Appendix to "Wealth Dynamics in a Bond Economy with Heterogeneous Beliefs": Detailed Analysis of Intermediate Economies

Timothy Cogley[*], Thomas J. Sargent[†], Viktor Tsyrennikov[‡]

February 20, 2013

1 Intermediate economies

Preferences, beliefs, and endowments are specified as in the main text. Two assets are traded, a risk-free real bond and a single Arrow security. We study three economies, indexed by $j \in \{1, 2, 3\}$, in which the market for Arrow security j is open and markets for other Arrow securities are closed. In economy j, agent i's flow budget constraint is

$$c^i(g^t) + q_b(g^t)b^i(g^t) + q_j(g^t)s^i_j(g^t) = e^i(g^t) + b^i(g^{t-1}) + s^i_j(g^{t-1}) \cdot 1_j(g_t). \tag{1}$$

On the left side, c^i, b^i, and s^i_j represent agent i's consumption and positions in bonds and Arrow securities, respectively. The bond price is denoted q_b and q_j is the price of Arrow security j. All depend on g^t, the history of aggregate-growth outcomes up to date t. Consumption and security purchases cannot exceed the sum of agent i's current endowment $e^i(g^t)$ plus the financial wealth he brings into the period, $b^i(g^{t-1}) + s^i_j(g^{t-1}) \cdot 1_j(g_t)$. The indicator function $1_j(g_t)$ equals 1 when $g_t = j$ and is zero otherwise.

So that an equilibrium exists, we also assume that consumers are subject to borrowing limits. As in our bond economy, we assume they can take a negative position in risk-free bonds up to a limit of twice their annual income:

$$b^i(g^t) \geqslant -By^i(g^t), \quad B = 2. \tag{2a}$$

[*]New York University, tim.cogley@nyu.edu
[†]New York University and Hoover Institution, thomas.sargent@nyu.edu
[‡]Cornell University, vt68@cornell.edu

Here consumers can also borrow by selling the Arrow security up to a limit of one annual income:

$$s_j^i(g^t) \geqslant -Sy^i(g^t), \quad S = 1. \tag{2b}$$

The combined borrowing capacity is thus state contingent: $By^i(g^t)$ when $g_t \neq j$ and $(B+S)y^i(g^t)$ when $g_t = j$.[1]

When the market for security j is open, the wealth share of agent i is:

$$\omega^i(g^t) = \frac{e^i(g^t) + b^i(g^{t-1}) + s^i(g^{t-1}) \cdot 1_j(g_t)}{e(g^t)}. \tag{3}$$

With two types of agents, $\omega^2(g^t) = 1 - \omega^1(g^t)$ and the distribution of wealth shares is conveniently summarized by the wealth share of the less-informed agent $\omega^1(g^t)$. In what follows, we refer to the wealth share of the less-informed agent as the economy's wealth distribution.

We restrict attention to wealth-recursive competitive equilibria. In a wealth-recursive equilibrium, individual decisions and the price system are functions of the wealth distribution ω, the current aggregate state g_t, and the parameters of agent 1's beliefs $n(g^t), m(g^t)$.

A *wealth-recursive competitive equilibrium* is a price system $(q_b(\omega, g, n, m), q_j(\omega, g, n, m))$ and a list of policy functions $(\rho_c^i(\omega, g, n, m), \rho_b^i(\omega, g, n, m), \rho_s^i(\omega, g, n, m))_{i=1}^2$ such that:
a) decision rules $(\rho_c^i, \rho_b^i, \rho_s^i)$ maximize agent i's subjective welfare given the price system;
b) goods and financial markets clear;
c) the evolution of the wealth distribution is consistent with individual decisions:

$$\omega(g^{t+1}) = \frac{e^1(g^{t+1}) + \rho_b^1(\omega(g^t), g_t, n_t, m_t) + \rho_s^1(\omega(g^t), g_t, n_t, m_t) \cdot 1_j(g_{t+1})}{e(g^t)}. \tag{4}$$

d) the evolution of agent 1's beliefs is consistent with the Bayes' Law:

$$n(g^{t+1}) = n(g^t) + 1(g_{t+1} = g_l),$$
$$m(g^{t+1}) = m(g^t) + 1(g_{t+1} = g_m).$$

Section (3) describes how the equilibrium is approximated and asseses the quality of approximation.

[1] The choice of S is restricted by the choice of B. The maximum amount of debt that an agent can end up with in state j is $B + S$. Since bond prices are sufficiently close to 1 rolling B units of debt is possible by selling B new bonds and selling a small portion of one's income. Repaying additional S units of debt may be impossible when j is a recession state and q_j is relatively low. The bound S must be relatively small to insure that both agents can always repay their debts. We experimented numerically with different values and found that $S \leqslant 0.80B$ is sufficient to insure full repayment for all choices of j. However, we set $S = 0.50B$ to keep solution accuracy at a satisfactory level. All the results are qualitatively equal for the two choices of S.

2 Simulations

The calibration is the same as in the text. We simulate 200,000 sample paths for g_t, each of length 100 years. This ensemble is held constant across economies. The driving force in this model is differences in estimates of p_d. Figure 1 plots the ensemble average of estimates of p_d by the less-informed agent 1, with the true value $p_d = 0.1$ shown as a horizontal dashed line. The estimate starts at $\hat{p}_d^1 = 0.50$ and converges gradually to 0.1. Convergence is slow, however; even after 400 periods the learning agent overestimates the probability of a deep recession by 0.057.

