm, $K(s, z, \delta)$ and $L(s, z, \delta)$, government policy, $H(s, z, \delta)$ and $B(s, z, \delta)$, as well as the pricing functions $r(s, z, \delta)$, $w(s, z, \delta)$, and $\bar{r}(s, z, \delta)$ 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 (3) and (4).

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

$$1 = \beta E_z\left[(1 + r(s', z', \delta'))\frac{u'(c_{g+1}(s', z', \delta'))}{u'(c_g(s, z, \delta))}\right] \quad \text{and} \tag{17}$$

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

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

3 Calibration

In order to measure generational risk, we first need to calibrate our model.

3.1 Endowments and Preferences

In the base model, as well as in the model with risk-inducing policy, risk aversion γ equals 2. In the extra-shocks model it equals 3. For sensitivity analysis purposes, we also consider the variants of the extra-shocks model where we set γ to 4 or 5, and the variant of the base model where we set γ to 15. Agents work for 45 periods and live for 80. We set the quarterly subjective discount factor, β, to 0.99, a standard value founded in the macroeconomics literature. This implies an annual value of 0.96 for β.

3.2 Technology

Base Model

In the base model, we calibrate the TFP process, z, based on Hansen (1985) and Prescott (1986), which has been widely adopted in the literature.[5] Both studies specify z as an AR(1) process in logs (see equation 5). Hansen estimates that the quarterly value for the autocorrelation coefficient of this process, ρ, is 0.95 and that the standard deviation, σ, of the innovation ϵ lies in the range $[0.007, 0.01]$. Prescott's (1986) estimates are 0.9 for ρ and 0.00763 for σ.

Our quarterly values for ρ and σ are 0.95 and 0.01, respectively. This implies annual values of 0.815 for ρ and 0.019 for σ. Over the 830 years of the simulation, this entails a mean value for the level of total factor productivity of 0.997 with a standard deviation of 0.033. Their ratio yields an annual coefficient of variation in the level of total factor productivity of 3.29 percent.

In our base model, we set the quarterly standard deviation of the depreciation shock to 0.0045 (implying an annual value of 0.018).[6] This is very close to the 0.0052 quarterly estimate of Ambler and Paquet (1994).

With this calibration, the wage displays a standard deviation of 0.081 around the mean of 1.928, for a coefficient of variation of 4.22 percent.

[5] 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).

[6] We interpret Y (equation 2) as the net production function, and hence set the mean of stochastic depreciation to zero.

Extra-Shocks Model

In the extra-shocks model, the TFP process is calibrated as above, but the quarterly standard deviation of the depreciation shock is increased to 0.0342 to reproduce the Sharpe ratio of the stock market.[7] This is in line with the standard deviation of the depreciation shock employed 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 measurement issues such as the time period used and the proxy used for the risk free rate.

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 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 the range of 13.9 to 15.8 percent.[8] 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 extra-shocks calibration generates the risk premium of 3.11 percent, the standard deviation of the returns to capital of 14.1 percent, for a Sharpe ratio of 0.221. Increasing the risk aversion parameter to 4 raises the risk premium to 3.89 percent and the Sharpe ratio to 0.277. And increasing risk aversion to 5 raises these values to 4.69 and 0.334, respectively. Clearly, these values are in line with the historical record of equity returns. They also accord with practitioner literature (see, e.g., Hasanhodzic, Lo, and Patel, 2009).

And with this calibration, the wage displays a standard deviation of 0.372 around the mean of 2.105. The implied coefficient of variation—17.7 percent—is four and a half times larger than that in the base model.

Other Models

In Section 4.5 we consider the variants of the base model with more persistent or volatile TFP process and depreciation shocks turned off. When the TFP process is more persistent, its autocorrelation coefficient equals 0.990 on a quarterly basis or 0.961 on an annual basis. A more volatile TFP process features the standard deviation of its innovation that is 5 times larger than in the base model.

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

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

3.3 Government

Social Security benefits are financed via a payroll tax, τ, of 15 percent. As for the risk-inducing policy, we set its parameters, a and b, to 0.1 and 0.4, respectively. The minimum and maximum values of w, $\underline{w}$ and $\bar{w}$, are estimated by iteratively simulating the model and using the minimum and maximum wage from a previous simulation until those values converge.[9] The resulting tax rates range from 9.52 percent to 40.0 percent of the wage. Note that (8) and (9), together with the above choice of parameters, imply that the risk-inducing policy is countercyclical, i.e., the correlation between the net wage and z is negative. Since this tax rate is applied to all wages and rises steeply with the wage, it can make good times less good for workers and turn good times into better times for retirees, whose transfers can rise dramatically in the context of good shocks.

4 Results

We now discuss the results. First, in Sections 4.1 and 4.2, we measure generational risk in our base and extra-shocks models, and assess the potential of Social Security to mitigate that risk. Then, in Section 4.3, we examine whether our generational risk measures are, indeed, capturing generation risk by simulating a model with the risk-inducing policy. After that, in Section 4.4, we study the ability of the bond market to share generational risk. Finally, in Section 4.5, we assess the robustness of our findings by also simulating models with higher risk aversion, higher persistence in TFP growth, and higher volatility in TFP growth.

4.1 Base Model

Figure 1 plots the evolution over 830 years of the capital stock, output, the rate of return, and the wage. These results are from our base model, which, recall, features 80 periods, bonds, a risk aversion coefficient γ equal to 2, and empirically relevant depreciation shocks. Each panel displays results for two cases—with and without Social Security.

Figure 1's main message is the secular impact of generational policy. Where the economy finds itself in the future is primarily the function of what the government does to it. As with non-stochastic life-cycle models (see Auerbach and Kotlikoff, 1987), redistributing 15 percent of wages annually to the elderly produces substantial reductions in capital, output, wages, and rates of return. Taking from young savers and giving to old spenders under the guise of Social Security or any other set of fiscal labels leads to more aggregate consumption,

[9] In starting this iteration we use the minimum and maximum wage from the no-policy model. The final minimum and maximum values of the wage are 1.048 and 3.531, respectively.

