NBER WORKING PAPER SERIES

SAFE ASSETS

Robert J. Barro
Andrew Mollerus

Working Paper 20652
http://www.nber.org/papers/w20652

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2014

We appreciate comments from Ben Broadbent, Ricardo Caballero, Xavier Gabaix, Oded Galor, Nicola Gennaioli, Matteo Maggiori, Ricardo Perez-Truglia, Andrei Shleifer, Jon Steinsson, Mike Woodford, Luigi Zingales, and participants of seminars at Columbia and Brown. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Safe Assets
Robert J. Barro and Andrew Mollerus
NBER Working Paper No. 20652
October 2014
JEL No. G1,E0,E2

ABSTRACT

A safe asset's real value is insulated from shocks, including declines in GDP from rare macroeconomic disasters. However, in a Lucas-tree world, the aggregate risk is given by the process for GDP and cannot be altered by the creation of safe assets. Therefore, in the equilibrium of a representative-agent version of this economy, the quantity of safe assets will effectively be nil. With heterogeneity in coefficients of relative risk aversion, safe assets may take the form of private bond issues from low-risk-aversion to high-risk-aversion agents. The model assumes Epstein-Zin/Weil preferences and log utility (intertemporal elasticity of substitution equal to one) and achieves stationarity by having agents die off and be replaced. We derive the steady-state quantity of safe assets and the shares of each agent in equity ownership and overall assets. In a baseline case, the risk-free rate is 1.0% per year, the unlevered equity premium is 4.2%, and the quantity of safe assets ranges up to 10% of economy-wide assets (comprising the capitalized value of the full GDP). A disaster shock leads to an extended period in which the share of wealth held by the low-risk-averse agent and the risk-free rate are low but rising and the ratio of safe to total assets is high but falling. In the baseline model, Ricardian Equivalence holds in that added government bonds have no effect on the risk-free rate and the net quantity of safe assets. The implied crowding-out coefficient for private bonds with respect to public bonds is around -0.5, a value found in some existing empirical studies.

Robert J. Barro
Department of Economics
Littauer Center 218
Harvard University
Cambridge, MA 02138
and NBER
rbarro@harvard.edu

Andrew Mollerus
Harvard University
mollerus@college.harvard.edu

In a Lucas-tree world (Lucas [1978]), the aggregate risk reflects the uncertainty in the process for GDP, which corresponds to the fruit that drops from the tree. This process may include rare macroeconomic disasters, which correspond to sharp and possibly permanent drops in the productivity or number of the trees. A safe asset in this world can be viewed as one whose real value is insulated from shocks, including the declines in GDP due to the rare disasters. However, if the GDP process is given, safe assets cannot mitigate overall risk but can only redistribute this risk across agents. In a representative-agent setting, the redistribution of aggregate risk cannot occur, and the economy's quantity of safe assets will effectively be nil.

To put this observation another way, it is possible to construct safe assets by issuing risk-free private bonds, by creating a financial structure with risk-free tranches, by entering into a variety of insurance contracts, and so on. However, the creation of any of these safe assets always goes along with a corresponding expansion in the riskiness of (levered) claims on the underlying asset, which is the Lucas tree. In equilibrium, the representative agent ends up holding the representative share of the overall risk, and this overall magnitude is unaffected by the various financial arrangements. The bottom line is that a meaningful analysis of safe assets requires heterogeneity across agents.

Differences in the degree of risk aversion are a natural form of heterogeneity for a study of safe assets. The present analysis relies on the simplest possibility, where there are two types of agents; group 1 has comparatively low risk aversion and group 2 has comparatively high risk aversion. Specifically, when each agent i has a constant coefficient of relative risk aversion, γ_i, we assume $0<\gamma_1\leq\gamma_2$, so that agent 1 is at least as willing as agent 2 to absorb risk.

We focus on a model in which the desire to redistribute risk across agents is the source of safe private assets. In equilibrium, the agent with relatively low risk aversion, agent 1, issues safe bonds (or equivalent claims) that are held by the agent with relatively high risk aversion, agent 2. Correspondingly, agent 1 owns a disproportionate share of risky assets, which are equity claims on the Lucas tree. The quantity of safe assets in this economy equals the magnitude of the bonds issued by agent 1 and held by agent 2. The equilibrium amount of these assets depends on differences in risk aversion across the agents, levels of risk-aversion coefficients, and the characteristics of the stochastic process (including rare disasters) that generates aggregate GDP.

The equilibrium requires an enforcement mechanism for repayments of safe claims; that is, agent 1 has to make payments of principal and interest to agent 2 even in bad states of the world, such as realizations of macroeconomic disasters. Repayment mechanisms may involve collateral, liquidity, and contractual features related to the legal system. However, these mechanisms are not the subject of the present analysis, which focuses on the underlying supply of and demand for safe private assets. Potentially complementary research that emphasizes liquidity, collateral, and asymmetric information includes Holmstrom and Tirole (1998) and Gorton and Ordoñez (2013). These issues could be considered in an extension of the present model.

A pure claim on the Lucas tree corresponds to unlevered equity. A match with the empirically observed high equity premium requires the expected rate of return on this equity to be substantially higher than the risk-free rate, which equals the rate of return on non-contingent, private bonds. Previous analyses with rare-disaster models, summarized in Barro and Ursúa (2012), found that the replication of this high equity premium requires two conditions: first, a coefficient of relative risk aversion, γ, in the range of 3-4 or more (for a representative agent) and, second, the presence of fat-tailed uncertainty, such as a non-negligible potential for drops in GDP in the short run by more than 10%. The present analysis incorporates these features.

With the familiar specification where utility is time-separable with a power form, a coefficient of relative risk aversion, γ, of 3-4 or more implies an intertemporal elasticity of substitution (IES) of 1/3-1/4 or less, which seems unrealistically low. Specifically, the high γ needed to generate a realistic equity premium precludes the case of log utility in the sense of IES=1, which has advantages in terms of tractability. More generally, in the standard utility formulation, it is impossible for all agents to have log utility (IES=1) when the γ_i differ across the agents.

As is well known, the Epstein-Zin/Weil form of recursive utility, based on Epstein and Zin (1989) and Weil (1990), allows for a separation between the coefficient of relative risk aversion and the IES. Typically, this benefit from EZW comes at the cost of analytical complexity, when compared with time-separable power utility. However, with heterogeneity in risk-aversion coefficients, the EZW specification allows for a simpler analysis. The key property of EZW is that it allows for high values of the γ_i that can differ across agents i, while maintaining values of the IES that are of reasonable magnitude and the same for each agent. More specifically, the EZW case admits the possibility of IES=1 for each agent, thereby gaining great simplifications from log utility. (The rate of time preference, ρ, is also assumed to be the same for each agent.)

Previous models of asset pricing with two types of agents distinguished by their coefficients of relative risk aversion include Dumas (1989), Wang (1996), Chan and Kogan (2002), Garleanu and Pedersen (2011), Longstaff and Wang (2012), Gennaioli, Shleifer, and Vishny (2012), and Caballero and Farhi (2014). These analyses assume time-separable power utility, augmented in Chan and Kogan (2002) to include an external habit in household utility. In Wang (1996), one agent has log utility and the other has square-root utility—coefficients of relative risk aversion of 1 and 0.5, respectively. In Garleanu and Pedersen (2011), one agent has log utility and the other has a coefficient of relative risk aversion greater than one. In the main analysis of Longstaff and Wang (2012), one agent has log utility and the other has squared utility—coefficients of relative risk aversion of one and two, respectively. Gennaioli, Shleifer, and Vishny (2012) assume that one agent is risk neutral and the other has infinite risk aversion, and Caballero and Farhi (2014) use an analogous setup.

Section I works out a baseline model that derives equilibrium holdings of equity claims and private bonds by the two types of agents, distinguished by their degrees of relative risk aversion. The analysis uses a tractable case with Epstein-Zin/Weil utility, where coefficients of relative risk aversion are "high" but utility is logarithmic in the sense that the intertemporal elasticity of substitution is one. The initial model with permanent differences in coefficients of relative risk aversion is non-stationary because, in the long run, the wealth share of the group with comparatively low risk aversion tends to approach one. We achieve stationarity by having agents continually replaced by new agents (possibly children) who are randomly assigned one of the two possible coefficients of relative risk aversion.

Section II carries out quantitative analyses based on specifications of the underlying parameters, including coefficients of relative risk aversion, the characteristics of the macro-disaster process, and the replacement rate. We focus on parameters that generate "reasonable" steady-state values of the risk-free interest rate (around 1.0% per year) and the unlevered equity premium (around 4.2%). When the replacement rate is 2% per year, the steady-state ratio of safe to total assets ranges up to 10%. We carry out dynamic analyses for two cases: the realization of a macroeconomic disaster and the experience of tranquility (no disasters) for 40 years.

Section III distinguishes further the gross quantity of private bonds from the net quantity corresponding to loans from group 2 to group 1. With infinitesimal transaction costs for paying interest and principal payments on bonds, agents in our model will not be simultaneously holding and issuing bonds. This condition pins down the equilibrium gross amount of safe assets.

Section IV introduces public debt. Added government bonds create more safe assets while simultaneously creating corresponding "safe liabilities" in the form of the present value of taxes. In the baseline setting, where the government and private sector are equally good at creating safe assets, Ricardian Equivalence holds, in the sense that changes in the quantity of government bonds do not affect the net quantity of safe assets or the risk-free rate. More specifically, the model predicts that an increase in government bonds by 1 unit crowds out private bonds by around 0.5 units. This prediction accords with some existing empirical evidence.

Section V discusses gold and other commodities that are often viewed as comparatively safe assets. Section VI relates the model's predictions on the quantity of safe assets to empirical estimates of this quantity. Section VII concludes.