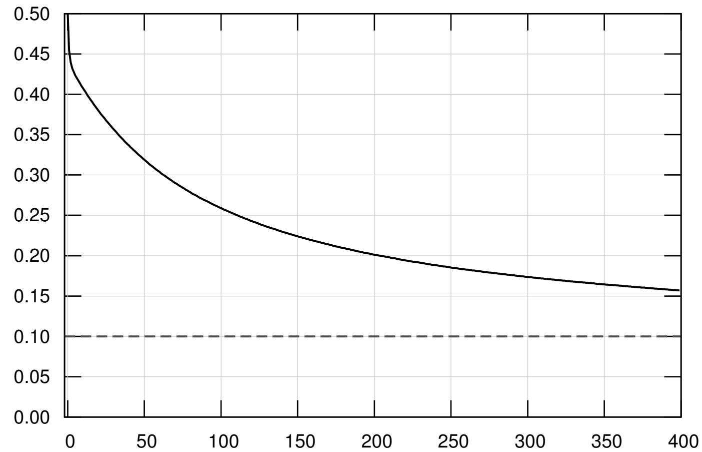

Figure 1: Dynamics of $\hat{p}_d^1$

Next we compute equilibrium prices and allocations for five economies that are identical in all respects except for their financial-market structures. Figure 2, reproduced from the text, summarizes our main result, showing how the ensemble average share of wealth for less-informed type-1 consumers depends on market structure. The complete-markets and the bond economies represent two extremes. In the complete markets economy, survival forces dominate and the less-informed agent's wealth approaches a lower bound determined by the borrowing limit. The opposite happens in the bond economy. Driven by a precautionary savings motive, the less-informed agent accumulates the maximum possible financial wealth. Three intermediate economies allow trading of one Arrow security along with a risk-free bond. The rate at which the less-informed agent accumulates wealth decreases as we move from the bond economy to one in which markets for expansion- or mild-recession-state securities are open (securities 1 and 2, respectively), and when a deep-contraction security is traded (security 3), survival forces return to the fore and type-1 consumers lose wealth.

Figures 3 and 4 record more details about the intermediate economies. Figure 3 portrays quantiles of the cross-sample path distribution of financial wealth for type-1 consumers along with their consumption share and position in the available Arrow

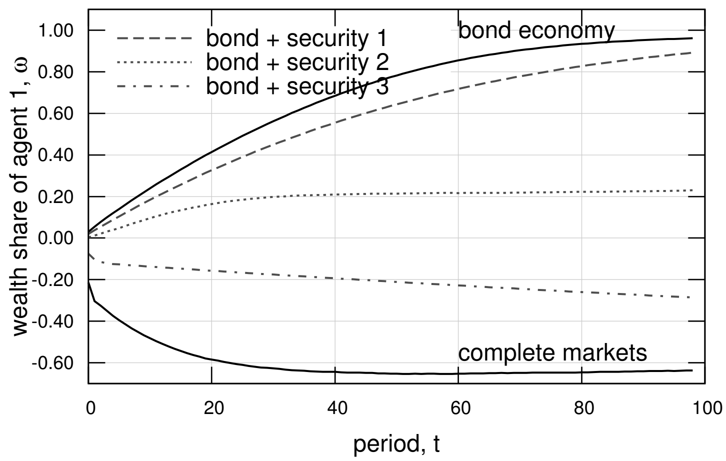

Figure 2: Average wealth share of the less-informed consumer under different financial market structures

security. Column j represents an economy in which a risk-free bond and Arrow security j are traded. Similarly, figure 4 portrays quantiles of the cross-sample path distribution for asset prices. The following sections explain the economic forces that generate these outcomes.

2.1 Economy 1: A risk-free bond plus the expansion Arrow security

We begin with an economy in which markets for the risk-free bond and expansion-state Arrow security are open. Markets for Arrow securities paying off in mild recessions and deep contractions are closed. Results for this economy are shown in the first column of figures 3 and 4.

Broadly speaking, the results resemble those of the bond economy. As shown in figure 3, pessimistic consumers accumulate wealth rapidly. After 30 periods, the learning agent's median financial wealth equals half of the economy's income, and it asymptotes near 100 percent. Their consumption share starts below 50% of aggregate income, but it grows quickly and asymptotes around 55 percent of total income. These asymptotes are reached when better-informed, type-2 consumers arrive at their borrowing limits. Finally, type-1 consumers initially sell the expansion-state Arrow security, but their position converges to zero as time goes on.

Figure 4 compares asset prices in our diverse-beliefs economies (shown as solid lines) with those in comparable economies populated entirely by well-informed type-2 consumers (dashed lines). As shown in the subplots in the first column, the risk-free bond price is higher than in an economy with homogenous beliefs, and the Arrow-