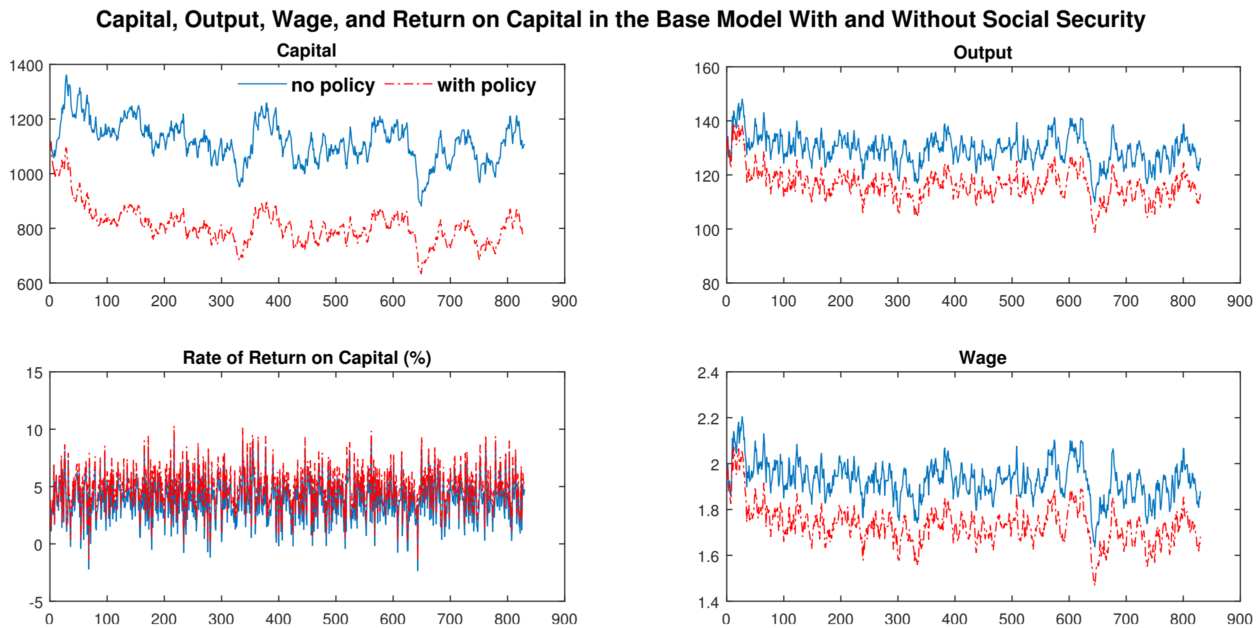

Figure 1: Capital, output, rate of return on capital, and wage rates through time in the base model with and without Social Security.

less national saving, and, in our closed economy, less domestic investment. The resulting crowding out of capital leaves wages lower and the return to capital higher.

Specifically, introducing Social Security reduces the average capital stock, calculated starting 75 years after the policy's transition, by 28.3 percent. It lowers long-run (after the 75 years) average output and average wage by 10.4 percent, and raises the long-run average return to capital by 24.8 percent.[10]

As with the calculations of average macro values, our measurements of generational risk in the presence of our Social Security policy are generated based on the model's data 75 years and beyond the start date of the policy.

In Tables 1 and 2 we report the aforementioned formal measures of generational risk, which we call expected and realized lifetime utility measures. Each newborn's expected lifetime utility measure is defined as the compensating consumption differential needed to achieve the average (expected) lifetime utility across newborns. To compute the differential for, say, the age group born in year x, we first calculate the average realized lifetime utility of generation x across different paths of draws of the productivity shock. Call this EU_x.

[10]The average values of the output, capital stock, the wage, and the rate of return on capital in the model without Social Security, calculated post transition, are 128.863, 1108.803, 1.919, and 0.039, respectively. The increase in the average return to capital may be hard to discern from the figure given the fluctuations associated with our two macro shocks.

Absolute Percent Adjustment to Achieve Average Expected Utility of Newborns Through Time
Base Model

	Mean	Standard Deviation
With Bond		
No Social Security	0.089	0.067
Social Security	0.122	0.076
No Bond		
No Social Security	0.066	0.067
Social Security	0.099	0.063

Table 1: Means and standard deviations of absolute percent adjustments needed to achieve average expected utility of newborns through time in the base model with and without bonds and with and without Social Security, post transition.

Absolute Percent Adjustment to Achieve Average Realized Utility of Newborns Through Time
Base Model

	Mean	Standard Deviation
With Bond		
No Social Security	1.978	1.625
Social Security	1.844	1.529
No Bond		
No Social Security	1.355	1.113
Social Security	1.398	1.167

Table 2: Means and standard deviations of absolute percent adjustments needed to achieve average realized utility of newborns through time in the base model with and without bonds and with and without Social Security, post transition.

Next, we divide the average of EU_t over all t from time 0 to 830 by EU_x and raise the ratio to the power $1/(1-\gamma)$. The result is the factor by which consumption of generation x needs to be multiplied in all possible states it might experience to achieve the same lifetime utility, on average, as other generations enjoy. Finally, we take the absolute value of the deviation of this factor from 1 to produce the absolute percent adjustment. 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. Given the stationary nature of the model, the expected lifetime utility differentials should be very small except for generations born while the policies are being introduced and the economy is heading toward its new stochastic steady state (ergodic distribution).

The realized utility measure is based on the particular state the generation is born into and the particular sequence of productivity shocks drawn 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 0 and 830. 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 830 generations experience on average. Finally, as above, we compute the absolute value of the deviation of this factor from 1.

As indicated above, under either measure, there is not much generational risk to be shared and Social Security these two measures virtually unchanged. Interestingly, Table 2 shows that the bond market increases generational risk experienced by newborns through time based on our realized utility measure. This makes sense. The realized utility measure compares generations who may not be contemporaneously alive. Obviously, the portfolio positions of such generations are not tailored to share risks among them. Indeed, they may vary greatly depending on the time (and, thus, the state) of birth, differentially exposing those holding them to risk.

Finally, we 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.

Absolute Percent Adjustment to Achieve Perfect Risk Sharing Base Model

	Min	Mean	Max
With Bond			
No Social Security	0.000	0.199	1.011
Social Security	0.000	0.329	1.013
No Bond			
No Social Security	0.000	0.408	2.249
Social Security	0.000	0.407	1.870

Table 3: Absolute percent adjustments to achieve perfect risk sharing in the base model with and without bonds and with and without Social Security, post transition. Minimum, mean, and maximum are taken across all cohorts and all time.

Tables 3 and 4 present these measures.[11] Table 3 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. Table 4 first reports summary statistics of pairwise correlations in annual percentage changes in consumption between different age groups among workers only, retirees only, and all agents.[12] Second, it examines the correlation of each agent's annual percentage change in consumption with the annual percentage change in per capita consumption.

Table 3 shows that the average (across agents) absolute percentage adjustments needed to achieve full risk sharing are less than 0.2 percent without Social Security, but with bonds. Hence, generational risk among contemporary cohorts is quite small even absent any government risk-sharing policy. Indeed, the maximum absolute adjustment is only 1.011 percent. Adding Social Security to the risk-sharing mix actually raises the average adjustment to 0.329.

Without bonds or Social Security, the absolute adjustments are bigger but still very small—0.408 percent on average. The maximum adjustment in this case—2.249 percent—is more than twice as big than the no-Social Security, with-bond value. Adding Social Security leaves the average absolute adjustment virtually unchanged. But it slightly reduces the maximum absolute adjustment—to 1.870 percent.

Note how even a one-period bond market is more effective than Social Security in eliminating generational risk. Compare, in this regard, the value of 0.199 percent with the value

[11] Values of 0.000 and 1.000 reflect rounding.

[12] For each age group g, we compute that age group's percentage change in consumption between t and $t-1$, $(c_{g,t} - c_{g-1,t-1})/c_{g-1,t-1}$, for all t. We then correlate the percentage change in the consumption time series for any pair of age groups.

of 0.407 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. Based on these results, the bond market is more than two times as efficient as Social Security in reducing risk among contemporaneous generations.