I. Baseline Model

A. Structure and First-Order Conditions

The model is set up for convenience in discrete time. For some purposes, the period length can be viewed as a year. However, in analyzing the equilibrium, we think of the period length as approaching zero.

Agent i, for i=1, 2, has an Epstein-Zin/Weil (EZW) utility function, given by:

(1) $$(1-\gamma_i)\cdot U_{it} = \left\{C_{it}^{1-\theta} + \left(\frac{1}{1+\rho}\right)\left[(1-\gamma_i)E_t U_{i,t+1}\right]^{(1-\theta)/(1-\gamma_i)}\right\}^{(1-\gamma_i)/(1-\theta)}.$$

The coefficients of relative risk aversion satisfy $0<\gamma_1\leq\gamma_2$; that is, agent 1 is the comparatively low-risk-aversion agent. The IES, $1/\theta>0$, and the rate of time preference, $\rho>0$, are the same for the two agents. We simplify by having $\theta=1$ in the main analysis. This condition corresponds to the usual case of log utility and makes the price of equity independent of parameters that describe expected growth and uncertainty.[1] Another result with log utility is that the consumption of each agent, C_1 and C_2, equals ρ multiplied by a measure of each agent's assets.

Parts of the structure parallel Longstaff and Wang (2012). The single Lucas tree generates real GDP of Y_t in year t. This GDP is consumed by the two agents:

(2) $$C_{1t} + C_{2t} = Y_t.$$

Ownership of the tree is given by K_{1t} and K_{2t}, which add to full ownership, normalized to one:

(3) $$K_{1t} + K_{2t} = 1.$$

We use a convention whereby K_{it} applies at the end of period t, after the payment of the dividend, $K_{i,t-1}\cdot Y_t$, to agent i in period t. However, this timing convention is unimportant when the length of the period becomes negligible. The price of the tree in period t in units of consumables is P_t.

The stochastic process that generates Y_t corresponds to previous rare-disaster models, except for the omission of a normally-distributed business-cycle shock, which is quantitatively unimportant. The probability of disaster is the constant p per year. With probability 1-p, real

[1]This result means that the expected rate of return on equity, r^e, is independent of uncertainty parameters. Therefore, with $\theta=1$, all of the effects of uncertainty parameters on the equity premium work through the risk-free rate, r^f, rather than r^e. We know from previous analyses of this i.i.d. setting with a representative agent, such as Barro (2009), that the equity premium is independent of the parameter θ. Therefore, in this context, the setting of $\theta=1$ would not affect the model's implications for the equity premium.

GDP grows over one year by the factor 1+g, where g>0 is constant. With probability p, a disaster occurs and real GDP grows over one year by the factor $(1+g)\cdot(1-b)$, where b>0 is the size of a disaster. In the present simplified setting, disasters last for only one "period" and have a single size. The expected growth rate of GDP, denoted g*, is given as the length of the period approaches zero by

$$g^* \approx g - pb. \tag{4}$$

The analysis can be extended to allow for a time-invariant size distribution of disasters, as in Barro and Ursúa (2012). With more complexity, the analysis can be modified to allow disasters to have stochastic duration and cumulative size and to be followed by a tendency for recovery in the sense of above-normal growth rates (Nakamura, Steinsson, Barro, and Ursúa [2013]).[2] Other feasible extensions include time variation in the disaster probability, p (as in Gabaix [2012]), and the growth-rate parameter, g (as in Bansal and Yaron [2004]).

The baseline calibration specifies p=0.04 per year. This probability corresponds to the empirical frequency of disasters—defined as short-term declines in real per capita GDP by at least 10%—in a long-term panel of countries. The effective disaster size—in the sense of the single value that generates an equity premium corresponding roughly to the full size distribution of disasters—is set at b=0.32. The growth-rate parameter, intended to correspond to the non-disaster growth rate of real per capita GDP or consumption, is set at g=0.025 per year.

We focus on two forms of assets. Aside from unlevered equity claims, K_{it}, on the tree, we consider a non-contingent, one-period private bond, B_{it}. The quantity B_{it} is negative for a borrower (issuer of a bond) and positive for a lender (holder of a bond). Since the analysis

[2]The recovery tendency lowers the effective size, b, of a disaster. Therefore, for some purposes, we could allow for recoveries within the present framework by adjusting b.

assumes a closed economy, the total quantity of these private bonds, when added up across the two types of agents, is always zero:

(5) $$B_{1t} + B_{2t} = 0.$$

The model assumes a perfect private credit market, in the sense of ignoring possibilities of default and neglecting any transaction costs associated with interest and principal payments; that is, with collecting on loans.[3] In this case, bonds pay off at the risk-free interest rate for period t, denoted r_t^f. The amount of principal and interest received or paid on bonds by agent i in period t is $(1 + r_t^f) \cdot B_{i,t-1}$.

The menu of assets and financial contracts could be extended to include levered equity, structured finance, stock options, macro-disaster insurance, etc., and to have risk-free bonds of varying maturities. However, the setup with unlevered equity and one-period risk-free private bonds is sufficient to characterize the equilibrium in the present model.[4]

Each agent's budget constraint for period t is:

(6) $$C_{it} + P_t K_{it} + B_{it} = (Y_t + P_t)K_{i,t-1} + (1 + r_t^f)B_{i,t-1}\ .$$

The choice for period t of C_{it} and the portfolio allocation, (K_{it}, B_{it}), occur when Y_t, P_t, and r_{t+1}^f are known but Y_{t+1} and P_{t+1} are unknown.

Let R_t represent the gross return on any asset (equity, risk-free bonds) between periods t and t+1. This return equals $(Y_{t+1}+P_{t+1})/P_t$ for equity and $(1+r_{t+1}^f)$ for bonds. Each agent seeks to maximize expected utility, given in equation (1), subject to the budget constraint in equation (6)

[3]In the cases considered, debtors always have sufficient total assets to make the prescribed principal and interest payments on bonds.
[4]This result also holds in Wang (1996, p. 80) and Longstaff and Wang (2012, p. 3175).

and the levels of initial assets. The first-order optimization conditions for each agent and each type of asset can be expressed by means of a perturbation argument for periods t and t+1 as:

$$(7) \qquad (E_t U_{i,t+1})^{(\frac{\theta-\gamma_i}{1-\gamma_i})} = (\frac{1}{1+\rho}) E_t \left[U_{i,t+1}{}^{(\frac{\theta-\gamma_i}{1-\gamma_i})} \cdot \left(\frac{C_{i,t+1}}{C_{it}}\right)^{-\theta} \cdot R_{t+1} \right].$$

This expression simplifies in straightforward ways under log utility, θ=1.

Previous analyses (Giovannini and Weil [1989], Obstfeld [1994]), Barro [2009]) found in a representative-agent model with i.i.d. shocks (as in the present setting) that the realized utility, U_{t+1}, can be expressed as a positive constant multiplying $(C_{t+1})^{1-\gamma}/(1-\gamma)$. This result suggests looking for an approximate solution to the present two-agent model in which $U_{i,t+1}$ is a positive constant (possibly different for each agent) multiplying the analogous object for agent i, $\frac{(C_{i,t+1})^{1-\gamma_i}}{1-\gamma_i}$. When this condition holds, equation (7) can be rewritten as:

$$(8) \qquad \left[E_t (\frac{C_{i,t+1}}{C_{it}})^{1-\gamma_i} \right]^{(\frac{\theta-\gamma_i}{1-\gamma_i})} = (\frac{1}{1+\rho}) \cdot E_t \left[(\frac{C_{i,t+1}}{C_{it}})^{-\gamma_i} \cdot R_{t+1} \right].$$

When R_{t+1} equals the risk-free return, $1+r_{t+1}^f$, and θ=1, equation (8) implies

$$(9) \qquad 1 + r_{t+1}^f = (1+\rho) \cdot \frac{E_t (\frac{C_{i,t+1}}{C_{it}})^{(1-\gamma_i)}}{E_t (\frac{C_{i,t+1}}{C_{it}})^{-\gamma_i}}.$$

Thus, a key implication of the first-order conditions is that, in equilibrium, the right-hand side of equation (9) has to be the same for each agent; that is, for γ_1 and γ_2, respectively. In other words, the prospective paths of uncertain consumption levels for the two agents have to accord with the differences in the coefficients of relative risk aversion.

Under log utility, θ=1, each agent's consumption in period t is approximately the multiple ρ of that agent's resources for period t: [5]

[5] See Giovannini and Weil (1989, Appendix B).

(10) $$C_{it} \approx \rho \cdot [(Y_t + P_t) \cdot K_{i,t-1} + (1 + r_t^f) B_{i,t-1}] \, .$$

Equations (10) and (9) jointly determine agent i's choices of consumption, C_{it}, and portfolio allocation, (K_{it}, B_{it}).

Adding up equation (10) for the two agents and using the conditions from equations (2), (3), and (5)—total consumption equals GDP, equity holdings add to 1, and bond holdings add to zero—leads to

(11) $$P_t = Y_t \cdot (1-\rho)/\rho \approx Y_t/\rho,$$

where the approximation assumes that the length of the period is negligible. Thus, under log utility, the equity price and, hence, the value of total assets is independent of parameters related to expected growth and uncertainty and the degree of risk aversion. This result implies that the expected rate of return on equity, r^e, is the dividend yield, ρ, plus the expected rate of capital gain, which equals g*, the expected growth rate of GDP and consumption:

(12) $$r^e \approx \rho + g^* \approx \rho + g - pb,$$

where g* is given in equation (4).