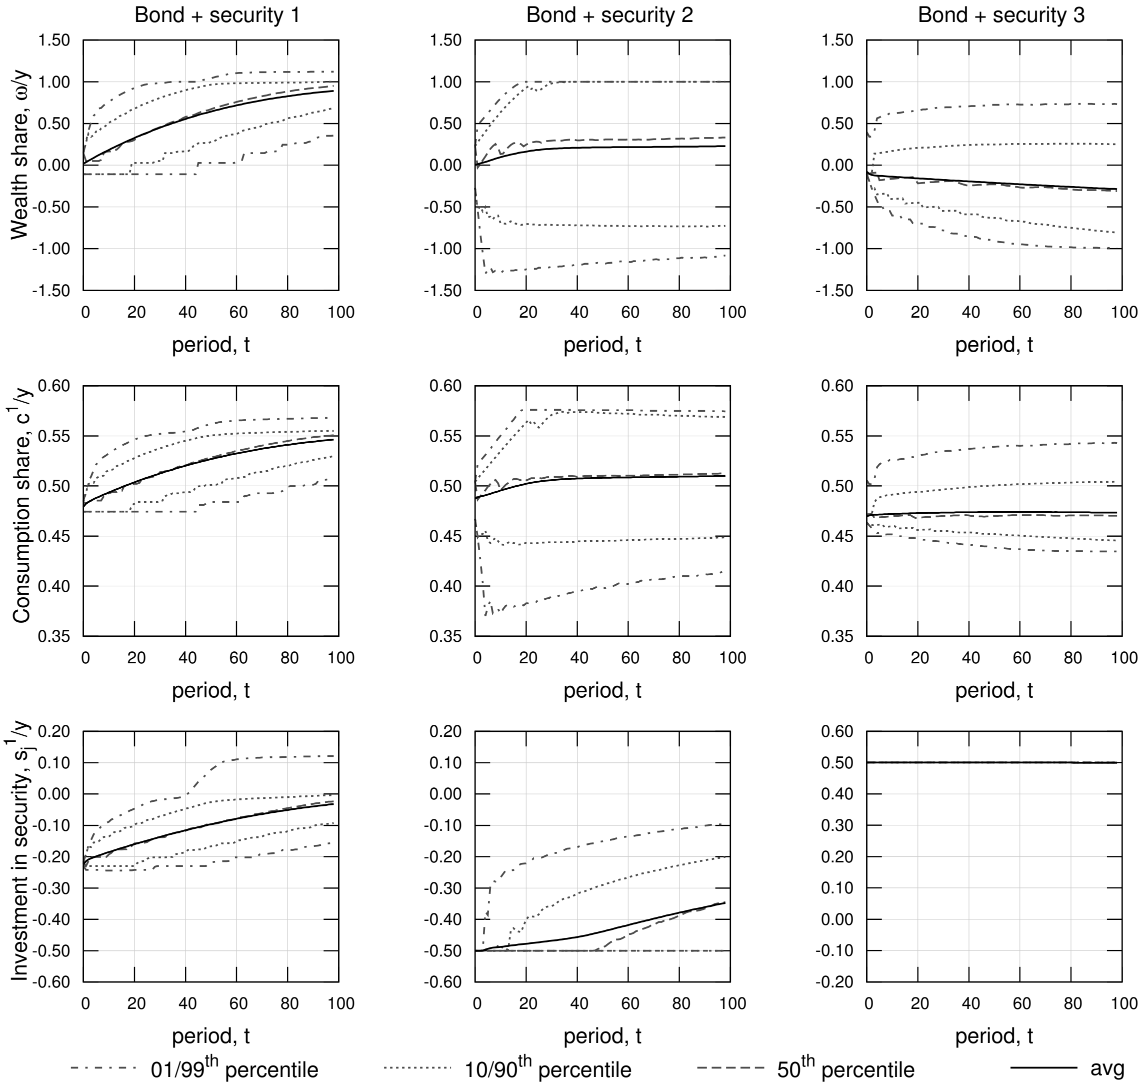

Figure 3: Dynamics of wealth share $\hat{\omega}^1$, consumption share $\hat{c}^1$ and income share invested in security $\hat{s}_j^1$

security price is lower. Both prices converge to their full-information valuations, with the Arrow-security price converging rapidly and the bond price converging slowly.

Intuition for outcomes in this economy can be developed by thinking about channels for precautionary saving. Since the learning agent is pessimistic, he wants to buy assets that pay off in deep contractions. In this economy, a positive payoff in deep contractions can be achieved only by holding the risk-free bond. Type-1 consumers therefore purchase the bond. This drives its price above its full-information value and induces type-2 consumers to sell. To afford a larger position in the risk-free bond, type-1 consumers sell the expansion-state Arrow security. Because there is no