Correlations in Percentage Changes in Consumption
Base Model With and Without Bond

	With Bond					
	No Social Security			Social Security		
	Min	Mean	Max	Min	Mean	Max
corr(% change in C worker i; % change in C worker j)	0.947	0.995	1.000	0.957	0.996	1.000
corr(% change in C retiree i; % change in C retiree j)	0.999	1.000	1.000	0.980	0.997	1.000
corr(% change in C agent i; % change in C agent j)	0.899	0.965	1.000	0.934	0.984	1.000
corr(% change in C agent i; % change in per capita C)	0.966	0.981	0.999	0.981	0.991	0.999
	No Bond					
	No Social Security			Social Security		
	Min	Mean	Max	Min	Mean	Max
corr(% change in C worker i; % change in C worker j)	0.777	0.968	1.000	0.813	0.973	1.000
corr(% change in C retiree i; % change in C retiree j)	0.999	1.000	1.000	0.990	0.999	1.000
corr(% change in C agent i; % change in C agent j)	0.717	0.917	1.000	0.774	0.943	1.000
corr(% change in C agent i; % change in per capita C)	0.859	0.953	0.997	0.887	0.968	0.998

Table 4: Summary statistics of pairwise correlations in percentage changes in consumption among different age groups in the base model with and without bonds and with and without Social Security, post transition. Minimum, mean, and maximum are taken across different subgroups of cohorts and all time. The cohort subgroups include all workers, all retirees, and all agents (workers and retirees).

Table 4 provides a different means to assess consumption co-movements. It presents minimum, mean, and maximum consumption correlation coefficients of annual percentage changes in consumption. The statistics are computed across 755 years (i.e., the first 75 years of secular transition are dropped) and across a) all concomitant workers, b) all concomitant retirees, and c) across all agents regardless of their work status. The table also shows these correlation statistics for all agents where the correlation is not with other agents' annual percentage consumption changes but with economy-wide per capita consumption change in the year in question. Results are shown in both the presence and absence of bond markets and Social Security.

Note that, without bonds, the smallest coefficient equals 0.717. This coefficient corresponds to the correlation in percent changes in consumption between the youngest worker and the oldest retiree. The presence of a bond market raises this value to 0.899. It also raises all other correlations. Including Social Security, on the other hand, raises the above mentioned coefficient only to 0.774. This is expected given our prior demonstration that the bond market can help hedge generational risk more successfully than Social Security. If bonds are present, rounded to one decimal place, all correlations equal 0.9, regardless of Social Security. This is another very strong demonstration that the risk the economy experiences is, for the most part, not generation-specific.

Standard Deviation of Output Percent Deviations from Trend

Data/Model	Standard Deviation (%)
Net National Product data, 1929-2014	3.865
Base model	3.892

Table 5: Standard deviation of percent deviations from trend in the Net National Product data from 1929 to 2014, and standard deviation of percent deviations from mean in output in the base model.

Table 5 shows that our base model matches the variability of output measured using the Net National Product data from 1929 to 2014 remarkably well. In particular, it produces the standard deviation of the output's deviations from mean of 3.892 percent compared to the corresponding value of 3.865 in the data.[13]

Furthermore, in our base model, the standard deviation of the rate of return to capital is 1.860 percent. Based on the U.S. national accounts and FED Flow of Funds aggregate wealth data from 1945 to 2013, we estimate that the standard deviation of the rate of return on the U.S. wealth is 0.511 percent.[14] Although this deviates from our simulated value, it does so in the direction that would overstate generational risk. I.e., our finding of minimal generational risk would be even stronger if the model were calibrated to produce an even lower variability in the rate of return to capital.

Under the base calibration, the equity premium, defined as the equilibrium gap between the mean return to capital and the return to risk-free bonds, is 0.068 percent. This is trivially small, but, as explained earlier, it makes sense. Since there is little generational risk to share, there is little demand for bonds relative to supply, which explains their low price. Low prices for safe bonds means, of course, high safe rates of return and a small risk premium. Indeed, in our model, the average rate of return on the risk-free bond, 3.797 percent, is almost as

[13]Following Prescott (1986), we measure the variability of output in the data as the standard deviation of deviations from trend. Our model abstracts from trend growth, so in this case deviations from mean are considered.

[14]Note that the risky asset traded in our model is the claim to the entire capital in the economy, rather than just to the corporate portion of it. One reason for the low volatility of the economy's aggregate wealth is the fact that a significant fraction of it is government infrastructure.

high as the 3.866 percent average return on capital. Consequently, the Sharpe ratio is only 0.037.

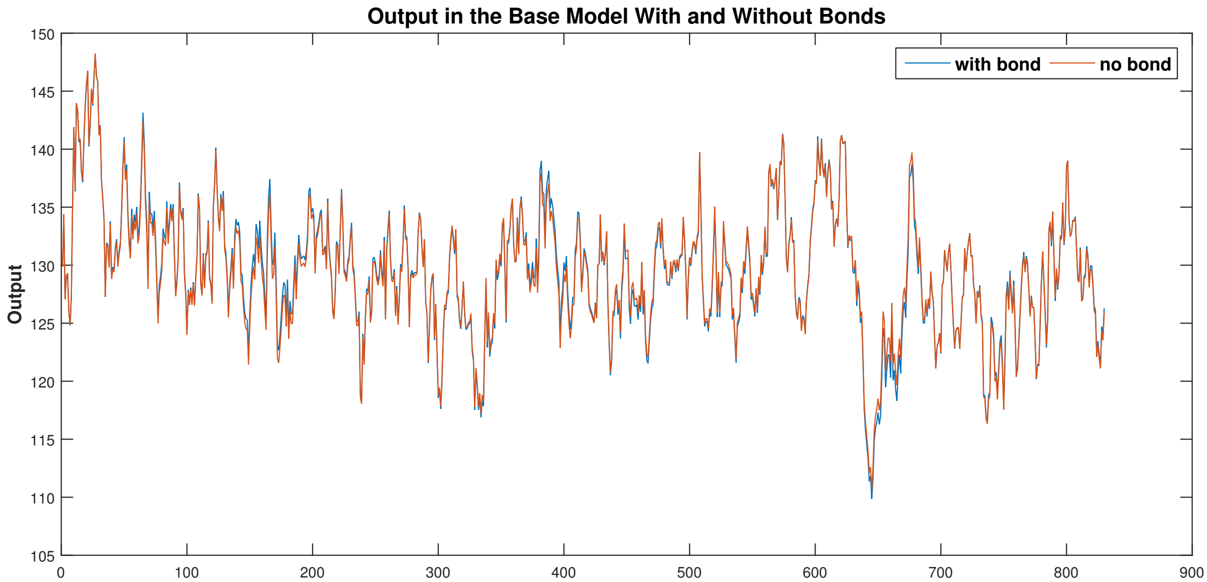

Figure 2: Output in the base model with and without bond.

Limiting the ability of the young to participate in financial markets in a setting where the young would naturally supply bonds is the mechanism used by Constantinides, Donaldson, and Mehra (2002) to illicit a higher risk premium. In their paper, like in this one, the young borrow to invest in capital. Precluding them from doing so limits the private supply of bonds, lowering the bond yield and raising the risk premium. Hasanhodzic (2014) examined this issue by introducing borrowing costs—a wedge between lending and borrowing rates—in a model very similar to ours. She showed that such borrowing costs can readily reduce equilibrium risk-free rates to essentially any desired level. However, while the presence of the bond market has strong implications for asset prices and the risk premium, it has virtually no effect on macroeconomic aggregates. Figure 2 illustrates this point in the case of output. This, together with the fact that our no-bond base model exhibits minimal generational risk, suggests that introducing borrowing constraints into our framework to generate a large risk premium would not materially change our generational risk results.