B. Market Equilibrium

To assess the market equilibrium, we think of agent 1's share of total wealth as the single state variable for each period. Agent i's wealth at the end of period t-1 is

$$W_{i,t-1} = P_{t-1}K_{i,t-1} + B_{i,t-1},$$

so that agent 1's wealth share at the end of period t-1 is

(13) $$\frac{W_{1,t-1}}{W_{t-1}} = K_{1,t-1} + \frac{\rho B_{1,t-1}}{(1-\rho) Y_{t-1}} \, .$$

(Note that total wealth, W_{t-1}, equals the equity price, P_{t-1}.)

The analysis requires an initial value for agent 1's wealth share. For example, we could assume that this share starts at 0.5 in period 0. Heuristically, if $\gamma_1<\gamma_2$, there is an incentive in the initial position for agent 1 to issue risk-free bonds, so that $B_{11}<0$ in period 1, and these bonds will be held by agent 2, so that $B_{21}>0$. That is, agent 1 borrows from agent 2 on a safe basis. Correspondingly, agent 1 uses its bond issue to increase its share of equity, so that $K_{11}>0.5$ and $K_{21}<0.5$. In a richer model, this process of safe credit creation would affect the equilibrium amount and composition of investment.

The pattern of bond and equity positions shifts risk from the high-risk-aversion agent 2 to the low-risk-aversion agent 1. However, the process does not entail complete risk shifting; rather, enough bond issue occurs so that the resulting stochastic paths of future consumption for each agent make the right-hand side of equation (9) the same for each agent i. This equation also determines r_1^f.

We can use equations (9) and (10), along with the agents' budget constraints, to find numerically the equilibrium values for period 1 of r_1^f, each agent's consumption, and each agent's allocation of assets between equity and bonds.[6] The realization for Y_1 (disaster or no disaster in the present case) then determines each agent's wealth at the end of period 1 and, hence, agent 1's wealth share at the end of period 1. This share is the state variable that determines the equilibrium values for period 2, and so on.

Using the budget constraint from equation (6) and the condition for consumption in equation (10), we can show that agent 1's wealth share at the end of period t relates to asset holdings from the end of period t-1 in accordance with

[6]We carried out this analysis numerically using periods of negligible length.

(13) $$\frac{W_{1t}}{W_t} = K_{1,t-1} + \frac{\rho(1+r_t^f)B_{1,t-1}}{Y_t}.$$

We can also show that the change in agent 1's wealth share from t-1 to t is

(14) $$\frac{W_{1t}}{W_t} - \frac{W_{1,t-1}}{W_{t-1}} \approx \frac{\rho B_{1,t-1}}{Y_t}(r_t^f - \rho - g_t),$$

where $g_t \equiv (Y_t/Y_{t-1} - 1)$ is the stochastic growth rate of GDP. Since the risk-free rate, r_t^f, will be less than $\rho + g^*$, which equals the expected rate of return on equity, r^e, the expectation of the right-hand side of equation (14) is positive if $B_{1,t-1} < 0$. In other words, the expected change in agent 1's wealth share is positive whenever agent 1 (the low risk-aversion agent) is borrowing in a risk-free manner from agent 2. The reason that agent 1's wealth share tends to rise over time is that this agent's wealth is relatively concentrated in risky equity, which has a higher expected rate of return than risk-free bonds. (Note that this calculation factors in the occasional macroeconomic disasters, which tend to reduce agent 1's wealth share.) Consequently, we can show that agent 1's wealth share asymptotically approaches one, and the ratio of risk-free bonds, B_{1t}, to total assets or GDP asymptotically approaches zero. In effect, there is a selection or survival effect, whereby wealth is concentrated asymptotically in the agent with relatively low risk aversion. Hence, the model behaves in the long run like a representative-agent economy with coefficient of relative risk aversion equal to γ_1.

The non-stationarity of the initial form of the model makes it unsatisfactory for studying the determination of wealth shares and quantities of safe assets.[7] To get a satisfactory analysis, the model has to be modified to achieve stationarity; in particular, to have the (expected) wealth share of agent 1 asymptotically approach a value less than one.

[7]This problem applies also to previous models with heterogeneous coefficients of relative risk aversion; see, for example, Longstaff and Wang (2012, p. 3208).

C. Replacement and Stationarity

A natural way to achieve stationarity is to have agents die off randomly, with replacement by new agents who have a random (50-50) chance of being type 1 or type 2; that is, having a coefficient of relative risk aversion of γ_1 or γ_2.[8] Since type-1 agents tend to have above-average wealth, this process tends to redistribute wealth back to type-2 agents.

The replacement agents might inherit the assets of their altruistic predecessors, who might be parents. Alternatively, as in Blanchard (1985), agents may leave no bequests and hold all of their bond-like assets as annuities, on which the returns factor in the probability of dying off. In either case, a full analysis requires the optimizing choices of consumption and asset holdings to take account of the possibility of dying with replacement by children whom one may or may not care about.

When agents are linked to their descendants via operative intergenerational transfers, the main effect in the model from death and replacement is the random change in the coefficient of relative risk aversion. Therefore, in this context, we can use a simpler metaphor in which no one literally dies off but where people randomly experience moments in which shifts occur in their coefficients of relative risk aversion to either γ_1 or γ_2. Each destination could have a 50-50 chance of being picked. This structure avoids dealing with inheritance and altruism but implies that optimizing decisions on consumption and asset holdings depend on the potential for future shifts in one's own attitude toward risk. Our conjecture is that, if the replacement rate, which we denote by ν, is very high, this model would function like a representative-agent model in which the coefficient of relative risk aversion is an average of γ_1 and γ_2. However, if ν is low, as we

---

[8]Chan and Kogan (2002) generate stationarity effectively by having each agent's coefficient of relative risk aversion, γ_i, be an increasing function of that agent's wealth share. The assumed sign of this effect is not obvious; that is, it is unclear that richer agents would have higher coefficients of relative risk aversion. In any event, this type of model functions in the steady state as a representative-agent model with a single coefficient of relative risk aversion.

assume, the main element in each agent's current optimization conditions would be the incumbent value of γ_i. That is, the optimizing conditions would be close to those already derived.[9]

Let ν≥0 be the rate at which agents are replaced in the sense of possibly changing their values of γ_i between the two alternatives. This process is stochastic at the individual level but roughly deterministic in the aggregate. We retained the first-order conditions derived from the model without replacement but modified the equilibrium analysis to factor in the shifting of wealth composition across the two types of agents. Specifically, the expression for the change in wealth share of agents of type 1 is modified from equation (14) to:

$$\text{(15)} \qquad \frac{W_{1t}}{W_t} - \frac{W_{1,t-1}}{W_{t-1}} \approx \frac{\rho B_{1,t-1}}{Y_t}\left(r_t^f - \rho - g_t - \nu\right) - \nu \cdot (K_{1,t-1} - 0.5).$$

Equation (15) takes account of how wealth is shifted between type-1 and type-2 agents when agents change their type. However, this shifting also means that individuals within the two groups will have to be heterogeneous in wealth. Typically, someone who just moved from type 1 to type 2 will have wealth above the mean of the existing type-2 agents, and vice versa for someone who just changed from type 2 to type 1. As this process evolves, the agents within groups will have a range of wealth levels, depending on their history of past transitions (and on how these changes related to the aggregate experience with regard to realizations of disasters).

The wealth distribution within groups might be interesting to analyze but is unimportant for the present analysis. For our purpose, what matters is the total wealth held by agents of each type and not the distribution of wealth within types. In particular, equations (9) and (10) imply that, for a given γ_i, higher wealth scales up proportionately the chosen values of consumption,

[9]Under log utility (θ=1), each agent's consumption is the multiple ρ of that agent's resources. This condition does not involve γ_i and would not change when γ_i has the potential to shift in the future. However, the first-order conditions in equation (9) would be affected.

without changing the proportionate amounts held of risky and risk-free claims. Thus, we can use equation (15) to gauge the changing wealth shares of groups 1 and 2, while neglecting effects from the differing levels of wealth within each group.

If ν=0, as before, the expectation of the right-hand side of equation (15) is positive if $B_{1,t-1}<0$. As agent 1's wealth share approaches one, $B_{1,t-1}/Y_t$ asymptotically approaches zero and, therefore, equation (15) implies that the expectation of the change in agent 1's wealth share asymptotically approaches zero. Another property of the equilibrium with ν=0 is that $K_{1,t-1}$ asymptotically approaches 1.

If ν>0, when $K_{1,t-1}$ is close to 1 and $B_{1,t-1}/Y_t$ is negligible, the term on the far right of equation (15) is negative and dominates in magnitude the first term on the right. It follows that the expected change in agent 1's wealth share will reach zero before $K_{1,t-1}$ gets close to one and $B_{1,t-1}/Y_t$ becomes negligible. For given values of ν and the other parameters, we find numerically the (unique) wealth share that makes the expectation of the right-hand side of equation (15) equal to zero. We think of this wealth share as the (stochastic) steady-state value, and we compute the associated steady-state values of r_t^f, $K_{1,t-1}$, and $B_{1,t-1}/Y_t$. This analysis should yield a satisfactory approximation to the steady-state equilibrium when the replacement rate, ν, is small.

II. Quantitative Analysis of Stationary Model

Aside from the coefficients of relative risk aversion, γ_1 and γ_2, the baseline parameter values, listed in the notes to Table 1, are ρ=0.04 per year (rate of time preference), g=0.025 per year (growth-rate parameter), p=0.04 per year (disaster probability), and b=0.32 (effective disaster size). These values accord with the prior empirical analysis summarized in Barro and Ursúa (2012). These parameter values imply from equation (4) that the expected growth rate is

(16) $$g^* = g - p \cdot b = 0.0122 \text{ per year.}$$

We also assume log utility, θ=1, and begin with a replacement rate of ν=0.02 per year. This value corresponds roughly to adult mortality rates.