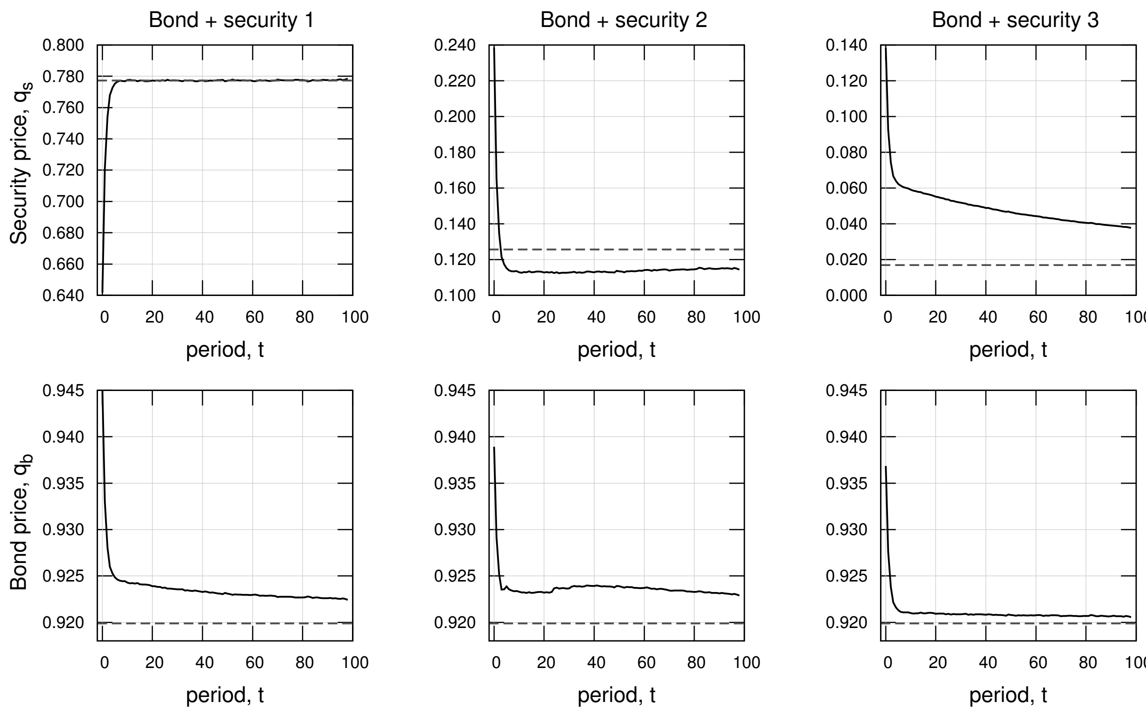

Figure 4: Bond and security price dynamics

disagreement about the expansion state, the less-informed agent can sell the security only at discount, thereby driving its price below its full-information value and inducing type-2 consumers to buy. As pessimism evaporates, the less-informed agent becomes less willing to sell at a discount and trade in the security converges to zero.

Table 1 records the sign of financial payoffs to agent 1 in each state, with this economy shown in column 2. The risk-free bond has a positive payoff in all three states, while the expansion-state Arrow security has a negative payoff in state 1 (because type-1 consumers are sellers) and zero payoff in the other states. The last row records the sign of the net payoff on their portfolio, positive in mild recessions and deep contractions and ambiguous in expansions.

	economy 1			economy 2			economy 3		
state	1	2	3	1	2	3	1	2	3
bond payoff	+	+	+	+	+	+	−	−	−
security payoff	−	0	0	0	−	0	0	0	+
portfolio payoff	?	+	+	+	?	+	−	−	?

Table 1: Optimal portfolio choice of the learning agent 1

As in the bond-only economy, precautionary savings are channeled into the risk-free bond. Less-well-informed investors therefore accumulate financial wealth, and better-informed agents accumulate debts. The direction in which wealth is transferred is the same as in the bond economy and opposite to that in the complete-markets economy. Relative to the bond economy, the quantitative effect of opening the market for the expansion-state Arrow security is to retard the rate at which type-1 consumers accumulate wealth. Since type-1 consumers sell this security, they must make a payment to type-2 consumers in every expansion, an outcome that occurs 86% of the time.

2.2 Economy 2: A risk-free bond plus the mild-recession Arrow security

Next we examine an economy in which an Arrow security paying off in mild recessions is traded along with a risk-free bond. Markets for Arrow securities paying off in expansions and deep contractions are closed. Results for this economy are depicted in the subplots in the second columns of figure 3 and 4. The patterns shown there are qualitatively similar to those for the first economy, as pessimistic consumers again accumulate wealth and enjoy increasing consumptions shares. Both grow more slowly than in economy 1, however, and they level off at lower values. For instance, median wealth and consumption asymptote at about 30 percent and slightly above 50 percent of aggregate income, respectively. There is also substantially more dispersion in the cross-sample-path distribution.

Intuition can again be developed by thinking about channels for precautionary savings. Pessimistic consumers want to purchase assets that pay off in deep contractions. Since no Arrow security contingent on deep contractions is traded, the risk-free bond is the only asset with a positive payoff in that state (see the third column of table 1.) The learning agent sells mild-recession Arrow securities partly to afford a larger position in the risk-free bond and partly because he believes the price is too low and therefore represents a good investment opportunity. The payoff on the optimal portfolio is positive in states 1 and 3 and ambiguous in state 2.[2]

Consumer 1's demand for the risk-free bond again drives up its price and induces better-informed agents to sell (see column 2, figure 4). The Arrow security price initially rises above its full-information value, then drops sharply and converges from below. The initial rise is due to the borrowing constraint, which binds at the beginning of the transition. Unconstrained knowledgeable agents would like to buy more than learning agents are allowed to sell, and competition for the limited quantity drives up

[2]Although the payoff in state 2 is ambiguous, it is likely to be negative. When agents buy a risk-free bond, they affect their financial wealth in every state tomorrow. But when Arrow security 2 is available, agents can choose their financial positions for state 2 independently of other states. Since the type 1 learning agent underestimates the probability of a transition into state 2, his consumption and wealth must decrease whenever state 2 occurs. Hence, wealth of the learning agent decreases whenever a mild recession occurs, $g_t = 2$, and increases otherwise, $g_t \in \{1, 3\}$.

the price. As belief differences grow smaller, the constraint slackens and this effect goes away. Bond-price dynamics are more intricate. Until period 40, the learning agent is constrained in the security market along more than 50% of sample paths. During this period, the demand for the bond is determined largely by the price of the security. Since the latter drops sharply and then recovers slowly, so does the bond price. After period 40, bond-price dynamics are driven by the declining pessimism of the less informed agent.

One difference between economies 1 and 2 is that borrowing constraints bind more often in economy 2. Type-1 consumers take a much larger negative position in the Arrow security than in economy 1, and their short positions hover close to the borrowing limit of 1 annual income (half of the aggregate income) for much of a typical simulation. Because type-1 consumers must occasionally make large payments on their short positions, this slows the rate at which they accumulate wealth. Furthermore, because type-1 consumers underestimate the probability of mild recessions, their bets against mild recessions go awry more often than they expect. Their large short positions and more-often-than-expected investment losses explain why their financial wealth and consumption increase more slowly and asymptote at lower levels than in economy 1 and also why the cross-sample path distribution has more dispersion.