4.2 Extra-Shocks Model

In this section we consider what happens to generational risk when shocks are made large enough to reproduce the variability of equity returns observed in the data. Here we follow

Risk Premium (%) and Sharpe Ratio
Extra-Shocks Model

Risk Aversion	Risk Premium (%)	Sharpe Ratio
3	3.110	0.221
4	3.893	0.277
5	4.694	0.334

Table 6: Risk premium and the Sharpe ratio in the extra-shocks model with different levels of risk aversion.

Krueger and Kubler (2006) and increase the volatility of depreciation shocks almost eightfold compared to the empirically plausible values. Table 6 shows the risk premium and the Sharpe ratio under this calibration for different values of risk aversion. The values fall within the empirically relevant range for the stock market. Notice how risk premium and the Sharpe ratio go up as risk aversion goes up from 3 to 4 and 5.

Why might an OLG model with macro shocks produce a larger risk premium than the standard RBC model for the same degree of risk aversion? The answer is not immediately clear. But it may be that there is more capacity to self-insure in the RBC model than in the OLG model. Certainly, older cohorts in the OLG model can't count on better returns or less depreciation to bail them out since they won't be around that long. And younger cohorts with very little wealth don't have a buffer stock to tide them over if times are bad and stay bad.

In contrast, in a standard RBC model, all agents will have the same infinite number of periods over which to hedge their risk.[15] In addition, they can all draw on the dynasty's assets in bad time. Both of these features make infinitely-lived agents effectively less risk averse and less in need of a premium.

Figure 3 and Table 7 show that generating a large risk premium via large shocks is at the expense of producing the variability of output at least four and a half times larger than that in the data.

Figure 4 plots the evolution over 250 years of the capital stock, output, the rate of return, and the wage.[16] As was the case for the corresponding figure in the base model (Figure 1), the main message here is the secular impact of generational policy. Specifically, introducing Social Security reduces the average capital stock, calculated starting 75 years after the policy's transition, by 42.7 percent. It lowers long-run (after the 75 years) average

[15] Of course, the time value of resources precludes perfect self-insurance.

[16] The simulation length is 830 years as before. We plot only the first 250 years to get a finer-grained picture.

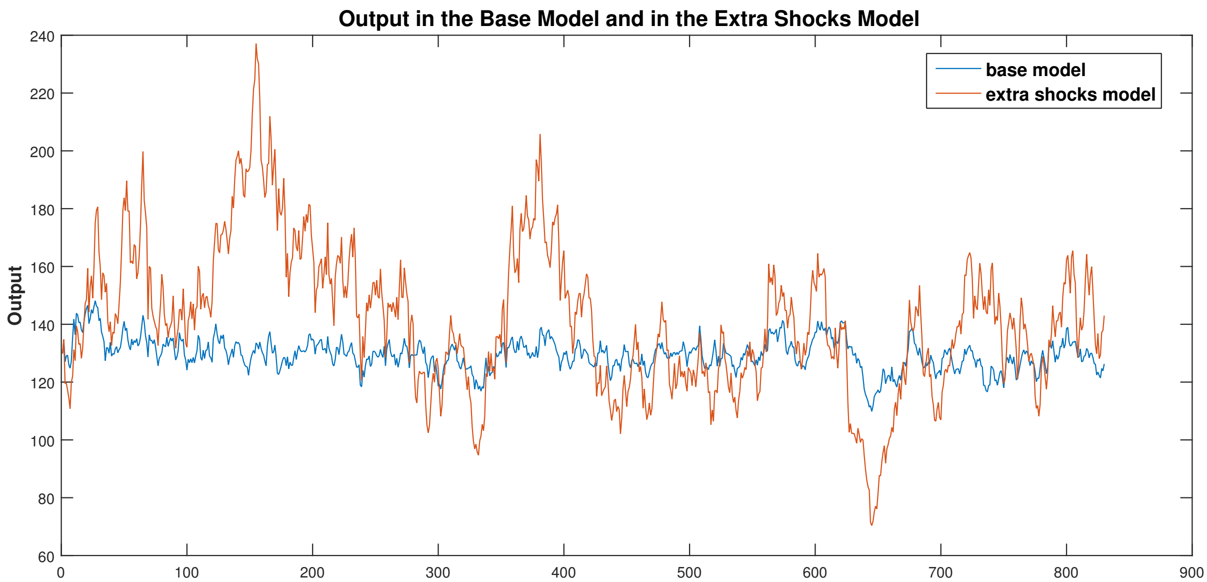

Figure 3: Output in the base model and in the extra-shocks model.

Standard Deviation of Output Percent Deviations from Trend
Extra-shocks model

Risk Aversion	Standard Deviation (%)
3	17.968
4	20.288
5	22.262

Table 7: Standard deviation of percent deviations from trend in the Net National Product data from 1929 to 2014, and standard deviation of percent deviations from mean in output in the base model.

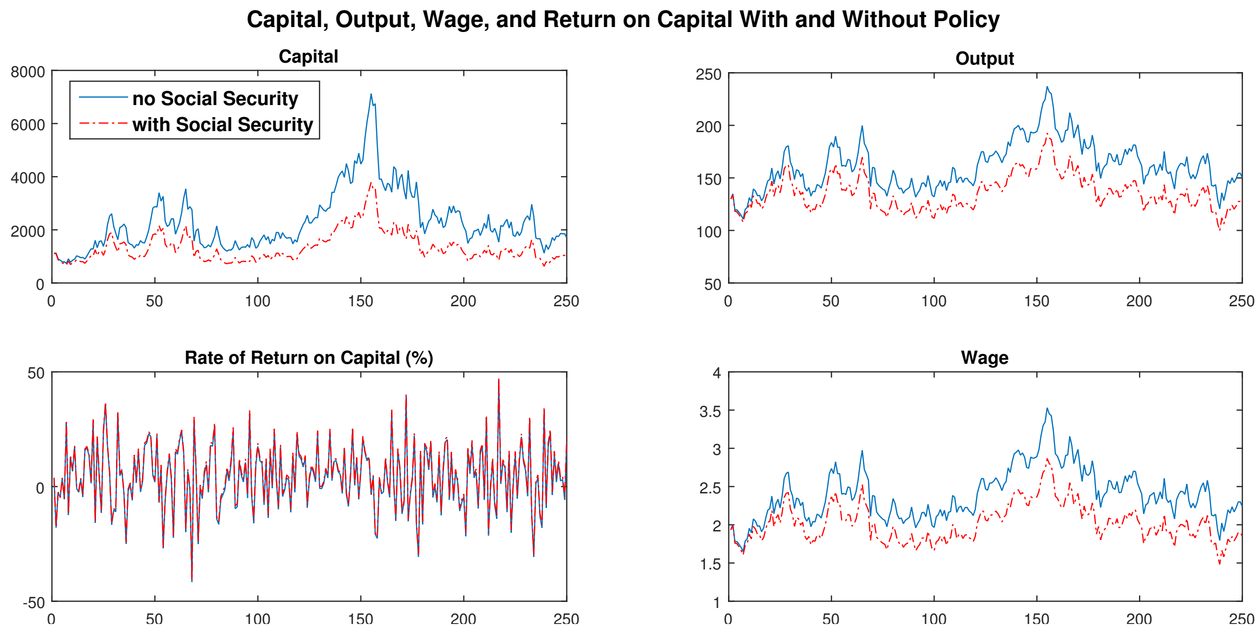

Figure 4: Capital, output, rate of return on capital, and wage rates through time in the extra-shocks model with and without Social Security.

output and average wage by 16.5 percent, and raises the long-run average return to capital by 40.8 percent.[17]

In Table 8 we report our realized lifetime utility measures in the extra-shocks model. The measures are much larger that their Table 2 counterparts. For example, under the extra shocks calibation, in the model with bonds and without Social Security, one would need to expand a newborn's consumption in all years of her life by 16.446 percent on average in order to achieve average realized utility across newborns. The correspodning value in the base model is only 1.978 percent. Thus, the extra shocks calibration is very effective in generating risk across newborns through time. Note again that since trading with the unborn is not possible, agents cannot insure this risk via the bond market. Indeed, as in 2, we see that the presence of the bond market increases generational risk according to this measure. However, with extra shocks, this increase is much more pronounced.