A. A Representative Agent

Table 1 considers a representative agent, where $\gamma_1=\gamma_2=\gamma$. In these cases, if we start with agent 1's wealth share at 0.5, B_{it} and K_{it} stay constant over time at 0 and 0.5, respectively, irrespective of the realizations of Y_t. Because of log utility, the expected rate of return on equity, r^e, is fixed at ρ+g*, where ρ=0.04 per year and g*=0.0122 (equation [16]), so that r^e=0.052 per year. A higher γ lowers the risk-free rate, r^f, and, thereby, raises the equity premium. Specifically, Table 1 shows that r^f ranges from 0.046 at γ=1 to -0.023 at γ=5 and -0.064 at γ=6.[10] An unlevered equity premium between 0.03 and 0.06 (corresponding roughly to historical data) requires γ to be between 3 and 4.5. For a given γ, r^f is fixed over time, regardless of the realizations of Y_t. This risk-free rate is a shadow rate in the sense that no risk-free borrowing and lending occur in equilibrium. That is, no safe assets are created in this representative-agent environment.

B. Heterogeneity in Risk Aversion

Table 2 allows for differences between γ_1 and γ_2. We show combinations of γ_1 and γ_2 that generate a (stochastic) steady-state risk-free interest rate of r^f=0.010 and, hence, a steady-state unlevered equity premium of 0.042. That is, these combinations of γ_1 and γ_2 accord roughly with empirically observed averages of risk-free returns and the equity premium. The table shows the corresponding (stochastic) steady-state values of the amount of safe assets, $|B_1|$, expressed relative to total assets or GDP, agent 1's share of risky assets, K_1, and agent 1's

[10] In the present model (which lacks risk-free and costless storage of final product), there is nothing special about a risk-free rate of zero.

wealth share, W_1/W (which equals agent 1's share of total consumption). Because economy-wide assets equal annual GDP times 25 ($1/\rho$) in this model, the amount of safe assets expressed relative to annual GDP is 25 times the ratio to total assets. Note that total assets correspond to the capitalization of the entire flow of GDP, effectively including human capital as well as physical capital.

The first row of Table 2 shows that, if $\gamma_1 = \gamma_2$, the value of γ_1 and γ_2 needed to generate r^f=0.01 is 3.78 (see Table 1). Columns 1 and 2 of Table 2 show that values of γ_1 below 3.78 require values of γ_2 above 3.78 to generate r^f=0.010. For example, γ_1=3.5 matches up with γ_2=4.2, γ_1=3.25 with γ_2=4.8, and γ_1=3.0 with γ_2=6.7. For still lower values of γ_1, the required value of γ_2 explodes, and there is a lower bound on γ_1 a little below 2.7 such that even an infinite γ_2 would not generate r^f=0.010. For example, γ_1=2.68 requires γ_2 =647, which is essentially infinity. Thus, given ν=0.02 per year, the model places a lower bound around 2.7 on the admissible γ_1.

In column 5, the steady-state share of risky assets held by agent 1, K_1, equals 0.50 when $\gamma_1=\gamma_2$, then rises toward 1.0 as γ_1 falls and γ_2 rises. When γ_1=2.68 and γ_2=647, K_1 is very close to 1.0.

Column 6 shows that the steady-state ratio of the magnitude of safe to total assets rises from 0 when $\gamma_1=\gamma_2$ to 2.5% when γ_1=3.50 (γ_2=4.2), 5.4% when γ_1=3.25 (γ_2=4.8), 9.5% when γ_1=3.0 (γ_2=6.7), and 15.4% when γ_1=2.75 (γ_2=24.4). The upper bound for this ratio around 18% applies as γ_1 approaches its lower bound a little below 2.7 and γ_2 approaches infinity. For subsequent purposes, we are particularly interested in the model's predictions about the size of safe assets in relation to total assets or GDP. From this perspective, an important result is that the predicted quantity of safe assets is less than 10% of economy-wide assets if γ_2 is less than 6.7, which is a high degree of relative risk aversion.

In column 7, the steady-state wealth share, W_1/W, starts at 0.50 when $\gamma_1=\gamma_2$, then rises toward an upper bound of 0.82 as γ_1 falls and γ_2 rises. This wealth share equals the fraction of total consumption received by agent 1.

Table 3 redoes the analysis for alternative settings of three of the parameters: the replacement rate, ν, is raised from 0.02 to 0.05 per year; the disaster probability, p, is lowered from 0.04 to 0.02 per year; and the rate of time preference, ρ, is decreased from 0.04 to 0.02 per year. In each case, the table shows steady-state values of r^f and the other variables for three of the combinations of (γ_1, γ_2) considered in Table 2.

In the upper panel of Table 3, where $\nu=0.05$, the risk-free rate, r^f, is somewhat below the value 0.010 from Table 2 (because the higher risk-aversion coefficient, γ_2, gets relatively more weight when ν is higher). The share of group 1 in risky assets and wealth is correspondingly reduced by the higher ν. For example, when $\gamma_1=3.00$ and $\gamma_2=6.66$, K_1 is 0.70, rather than 0.77, and W_1/W is 0.59, rather than 0.67. If γ_2 is constrained to be less than 6.7, the ratio of safe to total assets can be as high as 11.5% when $\nu=0.05$, compared to 9.5% when $\nu=0.02$. Overall, the most important finding is that the results do not change greatly when ν is 0.05, rather than 0.02.

In the middle part of Table 3, where p=0.02, the risk-free rate, r^f, is sharply higher than the value 0.010, which applies to p=0.04 in Table 2. Correspondingly, the equity premium becomes much too low in Table 3, compared with empirically observed averages. Thus, as in previous research, the model does not accord with regularities on mean rates of return unless the disaster risk is sufficiently high. A similar conclusion arises if the disaster size, b, is lowered substantially below its initial value of 0.32.

In the model, the ratio of the equity price, P_t, to GDP, Y_t (that is the price-dividend ratio), equals $1/\rho$ and is, therefore, independent of parameters related to expected growth, uncertainty, and risk aversion. This result depends on log utility (θ=1)—more generally, P_t/Y_t would depend on parameters other than ρ, though the signs would be determined by whether θ was less than or greater than one. In the context of θ=1, the effect from a shift in ρ is of interest because it changes P_t/Y_t, which equals the ratio of wealth to GDP. In particular, we can assess whether this change affects mostly the ratio of safe to total assets (wealth) or the ratio of safe assets to GDP.

The bottom part of Table 3 sets ρ=0.02 and thereby raises the wealth-GDP ratio to 50, compared to 25 in Table 2, where ρ=0.04. The decrease in ρ sharply lowers r^f—for example, when γ_1=3.00 and γ_2=6.66, r^f is -0.009, compared to 0.010 in Table 2. Although the equity premium implied by the lower ρ in Table 3 is similar to that found in Table 2, the overall level of rates implied by ρ=0.02 is counter-factual. Agent 1's shares of equity and wealth when ρ=0.02 (Table 3) are similar to those found when ρ=0.04 (Table 2).

The most interesting results from the change in ρ concern the amount of safe assets expressed relative to total assets or GDP. Table 3 shows that the ratio of safe to total assets when ρ=0.02 is slightly higher compared to that when ρ=0.04 (Table 2). For example, when γ_1=3.00 and γ_2=6.66, the ratio of safe to total assets is 10.5% when ρ=0.02 compared to 9.5% when ρ=0.04. In contrast, the ratio of safe assets to GDP more than doubles—it equals 5.2 when ρ=0.02 compared to 2.4 when ρ=0.04. In more general contexts, such as when $\theta \neq 1$, the wealth-GDP ratio would be sensitive to changes in various parameters, such as the expected growth rate, the probability and size distribution of disasters, and the levels of risk-aversion coefficients. Our results from the log-utility case suggest that the ratio of safe to total assets may not change greatly when the wealth-GDP ratio changes. That is, the ratio of safe assets to GDP may roughly move along with the ratio of total assets (wealth) to GDP. We use these results later when attempting to relate our theoretical findings to empirical measures of the quantity of safe assets.

C. Dynamics

The dynamics of the economy reflects the evolution of the single state variable, which is the share of agent 1 in total wealth, W1/W. Disaster shocks and long periods free of disasters affect this state variable and, thereby, have persisting influences on the risk-free interest rate, r^f, the ratio of safe to total assets, and other variables. We consider first the dynamic effects from a disaster and then examine the consequences from a long period free of disasters.

1. Aftermath of a Disaster. Figure 1 shows the dynamics of the economy starting from a steady state and assuming the realization of a disaster of size b=0.32 in period 0. (Periods in this exercise are extremely short.) The results correspond to the parameter combination (γ_1=3.00, γ_2=6.66) in Table 2. The paths of variables in Figure 1 assume no further disasters and are, therefore, deterministic. The variables considered over ten years are agent 1's wealth share, W1/W (which equals agent 1's consumption share), the risk-free interest rate, r^f, agent 1's share of total equity, K1, and the ratio of the magnitude of safe assets, B1, to total assets.

Because of agent 1's relatively high concentration in risky assets, this agent's wealth share falls with the disaster from 0.672 to 0.627. The share rises thereafter (in the absence of further disasters) but remains below the steady-state value even after 10 years, when the share reaches 0.656. Another way to look at this pattern is that relatively low inequality of wealth and consumption persist for a long time after a disaster shock. However, the recovery toward the steady state is accompanied by rising inequality. These patterns also appear in agent 1's share of equity, K1. This share falls on impact from its steady-state value of 0.767 to 0.732, then rises to 0.755 after 10 years.