Thus, less-well-informed investors still accumulate financial wealth because they channel precautionary savings into the risk-free bond. The rate at which their wealth increases is slower because they suffer greater losses from their speculation on the Arrow security. Despite that, the direction in which wealth is transferred is the same as in the bond economy and economy 1 and opposite to that in the complete-markets economy. Finally, results for versions of this economy in which consumers are less tightly constrained – and speculative losses are greater – are qualitatively similar. Precautionary accumulation of risk-free bonds remains the dominant force.

2.3 Economy 3: A risk-free bond plus the deep-contraction Arrow security

In our third economy, consumers can trade a risk-free bond and an Arrow security paying off in deep contractions. Trade in Arrow securities paying off in expansions and mild recessions is prohibited. Outcomes for this economy are shown in the third column of figures 3 and 4.

The wealth of type-1 consumers declines at a slow but significant pace. Their median debt equals 20 percent of aggregate income after 30 years and 30 percent after 100 years (40 and 60 percent of individual income, respectively). Relative to a complete-market economy with the same borrowing limits, the chief differences are that wealth declines more slowly and asymptotes at a higher level. This happens because the closure of other Arrow-security markets limits the extent to which better-informed investors can profit by trading with their less-well-informed counterparts. Similarly, the consumption share of type-1 consumers is lower than their income share, at 42 and 50 percent of aggregate income, respectively. Thus, type-1 consumers devote roughly 16 percent of individual income to debt service.

Opening a market for a deep-contraction security therefore changes wealth dynamics dramatically. Since the learning agent can now insure against deep-contraction risk by buying an Arrow security that pays off in that state, precautionary savings need no longer be channeled into the risk-free bond. On the contrary, type-1 consumers now sell the risk-free bond in order to afford larger purchases of the Arrow security. Type-1 consumers therefore sell safe assets in order to leverage their purchases of the risky asset. As shown, in the fourth column of table 1, their optimal portfolio has a positive payoff in the disaster state, $g_t = 3$, and a negative payoff otherwise, $g_t \in \{1, 2\}$. Because they over-estimate the probability of deep contractions and underestimate the probability of mild recessions, they suffer financial losses more often and reap gains less often than expected, thus losing wealth on average. In this respect, economy 3 resembles a complete-markets economy. The rate at which wealth is transferred is slower than in a complete-markets economy simply because fewer markets are open and fewer speculative opportunities exist.

The learning agent's demand for the deep-contraction security drives its price above its full-information valuation (see figure 4, column 3). Indeed, the security price is more than twice its full-information valuation during the whole sample. Better-informed consumers sell deep-contraction securities because they think they are overpriced. The opportunity is so attractive that they bump against their borrowing limits the entire time.[3] As in the complete-markets economy, better-informed consumers grow rich on average by selling 'overpriced' disaster insurance.

2.4 The role of the disaster risk

In section 4, we derive the following relation between the agents' intertemporal marginal rates of substitution (IMRS),

$$\begin{aligned}\Delta\text{IMRS} \equiv E_t^2 &\left[(g_c^2(g^{t+1}))^{-\gamma} - (g_c^1(g^{t+1}))^{-\gamma}\right] \\ &= (\pi^1(g_l|g^t) - \pi^2(g_l|g^t))[(g_c^1(g^t, g_l))^{-\gamma} - (g_c^1(g^t, g_m))^{-\gamma}]. \qquad (5)\end{aligned}$$

If the term on the right side is positive, then the marginal utility of the less-informed agent 1 is expected to decline, or equivalently, his consumption to grow. Because agent 1 (on average) overestimates the probability of deep recessions, $\pi^1(g_l|g^t) - \pi^2(g_l|g^t) > 0$, we get

$$\text{sign}(\Delta\text{IMRS}) = \text{sign}(g_c^1(g^t, g_m) - g_c^1(g^t, g_l)).$$

[3]If borrowing limits were loosened, type-1 consumers would lose wealth even more rapidly.

That is, the direction in which wealth is transferred depends on the relative size of agent 1's consumption growth in a mild and a deep recession states. We use this characterization to assess the role of disaster risk.

We seek to understand the forces that make the less-informed agent accumulate wealth in economy 2 (bond plus mild recession security) and decumulate wealth in economy 3 (bond plus deep recession security). The two suspects are the size of deep recessions, $g_m - g_l$, and the level of disagreement, $|\pi^1(g_l|g^t) - \pi^2(g_l|g^t)|/\pi^2(g_l|g^t)$. Consider economy 2. As $g_m - g_l$ decreases, precautionary motives grow weaker and individual consumption growth rates get closer to the aggregate growth rate. That is, $g_c^1(g^t, g_m)$ increases, $g_c^1(g^t, g_l)$ decreases, and ΔIMRS in (5) is less likely to be positive. So the *size* of deep recessions is an unlikely suspect to cause a wealth-dynamics reversal. But as the *relative* disagreement[4] about transition into a mild recession state increases, we should expect $g_c^1(g^t, g_m) \ll g_m$, while $g_c^1(g^t, g_l)$ cannot deviate much from g_l because the deep recession security market is closed. A sufficiently large disagreement can lead to $g_c^1(g^t, g_l) > g_c^1(g^t, g_m)$ and, hence, a reversal of wealth dynamics.[5] In turn, large disagreements are most likely to arise for rare events.