Table 9 presents our measures of generational risk among contemporaneous generations corresponding to Table 3 in the base case. It shows that the average (across agents) absolute percentage adjustments needed to achieve full risk sharing are less than half a percent without Social Security, but with bonds. Hence, generational risk among contemporary cohorts is quite small even absent any government risk-sharing policy. Indeed, the maximum absolute

[17]The average values of the output, capital stock, the wage, and the rate of return on capital in the model without Social Security, calculated post transition, are 140.352, 1585.809, 2.090, and 0.038, respectively.

Absolute Percent Adjustment to Achieve Average Realized Utility of Newborns Through Time
Extra Shocks Model

	Mean	Standard Deviation
With Bond		
No Social Security	16.446	13.124
Social Security	14.073	11.624
No Bond		
No Social Security	10.454	8.343
Social Security	7.774	6.760

Table 8: Means and standard deviations of absolute percent adjustments needed to achieve average realized utility of newborns through time in the extra-shocks model with and without bonds and with and without Social Security, post transition.

Absolute Percent Adjustment to Achieve Perfect Risk Sharing
Extra Shocks Model

	Min	Mean	Max
With Bond			
No Social Security	0.000	0.458	2.828
Social Security	0.000	0.256	3.529
No Bond			
No Social Security	0.000	2.988	24.043
Social Security	0.000	2.054	14.088

Table 9: Absolute percent adjustments to achieve perfect risk sharing, post transition, in the extra-shocks model with and without bonds and with and without Social Security, post transition. Minimum, mean, and maximum are taken across all cohorts and all time.

adjustment is only 2.828 percent.

Adding Social Security does not do much. It does reduce the average absolute adjustment to a small degree. But it also raises the maximum adjustment. Without bonds or Social Security, the absolute adjustments are bigger—3.00 percent on average, with a much larger maximum value of 24.0 percent. Social Security reduces the average absolute adjustment by roughly one third—to 2.05 percent. And it reduces the maximum absolute adjustment by over two-fifths—to 14.1 percent.

But, as was the case in the base model, even a one-period bond market is far more effective than Social Security in reducing generational risk among the living. To see this, compare the value of 0.458 percent with the value of 2.054 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. These results indicate that the bond market is more than four times as efficient as Social Security in reducing risk among contemporaneous generations.

Absolute Percent Adjustment to Achieve Average Realized Utility of Newborns Through Time

Extra Shocks Model With Higher Risk Aversion, Bonds, and No Social Security

	Mean	Standard Deviation
Risk aversion = 4	18.264	14.971
Risk aversion = 5	20.078	16.675

Table 10: Means and standard deviations of absolute percent adjustments needed to achieve average realized utility of newborns through time in the extra-shocks model with higher risk aversion and bonds and without Social Security, post transition.

Table 10 shows that increasing risk aversion to produce an even larger risk premium and the Sharpe ratio in the extra-shocks model further increases our realized utility measure of generational risk. Comparison with Table 7 shows that the increase in generational risk precisely coincides with the increase in variability of output and wages.

4.3 Generational Risk Inducing Policy

To ensure that our generational risk measures are, indeed, capturing generational risk, we next simulate our base model, but with risk-inducing policy and with depreciation shocks turned off to maximize the policy's ability to exacerbate TFP shocks. In considering this policy, bear in mind that it tends to persist because it's tied to the wage, which is tied to TFP, whose process is autocorrelated.

By construction, the risk-inducing policy differentially affects the young and old. It does so by making the good times less good for the young and even better for the old. The younger young and the older old are particularly affected. Indeed, we would expect smaller correlations in percentage changes in consumption the further agents are in age. Table 11 shows that this is precisely what we find.

Correlations of Age 2's Percent Changes in Consumption with Those of Other Age Groups

Age 2 correlated with age	3	25	34	35	36	37	38	39	55	80
Correlation	1.000	0.998	0.940	0.858	0.567	-0.189	-0.699	-0.858	-0.978	-0.980

Table 11: Correlations of age 2's percent changes in consumption with percent changes in consumption of other age groups.

Note that the specific risk-inducing policy that we consider exaggerates features that would arise under a Social Security policy that adjusted the system's tax rate up in bad times and down in good times in order to maintain benefit levels. This is the type of risk-making defined benefit pension system examined in various papers by Bohn and by Bovenberg and Uhlig (2008).

Measures of Generational Risk With Risk-Inducing Policy, With or Without Bond
Assumed TFP, No Stochastic Depreciation

	No Bond			With Bond		
	Min	Mean	Max	Min	Mean	Max
corr(% change in C worker i; % change in C worker j)	-0.975	0.383	1.000	1.000	1.000	1.000
corr(% change in C retiree i; % change in C retiree j)	0.996	0.999	1.000	0.994	1.000	1.000
corr(% change in C agent i; % change in C agent j)	-0.980	0.029	1.000	0.993	1.000	1.000
corr(% change in C agent i; % change in per capita C)	-0.977	0.118	0.985	0.990	0.998	0.998
all agents: absolute % adjustment for perfect risk sharing	0.000	0.899	7.478	0.717	0.924	1.146

Table 12: Measures of generational risk with risk-inducing policy with and without bonds, post transition. The measures are presented as summary statistics of pairwise correlations in percentage changes in consumption among different age groups (first four rows) and absolute percentage adjustments needed to achieve full risk sharing (last row).

Table 12, which incorporates the risk-inducing policy, displays results that are dramatically different from those in Tables 3 and 4. Take row 4, for example, which shows the minimum, mean, and maximum values of the correlation of each agent's percentage change in consumption with per capita consumption. These three values, in the no-bond-market case, are −0.977, 0.118, and 0.985. In the base case (Table 4), the corresponding values are 0.717, 0.917, and 1.000. Hence, we see what we expect—larger measured generational risk in the presence of larger generational risk.

The addition of bonds reverses this story. With bonds, the Table-12, row-4 values are 0.990, 0.998, 0.998 indicating, yet again, that the bond market is an effective means of hedging generational risk. This is also clear from the last row in Table 12. Without bonds, the maximum adjustment needed to achieve full risk sharing among living generations is 7.478 percent. With bonds, it's far less—1.146 percent.

4.4 Risk Sharing Via the Bond Market

We next examine how bonds enable agents to share generational risk. We start with the base model where this risk is intrinsic given the model's calibration. We then turn to the model with stochastic depreciation turned off, where generational risk is induced by the policy.