For the risk-free rate, r^f, we can view the disaster shock and consequent shift in relative wealth toward agent 2 as raising the demand for bonds (from agent 2) compared to the supply (from agent 1). In response to the shift in excess demand, r^f falls on impact from its steady-state value of 0.0100 to 0.0073. That is, the disaster leads to a low risk-free interest rate. In the recovery period, r^f rises but remains below its steady-state value. After 10 years, r^f reaches 0.0091.

The enhanced wealth share of agent 2 is accompanied on impact by a rise in the ratio of the magnitude of safe assets, B1, to total assets. This ratio increases initially from its steady-state value of 0.095 to 0.104. Thus, safe assets are comparatively large immediately after a disaster. The ratio then falls gradually and reaches 0.098 after 10 years.

To summarize, disasters generate low but rising wealth and consumption inequality, low but rising risk-free real interest rates, and high but declining ratios of safe to total assets. In particular, low inequality and risk-free interest rates and high safe-asset ratios are all symptoms of a gradual recovery from a serious adverse shock to the economy.

An important aspect of the disaster shock that we examined is that it disproportionately affects the low-risk-aversion agent, group 1, and, therefore, shifts the wealth share initially toward the high-risk-aversion agent, group 2. This pattern arises because the shock affects the value of equity, which is disproportionately held by group 1. Hart and Zingales (2014) argue that this kind of pattern characterizes some macro-financial shocks, such as the bursting of the Internet boom in 2000. They argue, however, that other shocks—notably the Great Recession of 2007-2009—feature the erosion in value of assets that were previously viewed as nearly safe. In the 2007-2009 case, this pattern applied particularly to claims associated with real estate, whose safety had been greatly exaggerated.

In our model, we could analyze the Hart-Zingales case by allowing for an unexpected decline in the value of the existing "safe" assets, which are the private bonds. That is, the zero-probability event of large losses on safe assets could be viewed as a one-time happening. In this case, agent 1's wealth share would initially shift discretely above its steady-state value.[11] The subsequent dynamics relates to that described in our next example.

2. Forty years of tranquility. Figure 2 deals with a scenario in which, starting from the steady state, the economy has a long period with no disasters ("40 years of tranquility"). This situation accords broadly with the U.S. experience from the 1950s up to the Great Recession of 2007-2009. The parametric assumptions for Figure 2 are the same as those for Figure 1.

In Figure 2, agent 1's wealth share rises gradually above its steady state value of 0.67. Conditional on no disasters, this ratio rises after 40 years to 0.71—and would asymptotically approach a higher value, which turns out to be 0.72, if no disaster ever occurred. The value 0.72 is a kind of steady-state wealth share (shown in quotes in the figure) in that it applies asymptotically conditional on the realization of no disasters. In contrast, the lower steady-state wealth share of 0.67 is defined inclusive of the occasional incidence of disasters.

The dynamic path of the wealth share in Figure 2 shows that sustained tranquility is accompanied by rising inequality, in the sense of growing wealth and consumption shares of group 1. The dynamics also features a rising risk-free rate, which increases above its steady-state value of 0.010 and eventually approaches 0.0125. The ratio of safe to total assets falls from its steady-state value of 0.095 and gradually approaches 0.084.

[11] A counter-vailing force in the real estate bust is that large financial institutions, including Lehman, experienced sharp losses in asset values. This aspect of the shock tends to lower group 1's wealth share and, thereby, works more like the disaster realization that was already analyzed.

In the paths shown in Figure 2, agent 1's wealth share would never rise above the "steady state" value of 0.72. However, a shock mentioned before—where the value of safe assets declines sharply because this safety had been exaggerated—could put agent 1's wealth share above its "steady-state" value of 0.72. In that case, the post-shock dynamic paths (conditional on no further disasters) would feature a gradually declining wealth share of agent 1, with this share asymptotically approaching from above the value 0.72. Correspondingly, the risk-free rate would rise initially above its "steady-state" value and then fall gradually, whereas the ratio of safe to total assets would fall initially below its "steady-state" value and then rise gradually.

III. Gross versus Net Lending

The bond holdings, B_1, shown in Table 2 correspond to net safe lending from the high-risk-aversion agent, group 2, to the low-risk-aversion agent, group 1. There is a sense, however, in which gross bond issuance is not pinned down, because the model would admit unlimited borrowing and lending within groups. That is, agent 1 could effectively issue an arbitrary amount of bonds to himself, and analogously for agent 2.

If the model were augmented to include an infinitesimal amount of transaction costs for bond issuance or collection of interest and principal, then borrowing and lending within groups would not occur in equilibrium in the present model. In this case, the quantity of bonds, B_1, shown in Table 2 would be the unique equilibrium for the gross amount outstanding.

If transaction costs associated with bonds are substantial, the quantity of net bond issuance and the risk-free rate might differ significantly from the values shown in Table 2. Moreover, the risk-free rate received by lenders (group 2) would deviate from that paid by borrowers (group 1). For example, if transaction costs were prohibitive, the results would correspond to autonomy for groups 1 and 2 and, therefore, to the results shown in Table 1. The quantity of net bond issuance would be 0, and the share of capital held by each group would be 0.5. As an example, if γ_1=3.0, the shadow risk-free rate for group 1 would be 0.024 (from Table 1) and if γ_2=5.0, the shadow risk-free rate for group 2 would be -0.023 (again from Table 1). That is, members of group 1 would be willing to pay a rate of 0.024 per year at the margin on risk-free borrowing, whereas members of group 2 would be willing to accept a rate of -0.023 per year at the margin on risk-free lending. However, no issue of safe debt occurs because of the prohibitive transaction costs.

IV. Government Bonds and Ricardian Equivalence

Suppose that the government issues one-period bonds with characteristics corresponding to those of private bonds. The interest rate on government bonds held from year t to year t+1 must then be r_{t+1}^f, the same as that on private bonds. The simplest approach is for the government to make a distribution in the form of lump-sum transfers of bonds in year t in the aggregate quantity B_t^g. This distribution is assumed to go 50-50 to members of groups 1 and 2. The aggregate principal and interest, $(1 + r_{t+1}^f)B_t^g$, is paid out to government bondholders in period t+1. This payout is financed by lump-sum taxes, levied equally in period t+1 on members of groups 1 and 2.

What is the impact of this government bond issue on private bond issue, the risk-free interest rate, and so on? The government bond issue does not affect the households' first-order conditions involving the risk-free rate, r_{t+1}^f, which appear in equation (9). There is also no effect on households' budget constraints in equation (6) (updated to apply to periods t and t+1), once one factors in the transfer payments in year t and the taxes levied in year t+1. Therefore, it is immediate that the equilibrium involves the same net borrowing and lending as before between groups 1 and 2, the same risk-free interest rate, r_{t+1}^f, the same equity price, $P_t=Y_t/\rho$, and the same expected rate of return on equity, $r^e = \rho + g^*$. That is, the equilibrium features Ricardian Equivalence with respect to (net) quantities of safe assets and the various rates of return.

There are multiple possibilities with respect to holdings of government bonds that support the equilibrium. One possibility is that groups 1 and 2 each hold 50% of the government bonds issued in year t. These quantities correspond to the present value of the (certain) tax liabilities imposed on each group. That is, government bonds and tax liabilities cancel with respect to creating safe assets; no safe assets are created on net. The quantity of net private borrowing and lending, corresponding to B_{1t}, is then the same as before, and the same risk-free interest rate supports this equilibrium. This outcome corresponds to the upper part of Table 4.

Alternatively, all of the government bonds could be held, in equilibrium, by group 2. Members of group 2 then have additional safe assets (in the form of government bonds) that exceed the present value of their added tax liabilities (which equal 50% of the government bonds). Correspondingly, members of group 1 hold no government bonds but have a present value of tax liabilities also equal to 50% of the bonds. Therefore, group 1 has additional tax liabilities (50% of the government bonds) that exceed their added safe assets in the form of government bonds (which are nil in this case). This result works if the net private borrowing by group 1 from group 2 falls by 50% of the government bonds that were issued. In this case, the net position of group 1 with respect to group 2 is the same as before. The only difference is that some of the borrowing and lending between the groups is purely private, while some works through the government as intermediary (collecting taxes from group 1 and using the proceeds to pay principal and interest on half of the government bonds held by group 2). This result corresponds to the lower part of Table 4.

As in the previous section, the indeterminacies with respect to gross debt are resolved if there is an infinitesimal amount of transaction costs for bond issuance or collection of interest and principal. These transaction costs are assumed at this point to be the same for private and public bonds. In this case, the unique equilibrium will be the one described in the lower part of Table 4, where group 2 holds all of the added government bonds, and the issue of government bonds crowds out private bonds with a multiplier of 0.5. (This result assumes that the gross quantity of private bonds outstanding was initially greater than 50% of the added government bonds.)

The finding of Ricardian Equivalence may not be surprising, but what is surprising is that this equivalence associates with a crowding-out coefficient for private bonds with respect to public bonds of -0.5, not -1.0. This result came from a model with a number of simplifying assumptions; notably, there were just two groups characterized by their coefficients of relative risk aversion, γ_i, and the incidence of the present value of taxes net of transfers associated with the government bond issue was the same for each group. However, the crowding-out coefficient around -0.5 does not depend on these assumptions holding precisely. For example, the restriction to two groups is unimportant.[12] The main assumption that seems to matter is that there is little correlation across groups between γ_i and the share of taxes net of transfers applying to group i.