Note the roles of pessimism and the "rareness" of the disaster state. Pessimism motivates the learning agent to take a positive position in the deep-contraction security. "Rareness" implies that the relative disagreement across agents is large; hence, speculative motives are strong. Together these forces lead the less-informed agent to borrow in order to leverage purchases of the over-priced disaster security. Figure 5 demonstrates the effect of disagreement. It plots average paths of the less-informed agent 1's wealth share in the economy 2 for different levels of the true probability of the disaster state p_d. The prior is kept the same. As we increase p_d, we reduce room for significant disagreement and tame the survival forces. When the probability mass shifts towards the deep recession state, mild recessions become the pivotal state. At $p_d = 0.90$, the roles of the recessions states are fully exchanged relative to the benchmark parametrization.[6]

The above experiment keeps the size of a deep recession fixed but varies its likelihood. We now fix the likelihood of a deep recession and vary the size. A caveat is that when we reduce the size of a deep recession (i.e., increase g_l) mild and deep recession states become "payoff-identical". Because two securities are traded, one can perfectly hedge income fluctuations. However, speculation possibilities are not fully unlocked because betting on one's beliefs is not possible without a third asset. Keeping this

[4]Relative disagreement about transition into state k is defined as $|\hat{\pi}^1(g_k|g^t) - \pi(g_k|g^t)|/\pi(g_k|g^t)$.

[5]Suppose that the less-informed agent 1 assigns zero probability to a mild recession state. In this case, we must have $g_c^1(g^t, g_m) = 0 < g_l \approx g_c^1(g^t, g_l)$ (for details see Tsyrennikov, 2011).

[6]Note that at $p_d = 0.5$ the learning agent is losing wealth. Given the prior, the less informed agent is equally often pessimistic and optimistic. When he overestimates the probability of a disaster he will buy bonds. Because he also sells the mild recession security his wealth cannot grow fast. When he underestimates the disaster probability he will sell bonds and buy mild recession securities. Because mild recessions are less frequent than imagined by the agent, he will lose wealth rapidly.

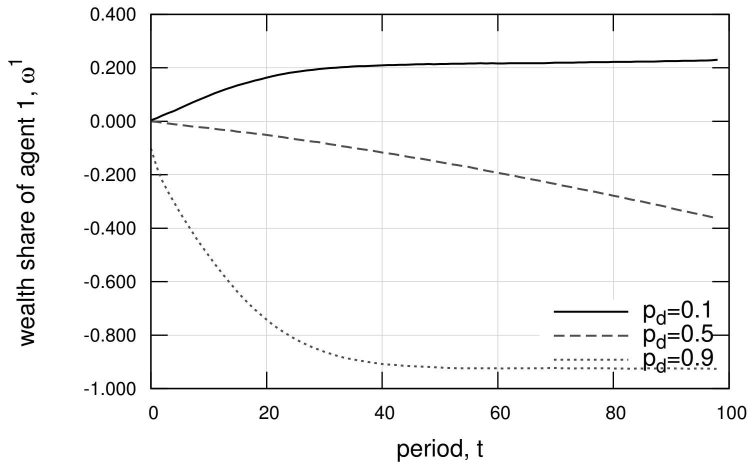

Figure 5: Wealth dynamics with different probabilities of the disaster state

in mind, figure 6 plots the evolution of wealth in economy 2 for different levels of g_l. The solid line corresponds to $g_l = 0.90$ as in the benchmark parametrization. As we increase g_l from 0.90 to 0.95 the less-informed agent accumulates wealth at a slower pace. For sufficiently shallow 'deep' recessions (e.g., $g_l = 0.97$) and in the limit when $g_l = 0.99 = g_m$ the less-informed agent 1 is loosing wealth. However, without a third asset, speculation is restricted and, hence, the speed is slower than predicted in Blume and Easley (2006). Finally, the effect is smaller when compared to that of changing p_d.

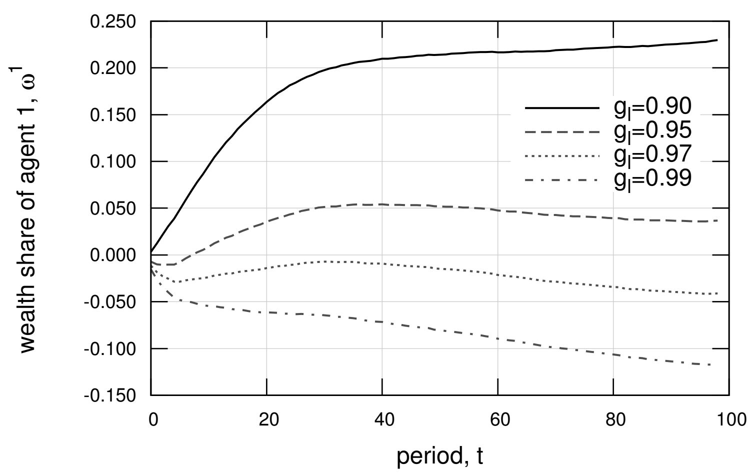

Figure 6: Wealth dynamics with different probabilities of the disaster state

3 Approximation methods

The solution consists of consumption $\rho^i_c(\hat{b}, n, m, s)$ and bond investment $\rho^i_b(\hat{b}, n, m, s)$ policy functions, the Lagrange multipliers associated with borrowing limits $\rho^i_\mu(\hat{b}, n, m, s)$ and the bond price function $q_b(\hat{b}, n, m, s)$. We solve for the policy functions iteratively using the system of equilibrium conditions. The stopping criterion is that the sup distance between consecutive policy function updates is less than $e_\rho = 10^{-6}$.