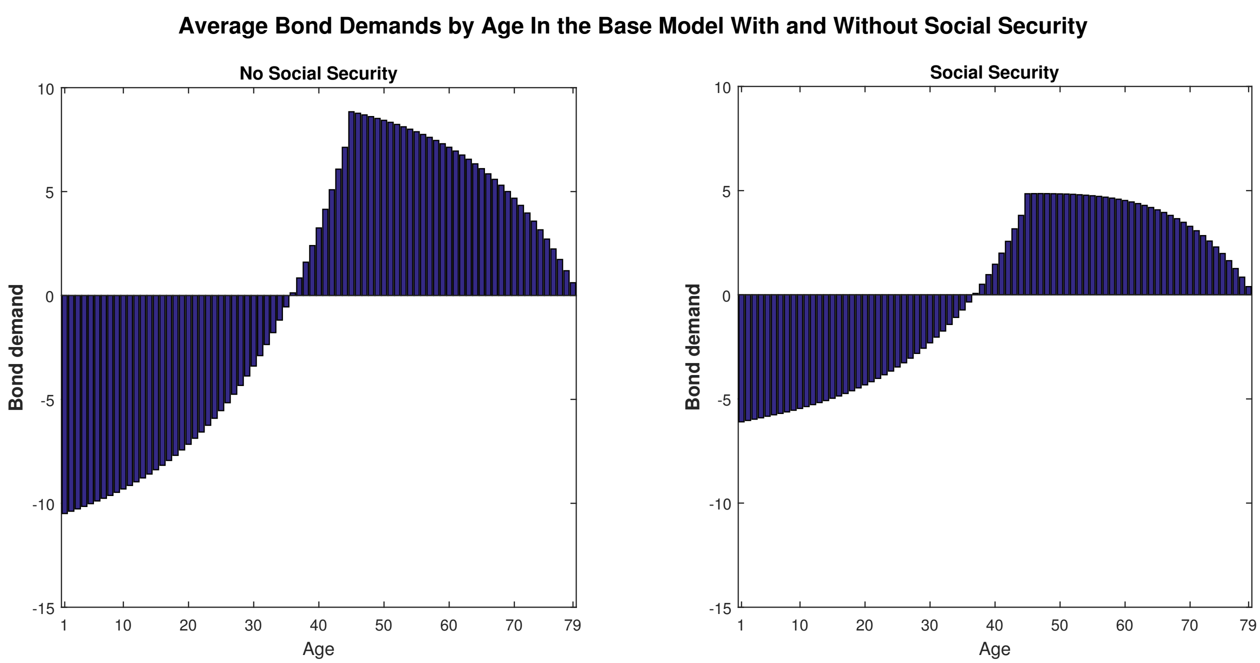

Figure 5: Average bond demands by age in the base model, with and without Social Security.

Figure 5 plots age-specific average bond demands, i.e., the time average of absolute amounts demanded in bonds by each age group. It does so in the base model with and without Social Security. It shows that, regardless of Social Security, this model exhibits a pattern of increasing bond demands with age. This pattern is intuitive since the old live off their principal and asset income. With stochastic depreciation, both the principal and the return on assets are uncertain and, thus, can be lost. On the other hand, the young's resources are primarily tied up in their labor income. Furthermore, wage rates and equity returns are uncorrelated, exhibiting a correlation coefficient of -1.92 percent, post transition. This falls within empirical estimates of Davis and Willen (2000). Hence, the young borrow against their future labor income to buy stock and thereby diversify their income sources.[18] In other words, they insure the old by selling them bonds (borrowing from them).[19] Figure 6 shows that the insurance purchased by the old is effective: the coefficient of variation of the oldest age group's consumption declines by over 36 percent, from 7.70 percent to 4.92 percent, when the bonds are introduced in the model.

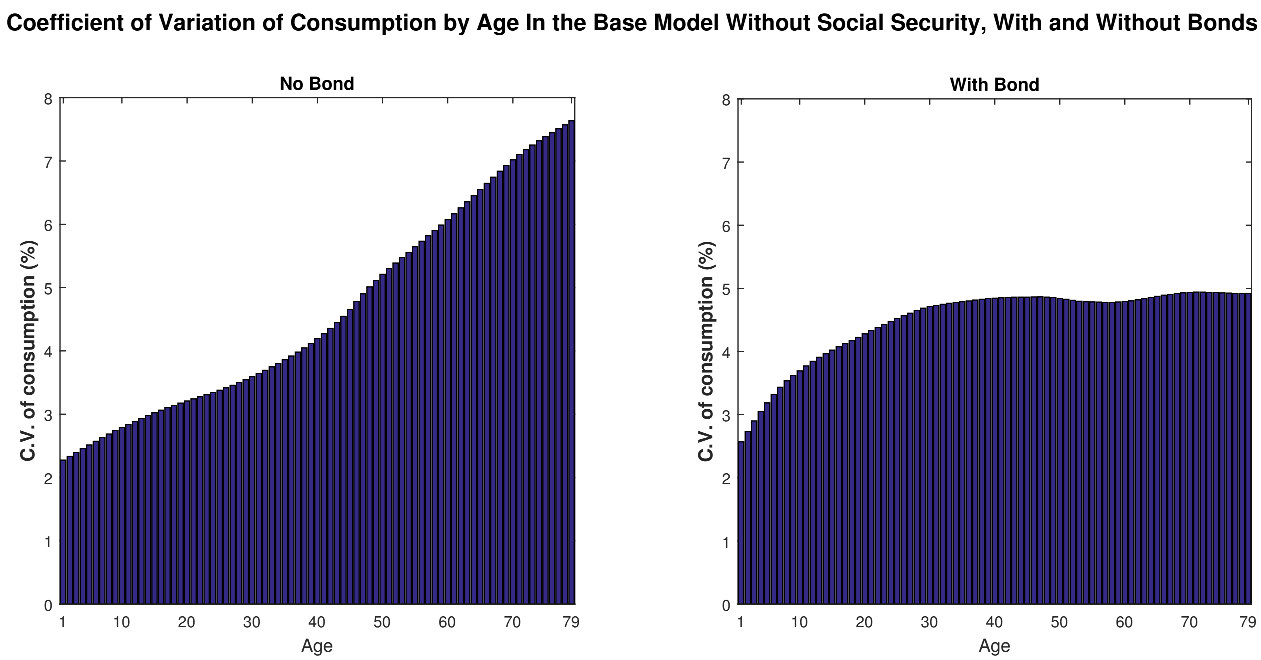

Figure 6: Coefficient of variation, in percent, of the age-specific consumption time series in the base model without Social Security, with and without bonds.

Figure 7 shows that turning off stochastic depreciation flips the pattern of average bond holdings by age observed in the base model. Now it is the young who buy bonds from the old. The old are in a position to absorb more risk from TFP shocks since the principal of their assets is unaffected. Again, Social Security does not alter this pattern. But when

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

[19]Despite their negative bond positions, the young in our base model do not show up with negative cash on hand.