[12] Suppose, for example, that there are four groups of agents, where $\gamma_1<\gamma_2<\gamma_3<\gamma_4$. Suppose further that the initial equilibrium involves private bond holdings of B_1=-100, B_2=-50, B_3=50, and B_4=100. Assume that the government issues 4 units of bonds, with the present value of taxes rising by 1 unit for each group. In this case, the two private borrowers go, in equilibrium, to B_1=-99 and B_2=-49, thereby preserving their positions for bond holdings net of tax liabilities (of 1 each) at -100 and -50, respectively. The two private lenders go, in equilibrium, to overall bond positions (inclusive of government bonds) of B_3=51 and B_4=101, thereby preserving their positions for bond holdings net of tax liabilities at 50 and 100, respectively. Note that the additional 4 units of government bonds crowd out the total of private bonds by 2 units.

The model's predicted crowding-out coefficient relates to the study by Krishnamurthy and Vissing-Jorgensen (2013, p. 1), who argued "that government debt … should *crowd out* the net supply of privately issued short-term debt." They tested this hypothesis on U.S. data for 1914-2011 and found (Table 4, Panel A) that an increase in the quantity of net U.S. government debt had a significantly negative effect on the net short-term debt created by the private financial sector. Remarkably, their estimated coefficient was close to -0.5, the value predicted by our model.[13] Similarly, Gorton, Lewellen and Metrick (2012, Table 1) estimated a crowding-out coefficient close to -0.5 for a broad concept of private financial-sector liabilities (their "high estimate") in the United States for 1952-2011.[14]

Ricardian Equivalence may not hold exactly in our model if the government is superior to the private sector in the technology of creating safe assets.[15] In particular, the government might be able to commit better than private agents to honoring payments of principal and interest on its bonds and can also use the coercive power of the tax system to ensure the financing of these payments. On the other hand, a private lending arrangement requires only that group 1 make principal and interest payments in period t+1 to group 2, whereas the public setup entails the government collecting taxes in period t+1 from group 1 and then using the proceeds to pay off

[13] Krishnamurthy and Vissing-Jorgensen (2013, p. 23) say: "These results suggest that a one-dollar increase in Treasury supply reduces the net short-term debt issued by the financial sector by 50 cents." In their theory (Section 3), they derive a crowding-out hypothesis from a model in which Ricardian Equivalence fails. However, they do not mention that their empirical results are consistent with a model in which Ricardian Equivalence holds.

[14] Gorton, Lewellen and Metrick (2012, p. 103) say: "These results suggest that financial liabilities and government liabilities may be substitutes."

[15] Caballero and Farhi (2014, p.3) make this assumption, although they do not clarify the elements that underlie the government's superior technology: "Public debt … plays a central role … as typically the government owns a disproportionate share of the capacity to create safe assets while the private sector owns too many risky assets. … The key concept then is that of *fiscal capacity*: How much public debt can the government credibly pledge to honor should a major macroeconomic shock take place in the future?" They also do not consider that public debt issue creates additional "safe liabilities" in the form of taxes that match the added safe assets in a present-value sense.

group 2. Once the distorting influences from taxation are considered, it is not obvious that the public process entails lower "transaction costs" overall.[16]

An additional consideration is that the expansion of public debt and the associated taxation are poorly targeted: 50% of the added government bonds—held by group 2—match the added present value of tax liabilities for this group and, therefore, do not serve to shift risk toward group 1. Only the remaining 50% of government bonds corresponds to this shifting of risk. In contrast, all private bonds issued by group 1 and held by group 2 associate with risk shifting.

To highlight a case where the issue of public debt is important, suppose that the private sector's technology for creating safe assets is so poorly developed that no issue of private bonds occurs in the initial equilibrium (where no government bonds exist). In this case, analyzed at the end of Section III, groups 1 and 2 are effectively autonomous, and the equilibrium for each group is the one that would apply in the corresponding representative-agent economy. The risk-free interest rate for group 1 can then diverge substantially from that for group 2.

In this environment, the government's issue of bonds can substitute for the private lending that would have occurred if the private sector had possessed the technology to create safe assets. (The assumption here is that the government can use the tax system to generate these safe assets.) In this setting, Ricardian Equivalence fails, and the government's debt issue moves the economy toward a more efficient outcome, where risk is shifted from type-2 to type-1 agents and the risk-free interest rates of the two groups converge.

[16]Even if the interest rate on government bonds is lower than that on private bonds, the overall transaction costs—including the distorting effects from taxation—associated with the public process might exceed that for the private process.

We can assess how much public debt is required to get the economy into the equilibrium where private debt was available (with negligible transaction costs). The answer—related to the -0.5 crowding-out coefficient discussed before—is that the required quantity of public debt is twice the level of private bonds that arose in the initial setting. Moreover, if public debt expands beyond this quantity, it has no further effect on the equilibrium. That is, Ricardian Equivalence holds in this range at the margin even though private bonds are assumed to be absent.

V. Gold and other Commodities as Safe Assets

Gold and other precious durable commodities, such as silver and platinum, are often viewed as forms of comparatively safe assets. However, as noted in Barro and Misra (2013), real returns on gold are as volatile as stocks in the period since 1975, although real gold returns have covariances with growth rates of GDP and consumption that may be negative and are surely much smaller than those for stock returns.

In the model, the underlying economy-wide risk corresponds to the uncertainty in the process for GDP, Y_t, which corresponds to the fruit from the Lucas tree. Any application of this model to macroeconomic data would require Y_t to be identified with a composite of a variety of goods and services. In this sense, the services from gold or other commodities would constitute forms of goods and services that are already factored into the aggregate variable Y_t. Then, depending on how the values of these commodities covary with the rest of GDP, the variable Y_t inclusive of the commodity service flows might be more or less volatile than the variable gauged exclusively of these flows. In any event, the analysis would be the same as that already carried out. In particular, the determinants of the risk-free rate and the net quantity of safe assets would still follow the form of the analysis in Table 2.

VI. The Quantity of Safe Assets

In the model, the net quantity of safe assets is well defined and corresponds to the shifting of risk from the high-risk-averse agent, group 2, to the low-risk-averse agent, group 1. Table 2 shows that, for reasonable parameter values, the steady-state quantity of safe assets ranges up to 10% of total assets.

Using data to match the model's predictions for the quantity of safe assets is challenging because it is unclear how to measure the amounts of these assets. Gorton, Lewellen, and Metrick (2012) (henceforth, GLM) define safe assets to comprise mostly liabilities of the government and the private financial sector. After making a number of adjustments—for example, to eliminate U.S. government securities held by federal trust funds and to deduct 15% of long-term debt issued by the financial sector—they focus on a "high estimate" of the amount of safe assets. Using U.S. data from 1952 to 2010, GLM report two major findings. First, the ratio of their measure of safe assets to a concept of total assets remained relatively stable over time, ranging between 30% and 35%. Second, the ratio of total assets to annual GDP rose sharply from around 4.0 in 1952 to about 10.3 in 2010. By implication, the ratio of safe assets to GDP rose from about 1.4 in 1952 to around 3.3 in 2010.

When comparing with our theoretical predictions, the first observation is that the average ratio of safe to total assets computed by GLM—30-35%--exceeds the steady-state values predicted by our model—on the order of 10%. However, the stability of the ratio found by GLM accords with our model. Notably, we found in Table 3 that a change—a reduction in the rate of time preference, ρ—that sharply raised the ratio of total assets to GDP was consistent with rough stability in the ratio of safe to total assets. In the log-utility version of the model, the observed sharp increase in the ratio of total assets to GDP from 1952 to 2010 could be explained only by a reduction in the rate of time preference, ρ. However, in extended versions of the theory, the change could also derive from shocks to expected growth, uncertainty, and the extent of risk aversion.

One reason that GLM's measured ratio of safe to total assets would diverge from our theoretical concept concerns the denominator, total assets. In our theory, total assets comprise the discounted value of GDP, which equals consumption. In effect, the theoretical measure of total assets includes human capital as well as physical capital. In contrast, GLM's concept of total assets corresponds more closely to the value of physical capital, though also including the value of government bonds. These considerations help to explain why GLM's average ratio of safe to total assets is above the range predicted by our model. For example, if income from capital constitutes about one-third of GDP, then total assets based on the value of capital would be around one-third of the capitalized value of GDP. In this case, the model's predicted ratio of safe to total assets would be about 30%, close to the numbers calculated by GLM.

There are also reasons why GLM's measure of safe assets would diverge from our theoretical concept, which relates to net lending from group 2 (high risk aversion) to group 1 (low risk aversion). One issue is that the GLM measure does not compute a net figure for liabilities of financial institutions; that is, there is no deduction for safe assets held by these institutions. For example, in 2007-2008, Lehman Brothers issued bonds and commercial paper and also held U.S. government securities and liabilities of other financial firms.[17] On this ground, GLM's measured liabilities of government and financial institutions would overstate the net quantity of safe assets.

[17]Our model could be extended to account for this kind of borrowing and lending within groups. These patterns might arise because of idiosyncratic shocks that affect individual agents within groups, still defined by coefficients of relative risk aversion.

Another consideration is that an array of financial arrangements—including structured finance, stock options, and insurance contracts—can be used to convert risky assets into relatively safe assets. On this ground, the measured liabilities of governments and financial institutions might understate the quantity of safe assets.

GLM also include government liabilities as safe assets but do not include any portion of capitalized future taxes as "safe liabilities," even at the margin. Although it is true that tax liabilities cannot be directly traded, it is also true that these liabilities—and how they vary along with changes in the quantity of government bonds—affect economic analyses of public debt. To the extent that future taxes are factored in by agents, the gross public debt would overstate a meaningful measure of safe assets.

Despite various caveats that affect interpretations of GLM's computed average ratio of safe to total assets, it is possible that their findings on the rough stability of the safe-asset ratio would be valid. That is, the proportionate errors that generate deviations between the measured and "true" ratios of safe to total assets might stay roughly constant over time.