We verify the computed solution on a grid 5 times denser than the one used to compute policy functions. The verification procedure consists of computing the following error functions:

$$e_1(\omega, n, m, s) =$$
$$1 - \frac{1}{\rho^1_c(\omega, n, m, s)} \left[\frac{q_b(\omega, n, m, s)}{E^1[(\rho^1_c(\omega', n', m', s')g(s'))^{-\gamma}] + \rho^1_{\mu,b}(\omega, n, m, s)} \right]^{1/\gamma},$$
$$e_2(\omega, n, m, s) =$$
$$1 - \frac{1}{\rho^2_c(\omega, n, m, s)} \left[\frac{q_b(\omega, n, m, s)}{E^2[(\rho^2_c(\omega', n', m', s')g(s'))^{-\gamma}] + \rho^2_{\mu,b}(\omega, n, m, s)} \right]^{1/\gamma},$$
$$e_3(\omega, n, m, s) =$$
$$1 - \frac{1}{\rho^1_c(\omega, n, m, s)} \left[\frac{q_s(\omega, n, m, s)}{\beta\pi^1(j|s, n, m)\rho^1_c(\omega', n', m', j)g(j)^{-\gamma} + \rho^1_{\mu,s}(\omega, n, m, s)} \right]^{1/\gamma},$$
$$e_4(\omega, n, m, s) =$$
$$1 - \frac{1}{\rho^2_c(\omega, n, m, s)} \left[\frac{q_s(\omega, n, m, s)}{\beta\pi^2(j|s)\rho^2_c(\omega', n', m', j)g(j)^{-\gamma} + \rho^2_{\mu,s}(\omega, n, m, s)} \right]^{1/\gamma},$$
$$e_5(\omega, n, m, s) =$$
$$1 - \frac{\omega - q_b(\omega, n, m, s)\rho^1_b(\omega, n, m, s) - q_s(\omega, n, m, s)\rho^1_s(\omega, n, m, s)}{\rho^1_c(\omega, n, m, s)}.$$

These error functions are designed to answer the following question: "what fraction should be added/subtracted from an agent's consumption so that an equilibrium condition holds exactly?" The first two equations are the consumption Euler equations for agent 1's and 2's bond position, respectively. The next two equations are the consumption Euler equations for agent 1's and 2's security position respectively. The fifth equation is the budget constraint of agent 1. Because feasibility constraint is imposed, the budget constraint of agent 2 holds exactly.

We started with 100 grid points for the wealth share ω and increased this number until a sufficient level of accuracy was achieved. With 1,000 grid points the errors are smaller than 0.38% of average consumption.[7] For comparison, the statistical

[7]This amounts to 38$ for every 10,000$ of consumption. This error decreases linearly with the

discrepancy in the U.S. NIA between 1929 and 2010 averaged 0.54% of total income.[8]

Figure 7 plots *maximal* errors (over all possible vectors of counters and growth states) for each ω and 1000 grid points. These errors are largest when the markets for recession state securities, both mild and severe, are open. The maximal errors for economy 1, 2 and 3 are respectively 0.138%, 0.378% and 0.305%.

4 IMRS criterion

Define $r(g^{t+1}) = \pi^1(g_{t+1}|g^t)/\pi^2(g_{t+1}|g^t)$, and notice that $E_t^2[r(g^{t+1})] = 1$. Then the bond Euler equation implies

$$\begin{aligned} E_t^2\left[(g_c^2(g^{t+1}))^{-\gamma}\right] &= E_t^1\left[(g_c^1(g^{t+1}))^{-\gamma}\right], \\ &= E_t^2\left[r(g^{t+1})(g_c^1(g^{t+1}))^{-\gamma}\right], \\ &= E_t^2\left[(g_c^1(g^{t+1}))^{-\gamma}\right] + E_t^2\left[(r(g^{t+1})-1)(g_c^1(g^{t+1}))^{-\gamma}\right]. \end{aligned}$$

Because agents agree on transitions into the expansion state, $r(g^t, 1) = 1$,[9]

$$\begin{aligned} &E_t^2\left[(r(g^{t+1})-1)(g_c^1(g^{t+1}))^{-\gamma}\right] \\ &\quad = \sum_{g_j}(r(g^t,g_j)-1)(g_c^1(g^t,g_j))^{-\gamma} = \sum_{g_j\in\{g_m,g_l\}}(r(g^t,g_j)-1)(g_c^1(g^t,g_j))^{-\gamma} \\ &\quad = (\pi^1(g_m|g^t)-\pi^2(g_m|g^t))(g_c^1(g^t,g_m))^{-\gamma} + (\pi^1(g_l|g^t)-\pi^2(g_l|g^t))(g_c^1(g^t,g_l))^{-\gamma} \\ &\quad = (\pi^1(g_l|g^t)-\pi^2(g_l|g^t))[(g_c^1(g^t,g_l))^{-\gamma}-(g_c^1(g^t,g_m))^{-\gamma}]. \end{aligned}$$

Combining all the above results we get:

$$E_t^2\left[(g_c^2(g^{t+1}))^{-\gamma}-(g_c^1(g^{t+1}))^{-\gamma}\right] = (\pi^1(g_l|g^t)-\pi^2(g_l|g^t))[(g_c^1(g^t,g_l))^{-\gamma}-(g_c^1(g^t,g_m))^{-\gamma}].$$

number of grid points. However, (slightly more than) 1000 grid points is the operating system's permissible maximum memory.

[8] Model errors are smaller when normalized by the total income.