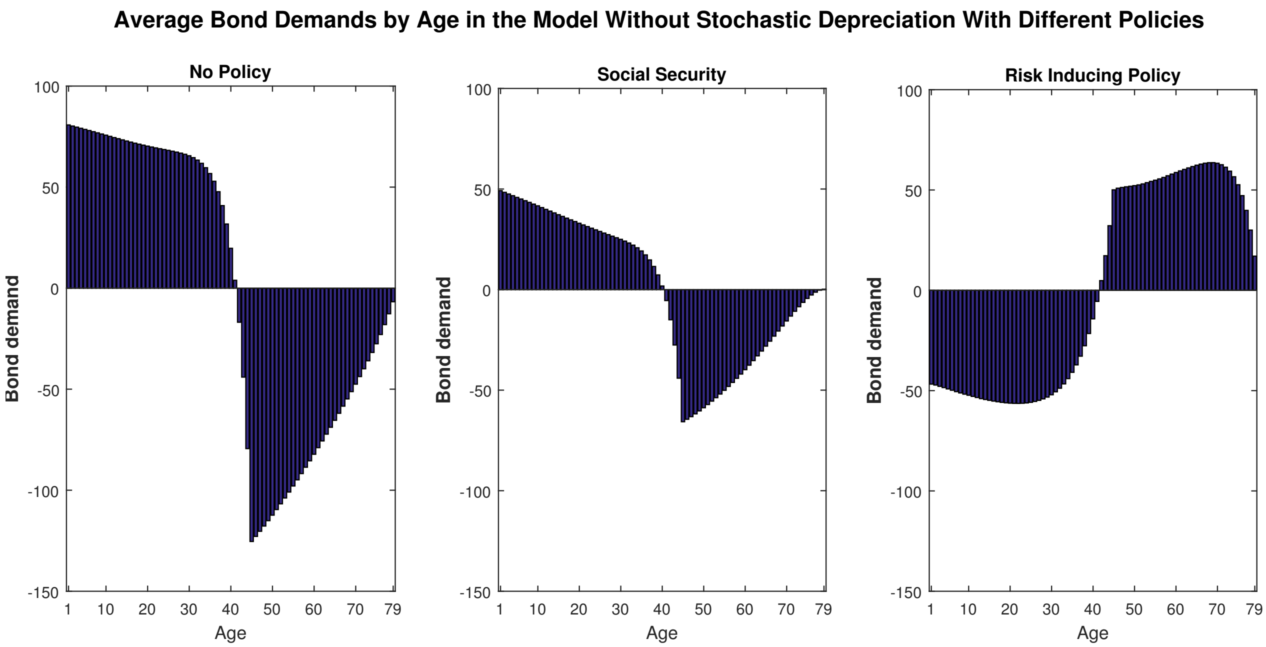

Figure 7: Average bond demands by age in the model without stochastic depreciation under various policy assumptions (without any policy, with Social Security, and with risk-inducing policy).

the risk-inducing policy is introduced, bond positions resemble those observed in the base model. Here, like there, the young sell bonds to the old in order to buy additional stock. This makes sense. When the stock market does well (its return is high due to a positive TFP shock), the young end up with smaller increases in their net wages whereas the old receive not just higher returns on their stocks, but also larger transfer payments. To improve the intergenerational risk allocation, the young borrow from the old to buy stock.

At first sight, the positive bond holdings of the young may seem large, but they are reasonable. To see this, note that the young short stocks to insure against an adverse shock in z and the resulting decline in wage. Abstract from any government policy and consider two scenarios. In one, the beginning-of-period capital is equal to the average capital stock over 830 periods, 1113.822, and z is equal to the average z, 0.999, implying a wage of 1.929 and a rate of return on capital of 0.038 per period. In the other, capital is again equal to its average value and z is one standard deviation below average, at 0.967, implying a wage of 1.869 and a rate of return on capital of 0.037. One measure of the young's potential loss in wages is the difference in wage between the two scenarios, 0.061. Since the difference in the rates of return on capital between the two scenarios is 0.001 and the average bond demand of the youngest age group is 80.763, the youngest gain 0.081 in consumption units when the adverse shock hits. Hence the potential capital gain covers the loss.

4.5 Sensitivity Analysis

Absolute Percent Adjustment to Achieve Average Realized Utility of Newborns Through Time

Models With More Volatile or Persistent TFP or Higher Risk Aversion

TFP	Risk Aversion	Mean	Standard Deviation
More persistent	2	4.461	3.656
More volatile	2	6.620	5.002
Assumed	15	1.375	0.954
Assumed	2	1.282	0.992

Table 13: Means and standard deviations of absolute percent adjustments needed to achieve average realized utility of newborns through time, post transition, in the model with high risk aversion, or with high persistence or volatility of the TFP process. All models feature a bond market and no depreciation shocks or Social Security.

To assess the robustness of our finding of small generational risk, we also simulate models with high risk aversion (15), or with high persistence or volatility of the TFP process. Recall that, when the TFP process is more persistent, its autocorrelation coefficient equals 0.990 on a quarterly basis or 0.961 on an annual basis (in the base model, the corresponding values are 0.950 and 0.814, on quarterly and annual bases respectively). A more volatile TFP process features the standard deviation of its innovation that is 5 times larger than in the base model. All models feature a bond market and no depreciation shocks or Social Security. Table 13 shows that increasing the volatility or the persistence of the TFP process increases our realized utility measure of generational risk. Table 14 shows that this coincides with the increase in the variability of output (and, thus, the wages) beyond the empirically relevant values. Under the standard calibration of TFP shocks, increasing the risk aversion to 15 has minimal effect on both output variability and generational risk. Table 15 shows that each of the above variants confirms our finding in Sections 4.1 and 4.2 of very small generational risk across contemporaneous generations in the presence of the bond market.

Standard Deviation of Output Percent Deviations from the Mean

Model	Risk Aversion	Standard Deviation (%)
Base model	15	3.662
More persistent TFP	2	8.739
More volatile TFP	2	18.644

Table 14: Standard deviation of percent deviations from the mean in output, post transition, in the model with high risk aversion, or with high persistence or volatility of the TFP process. When risk aversion is high, it equals 15, otherwise it equals 2. When the TFP process is more persistent, its autocorrelation coefficient equals 0.99 on a quarterly basis (0.96 on an annual basis). When it is more volatile, the standard deviation of its innovation is 5 times larger than that assumed in the base model. All models include bonds and do not include depreciation shocks.

Generational Risk With More Volatile or More Persistent TFP Process, and High Risk Aversion

	Assumed TFP Risk aversion 15			More persistent TFP Risk aversion 2		
	Min	Mean	Max	Min	Mean	Max
corr(% change in C worker i; % change in C worker j)	1.000	1.000	1.000	1.000	1.000	1.000
corr(% change in C retiree i; % change in C retiree j)	1.000	1.000	1.000	1.000	1.000	1.000
corr(% change in C agent i; % change in C agent j)	1.000	1.000	1.000	1.000	1.000	1.000
corr(% change in C agent i; % change in per capita C)	0.997	0.997	0.997	0.999	0.999	0.999
all agents: absolute % adjustment for perfect risk sharing	0.000	0.036	0.092	0.005	0.129	0.234

	More volatile TFP Risk aversion 2			Assumed TFP Risk aversion 2		
	Min	Mean	Max	Min	Mean	Max
corr(% change in C worker i; % change in C worker j)	1.000	1.000	1.000	1.000	1.000	1.000
corr(% change in C retiree i; % change in C retiree j)	1.000	1.000	1.000	1.000	1.000	1.000
corr(% change in C agent i; % change in C agent j)	0.999	1.000	1.000	1.000	1.000	1.000
corr(% change in C agent i; % change in per capita C)	0.998	0.999	0.999	0.999	0.999	0.999
all agents: absolute % adjustment for perfect risk sharing	0.000	0.133	0.495	0.076	0.129	0.176

Table 15: Measures of generational risk in the model with high risk aversion, or with high persistence or volatility of the TFP process, post transition. All models feature a bond market and no depreciation shocks Social Security. The measures of generational risk include summary statistics of pairwise correlations in percentage changes in consumption among different age groups (first four rows) and absolute percentage adjustments needed to achieve full risk sharing (last row).