VII. Conclusions

We constructed a model with heterogeneity in risk aversion to study the determination of the equilibrium quantity of safe assets. The model achieves tractability and transparency by assuming two types of agents with Epstein-Zin/Weil utility. The agents differ by coefficients of relative risk aversion but have the same intertemporal elasticity of substitution (IES) and rate of time preference. In the main analysis, each agent has log utility, in the sense of IES=1. We focused on a stationary version of the model in which agents are periodically replaced by new agents (possibly children) who receive a random assignment of coefficient of relative risk aversion. In the baseline setting, Ricardian Equivalence holds in that the quantity of government bonds does not affect the risk-free interest rate or the net quantity of safe assets. The predicted crowding-out coefficient for private bonds with respect to government bonds is -0.5, in line with some existing empirical estimates.

We generated quantitative implications for the quantity of safe assets by calibrating the model with sufficient disaster risk to get the model's predictions into the right ballpark for the average equity premium and risk-free rate. In a benchmark case, the magnitude of safe assets would be around 10% of total assets. We noted that this value could be reconciled with an existing estimate that found the ratio of safe to total assets in the United States to be nearly stable over time at a value between 30% and 35%.

The basic structure of the model with heterogeneity in coefficients of relative risk aversion can be applied to other economic problems. For example, the framework can incorporate issues related to credit-market imperfections, including the necessity for enforcement mechanisms to ensure repayment of private debts. This extension relates to issues concerning collateral, liquidity, and asymmetric information. This type of extension would be important for assessing implications for the magnitude and composition of investment.

References

Bansal, R. and A. Yaron (2004). "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, August, 1481-1509.

Barro, R.J. (2009). "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, March, 243–264.

Barro, R.J. and S. Misra (2013). "Gold Returns," unpublished, Harvard University, October.

Barro, R.J. and J.F. Ursúa (2012). "Rare Macroeconomic Disasters," *Annual Review of Economics*, 83-109.

Blanchard, O.J. (1985). "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, April, 223-247.

Caballero, R.J. and E. Farhi (2014). "The Safety Trap," unpublished, MIT, May.

Chan, Y.L. and L. Kogan (2002). "Catching Up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices," *Journal of Political Economy*, December, 1255-1285.

Dumas, B. (1989). "Two-Person Dynamic Equilibrium in the Capital Market," *Review of Financial Studies*, no. 2, 157-188.

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

Gabaix, X. (2012). "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, May, 645-700.

Garleanu, N. and L. Pedersen (2011). "Asset Prices and Deviations from the Law of One Price," *Review of Financial Studies*, June, 1980-2022.

Gennaioli, N., A. Shleifer, and R. Vishny (2012). "Neglected Risks, Financial Innovation, and Financial Fragility," *Journal of Financial Economics*, June, 452-268.

Giovannini, A. and P. Weil. (1989). "Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model," National Bureau of Economic Research Working Paper 2824, January.

Gorton, G., S. Lewellen, and A. Metrick (2012). "The Safe-Asset Share," *American Economic Review*, May, 101-106.

Gorton, G. and G. Ordoñez (2013). "The Supply and Demand for Safe Assets," unpublished, Yale University, August.

Hart, O. and L. Zingales (2014). "Banks Are where the Liquidity Is," unpublished, Harvard University, May.

Holmstrom, B. and J. Tirole (1998). "Private and Public Supply of Liquidity," *Journal of Political Economy*, February, 1-40.

Krishnamurthy, A. and A. Vissing-Jorgensen (2013). "Short-term Debt and Financial Crises: What we can learn from U.S. Treasury Supply," unpublished, Northwestern University, May.

Longstaff, F.A. and J. Wang (2012). "Asset Pricing and the Credit Market," *Review of Financial Studies*, November, 3169-3215.

Lucas, R.E., Jr. (1978). "Asset Prices in an Exchange Economy," *Econometrica*, November, 1429-1445.

Nakamura, E., J. Steinsson, R.J. Barro, and J.F. Ursúa (2013). "Crises and Recoveries in an Empirical Model of Consumption Disasters," *American Economic Journal: Macroeconomics*, July, 35-74.

Obstfeld, M. (1994). "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, August, 1471-1486.

Wang, J. (1996). "The Term Structure of Interest Rates in a Pure Exchange Economy with Heterogeneous Investors," *Journal of Financial Economics*, May, 75-110.

Weil, P. (1990). "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics*, February, 29-42.

Table 1

Representative-Agent Economy

(Single Coefficient of Relative Risk Aversion)

$\gamma_1=\gamma_2=\gamma$	r^e	r^f
1	0.052	0.046
1.5	0.052	0.042
2	0.052	0.037
2.5	0.052	0.031
3	0.052	0.024
3.5	0.052	0.016
4	0.052	0.005
4.5	0.052	-0.008
5	0.052	-0.023
5.5	0.052	-0.042
6	0.052	-0.064

When the coefficients of relative risk aversion are the same for the two agents, $\gamma_1=\gamma_2=\gamma$, the equilibrium quantities of bonds, B_1 and B_2, are zero and the ownership of equity is evenly distributed, $K_1=K_2=0.5$. The table shows the equilibrium risk-free rate, r^f, for each value of γ. The calculations assume that the growth-rate parameter is g=0.025 per year, the rate of time preference is ρ=0.04 per year, the disaster probability is p=0.04 per year (corresponding in the historical data to contractions of per capita GDP by at least 10%), and the effective disaster size is b=0.32. The expected growth rate is $g^*=g-p\cdot b=$ 0.0122 per year. The reciprocal of the IES is $\theta=1$. The expected rate of return on equity, given $\theta=1$, is $r^e=\rho+g^*=0.052$ per year, which is independent of γ. The price of equity is $P=Y/\rho=25\cdot Y$.

In this representative-agent case, the risk-free rate can be written in closed form, if $\gamma\neq1$, as:

$$r^f=\rho+\theta g+p\left(\frac{\theta-1}{\gamma-1}\right)-p(1-b)^{-\gamma}+p\left(\frac{\gamma-\theta}{\gamma-1}\right)(1-b)^{1-\gamma}\ .$$

If $\theta=1$, as γ approaches 1, r^f approaches $\rho+g-pb/(1-b)$.

Table 2

Steady-State Equity Ownership, Safe Assets, and Wealth Share

Alternative Values of γ_1 and γ_2 that Generate r^f=0.010 when ν=0.02 per Year

(1) γ_1	(2) γ_2	(3) r^e	(4) r^f	(5) K_1	(6) $\lvert B_1 \rvert$/assets [GDP]	(7) W_1/W
3.78	3.78	0.052	0.010	0.50	0.00 [0.00]	0.50
3.50	4.16	0.052	0.010	0.57	0.025 [0.62]	0.54
3.25	4.84	0.052	0.010	0.65	0.054 [1.36]	0.60
3.00	6.66	0.052	0.010	0.77	0.095 [2.38]	0.67
2.75	24.4	0.052	0.010	0.93	0.154 [3.86]	0.78
2.68	647	0.052	0.010	0.996	0.176 [4.41]	0.82

This analysis assumes that the "replacement rate" for agents is ν=0.02 per year. The coefficients of relative risk aversion for the two agents, γ_1 and γ_2, are the values that generate a steady-state risk-free interest rate, r^f, of 0.010, given the other parameter values used in Table 1. Table 2 shows the corresponding steady-state values of agent 1's share of equity ownership, K_1, the ratio of the magnitude of safe assets, B1, to total assets and GDP, and agent 1's share of total assets, W_1/W (which equals agent 1's share of total consumption). Given ρ=0.04 per year, the ratio of safe assets to GDP equals 25 times the ratio to total assets.

Table 3

Steady-State Equity Ownership, Safe Assets, and Wealth Share

Alternative Parameter Values

(1) γ_1	(2) γ_2	(3) r^e	(4) r^f	(5) K_1	(6) $\lvert B_1 \rvert$/assets [GDP]	(7) W_1/W
ν=0.05 per year						
3.50	4.16	0.052	0.0096	0.54	0.025 [0.63]	0.52
3.25	4.84	0.052	0.0083	0.60	0.057 [1.44]	0.54
3.00	6.66	0.052	0.0050	0.70	0.111 [2.77]	0.59
p=0.02 per year						
3.50	4.16	0.052	0.0373	0.54	0.020 [0.49]	0.52
3.25	4.84	0.052	0.0364	0.59	0.046 [1.14]	0.54
3.00	6.66	0.052	0.0342	0.69	0.091 [2.28]	0.60
ρ=0.02 per year						
3.50	4.16	0.032	-0.0099	0.58	0.029 [1.45]	0.55
3.25	4.84	0.032	-0.0094	0.67	0.062 [3.11]	0.61
3.00	6.66	0.032	-0.0087	0.79	0.105 [5.25]	0.68

These results are the same as those shown in Table 2 for the indicated values of γ_1 and γ_2, except for changes in the indicated parameter value. The first three lines use the replacement rate ν=0.05 per year, the next three use the disaster probability p=0.02 per year, and the last three use the rate of time preference ρ=0.02 per year.

Table 4

Changes in Safe Assets when the Government Issues Bonds

	Agent 1	**Agent 2**	**Total**
Case 1: Government bonds up by 100, held 50-50			
Private bond holdings, B	0	0	0
Government bond holdings, B^g	+50	+50	+100
Taxes (present value)	+50	+50	+100
Net safe assets in model	0	0	0
Net safe assets as measured by Gorton, et al. (2012)	+50	+50	+100
Case 2: Government bonds up by 100, all held by agent 2			
Private bond holdings, B	+50*	-50	0
Government bond holdings, B^g	0	+100	+100
Taxes (present value)	+50	+50	+100
Net safe assets in model	0	0	0
Net safe assets as measured by Gorton, et al. (2012)	0	+50	+50

Note: In all cases, the government issues 100 of bonds, B^g, and transfers these bonds 50-50 to agents 1 and 2. The present value of taxes rises by 100, divided 50-50 between agents 1 and 2. In case 1, the added government bonds are held 50-50 by agents 1 and 2. In case 2, all of the added government bonds are held by agent 2.