[9] Agreement about state 1 implies the following relations:

$$\begin{aligned} E_t^2\left[(r(g^{t+1})-1)(g_{ct+1}^1)^{-\gamma}\right] &= \sum_{j\in\{1,2,3\}}(r(g^t,g_j)-1)(g_c^1(g^t,g_j))^{-\gamma} \\ &= \sum_{j\in\{2,3\}}(r(g^t,g_j)-1)(g_c^1(g^t,g_j))^{-\gamma} \\ &= (\pi^1(2|g^t)-\pi^2(2|g^t))(g_c^1(g^t,g_m))^{-\gamma} + (\pi^1(3|g^t)-\pi^2(3|g^t))(g_c^1(g^t,g_l))^{-\gamma} \\ &= (\pi^1(3|g^t)-\pi^2(3|g^t))[(g_c^1(g^t,g_l))^{-\gamma}-(g_c^1(g^t,g_m))^{-\gamma}]. \end{aligned}$$

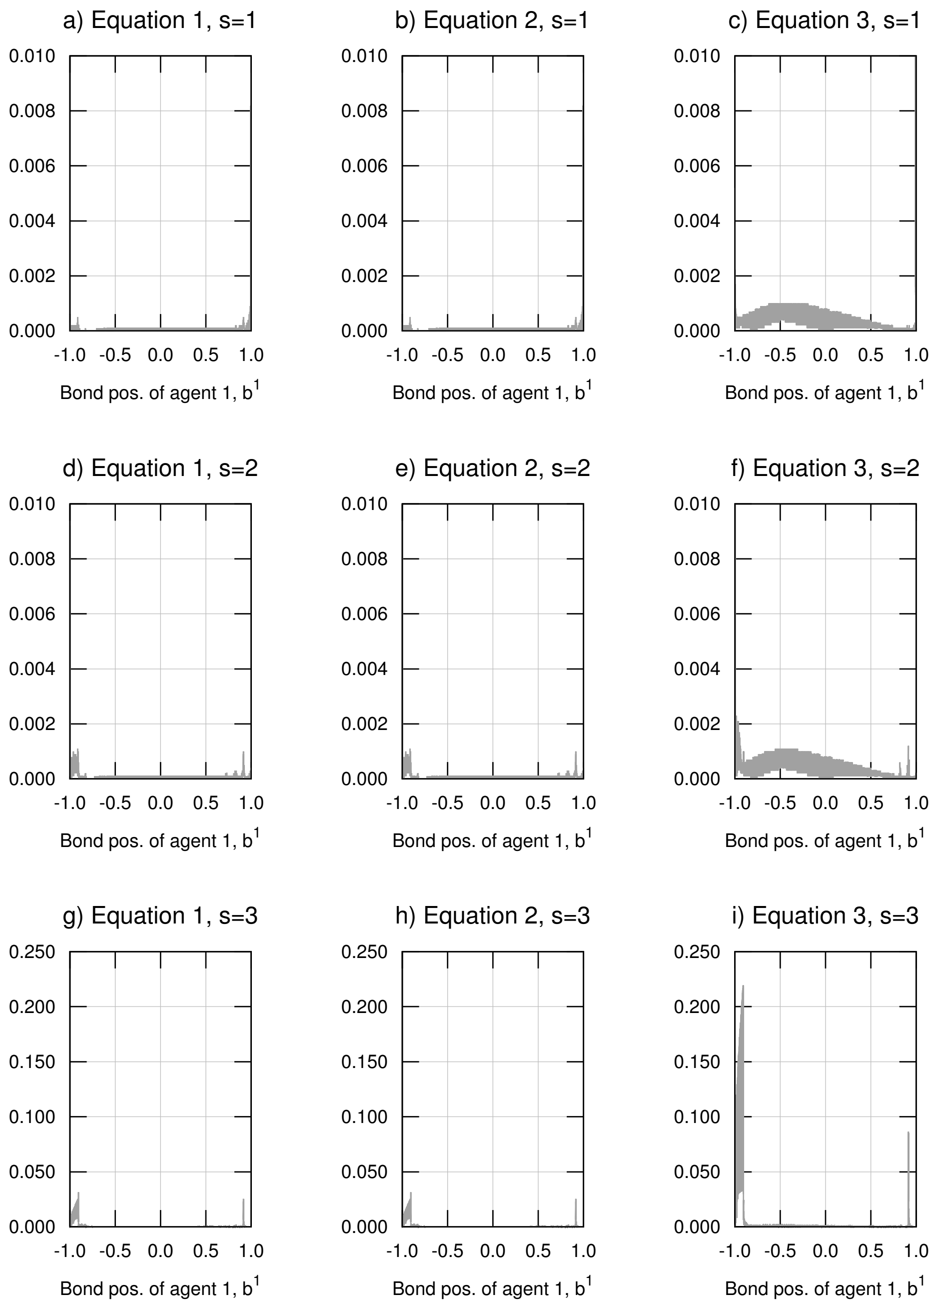

Figure 7: Solution errors

5 References

Blume, L. and D. Easley, 2006, If You'Re So Smart, Why Aren't You Rich? Belief Selection In Complete And Incomplete Markets, *Econometrica* 74(4) 929-966.

Cogley, T., T.J. Sargent, and V. Tsyrennikov, 2013, Wealth Dynamics in a Bond Economy with Heterogeneous Agents, unpublished manuscript, Cornell University and NYU.

Tsyrennikov, V., 2011, Mechanics of Wealth Dynamics with Heterogeneous Beliefs, unpublished manuscript, Cornell University.