5 Conclusion

Policy aside, when the model is calibrated using empirically informed parameter values to reproduce the variability of key macro variables, generational risk appears to be remarkably small. This risk is smaller still if generations trade bonds. Once in place, pay-go Social Security can even exacerbate generational risk. Moreover, its introduction as well as the adoption of any other major generational policy that systematically redistributes from one generation to another independent of their needs represents a major source of generational risk to all other generations who will need to cover the associated cost.

While our model features a risk-free rate that is almost as high as the return to capital, it can be modified, via hard or increasing borrowing costs, to produce essentially any desired level of equilibrium risk-free rates. Since this would have little or no impact on the economy's key macro variables, we believe that it would also have minimal effect on generational risk. We leave the implementation of these costs to future research.

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. Our solution algorithm is simple, fast, and highly robust based on a range of simulations including those featuring very large and highly persistent shocks, as well as very high degrees of risk aversion.

A Appendix

A.1 Algorithm

On a high level, our algorithm consists of an outer loop and an inner loop. In the outer loop we solve for consumption functions of each generation. In the inner loop we use 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.

We now give a more detailed overview of the algorithm. Further below is a step-by-step description. First, recall that the state vector consists of cash-on-hand variables, $x_{g,t}$, of generations 1 through $G-1$ and exogenous shocks z_t and δ_t. Given the state, at each time t agents decide how much of their cash on hand to consume, $c_{g,t}$ (or, equivalently, how much to save, $x_{g,t}-c_{g,t}$). 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. Next, we take a draw of the path of shocks for T periods. We then run the model forward for T periods using the 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 $t+1$. 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. At each iteration of Broyden's method, we use the Gauss-Seidel algorithm to solve the system of simultaneous $G-1$ generation-specific equations for the $G-1$ α's. In turn, at each iteration of Gauss-Seidel, we use Newton's method to solve each of the $G-1$ equations in one α.

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. Following GSSA, we then regress these time series of generation-specific consumption functions on the state variables and use the new regression estimates to update 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.

The following is the 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, $\bar{s} = (\bar{x}_1, \ldots, \bar{x}_{G-1})$. 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: $c_1(s, z, \delta) = \phi_1(s, z, \delta; b_1), \ldots, c_{G-1}(s, z, \delta) = \phi_{G-1}(s, z, \delta; b_{G-1})$, where $b_1, \ldots, b_{G-1}$ are polynomial coefficients. We use degree 1 polynomials. To start iterations, 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 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 (7) 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 the Gaussian quadrature.[20] 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. 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.[21] 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 (7) 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.

[20] We use 4 nodes in the quadrature, using more does not change the results.

[21] As the starting point for Gauss-Seidel we use the α's computed at time t-1.

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

$$c_g(s_t, z_t, \delta_t) = \left\{ \beta E_z \left[u'(c_g(s_{t+1}, z_{t+1}, \delta_{t+1}))(1 + r(s_{t+1}, z_{t+1}, \delta_{t+1})) \right] \right\}^{-\frac{1}{\gamma}}. \tag{A.1}$$

  Denote the right-hand-side of (A.1) by $y_g(s_t, z_t, \delta_t)$ and evaluate the expectation using Gaussian quadrature.

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

- 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 $\epsilon \in [10^{-7}, 10^{-13}]$.

We implemented our algorithm in Matlab and ran it on a Dell Precision laptop. All base-model simulations converged within half an hour. We did not utilize any low-level optimizations.

A.1.1 Out-of-Sample Deviations from the Perfect Satisfaction of Euler Equations

A satisfactory solution requires the generation-specific Euler equations (17) hold out of sample. Hence, to test the accuracy of solutions, for each model considered we draw a fresh sequence of shocks. We then simulate the model forward on the new path of 1000 years of shocks, using the original consumption functions, c_g, and clearing the bond market by rerunning the inner loop. We calculate the out-of-sample unit-free deviations from full satisfaction of the Euler equations,

$$\epsilon(s, z, \delta) = \frac{\beta E_z \left[(1 + r(s', z', \delta')) u'(c_{g+1}(s', z', \delta')) \right] - u'(c_g(s, z))}{u'(c_g(s, z))}, \tag{A.2}$$

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 Euler equations for each generation.

Deviations from Euler Equations

Model					Mean Absolute Euler Eq. Deviations		
γ	Policy	Bond	TFP	Stochastic Depreciation	Min	Mean	Max
2	None	Yes	Base	Base	0.0000	0.0001	0.0001
2	Social Security	Yes	Base	Base	0.0000	0.0001	0.0002
2	None	No	Base	Base	0.0000	0.0000	0.0001
2	Social Security	No	Base	Base	0.0000	0.0001	0.0002
3	None	Yes	Base	Extra Shocks	0.0008	0.0022	0.0036
3	Social Security	Yes	Base	Extra Shocks	0.0011	0.0020	0.0031
3	None	No	Base	Extra Shocks	0.0007	0.0017	0.0024
3	Social Security	No	Base	Extra Shocks	0.0012	0.0021	0.0026
2	Risk Inducing	No	Base	No	0.0000	0.0000	0.0001
2	Risk Inducing	Yes	Base	No	0.0000	0.0001	0.0005
2	None	Yes	Base	No	0.0000	0.0000	0.0000
2	Social Security	Yes	Base	No	0.0000	0.0000	0.0000
2	None	Yes	More volatile	No	0.0000	0.0004	0.0007
2	None	Yes	More persistent	No	0.0000	0.0000	0.0001
15	None	Yes	Base	No	0.0000	0.0000	0.0000
4	None	Yes	Base	Extra Shocks	0.0007	0.0022	0.0034
5	None	Yes	Base	Extra Shocks	0.0009	0.0025	0.0044

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.

Table A.1 reports the summary statistics, across generations, of their average absolute deviations from Euler equations for each model considered.[22] It shows that in all cases these deviations are at most 0.004. And in most cases, they are in the fourth decimal place.

The portfolio choice equations (18) and the bond market clearing condition (7) hold almost perfectly by construction, since α's and $\bar{r}$ that satisfy them are calculated in the inner loop with a high degree of precision. In particular, the deviations from these equations are at most 10^{-6} and 10^{-7}, respectively.

[22] 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 $r(s', z', \delta')$). Rather, they reference the discrepancy in period t between the marginal utility and its properly weighted expectation.

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.

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

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., 2005, "Who Bears What Risk? An Intergenerational Perspective", Working paper, Department of Economics, University of California at Santa Barbara.

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.

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.

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

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

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

Gourinchas, P.-O., 2000, "Precautionary Savings, LifeCycle 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 "Institutional Foundations of Public Finance", Harvard University Press, 241–256.

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

Hasanhodzic, J., 2014, "Increasing Borrowing Costs and the Equity Premium", SSRN 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 Papers 19179, National Bureau of Economic Research.

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

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

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.

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

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.

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., forthcoming, "Solving OLG Models with Many Cohorts, Asset Choice and Large Shocks", personal correspondence.

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.