***Borrowing by agent 1 goes down by 50.

Figure 1

Dynamic Paths Following a Disaster

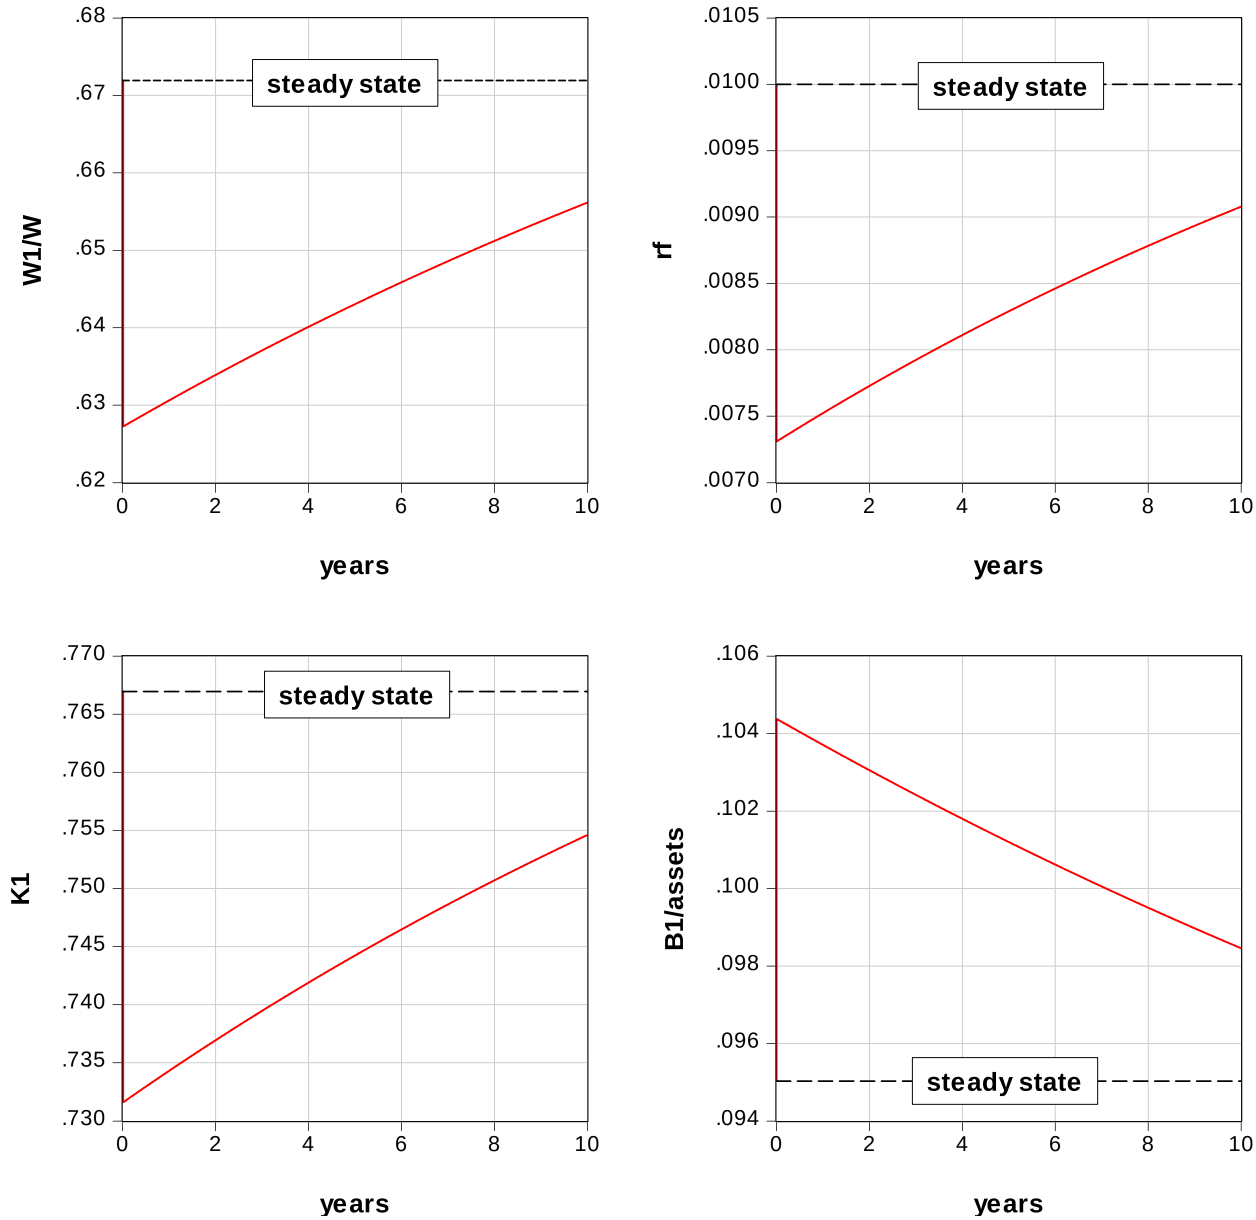

This analysis corresponds to the case where γ_1=3.00 and γ_2=6.66 in Table 2. The simulated paths start from the steady state value of W1/W, 0.672, then assume that a disaster of proportionate size 0.32 materializes in period 1. Subsequently, no further disasters occur. The panels show the dynamic paths after period 1 for agent 1's wealth share, W1/W, the risk-free interest rate, r^f, agent 1's share of total equity, K1, and the ratio of the magnitude of safe assets, B1, to total assets.

Figure 2

Dynamic Paths for 40 Years of Tranquility

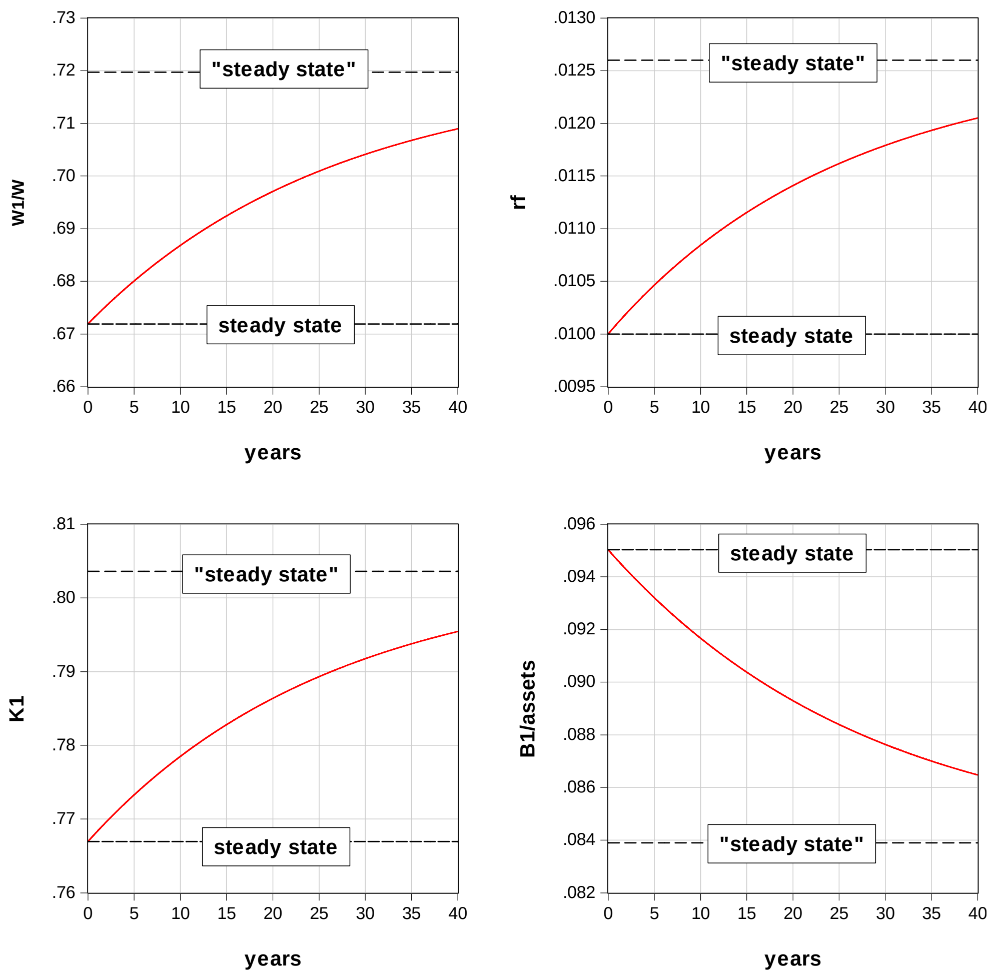

This analysis corresponds to the case where γ_1=3.00 and γ_2=6.66 in Table 2. The simulated paths start from the steady state value of W1/W, 0.672, then assume that no disasters occur over the next 40 years. The panels show the dynamic paths after period 1 for agent 1's wealth share, W1/W, the risk-free interest rate, r^f, agent 1's share of total equity, K1, and the ratio of the magnitude of safe assets, B1, to total assets. The lines marked as "steady states" are values that would be approached asymptotically conditional on disasters never happening.