

Wealth Dynamics in a Bond Economy with Heterogeneous Beliefs*

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

November 12, 2011

Abstract

We study an economy in which two types of agents have diverse beliefs about the law of motion for an exogenous endowment. One type knows the true law of motion, and the other learns about it via Bayes's theorem. Financial markets are incomplete, the only traded asset being a risk-free bond. Borrowing limits are imposed to ensure the existence of an equilibrium. We analyze how financial-market structure affects the distribution of financial wealth and survival of the two agents. When markets are complete, the learning agent loses wealth during the learning transition and eventually exits the economy (Blume and Easley 2006). In contrast, in a bond-only economy, the learning agent accumulates wealth, and both agents survive asymptotically, with the knowledgeable agent being driven to his debt limit. The absence of markets for certain Arrow securities is central to reversing the direction in which wealth is transferred.

JEL CLASSIFICATION: D52, D53, D83, D84

KEY WORDS: Wealth dynamics, survival, incomplete markets, diverse beliefs, learning, speculation

1 Introduction

We study how market incompleteness affects the distribution of wealth when beliefs are heterogeneous, extending research by Blume and Easley (2006) and Cogley and Sargent (2009), among others. Blume and Easley describe conditions under which

*For comments and suggestions, we thank seminar participants at Georgetown and the September 2011 Cornell/Penn State Macroeconomics Conference.

[†]New York University. Email: tim.cogley@nyu.edu

[‡]New York University and Hoover Institution. Email: thomas.sargent@nyu.edu

[§]Cornell University. Email: vt68@cornell.edu

Friedman's (1953) survival hypothesis holds. Roughly speaking, when markets are complete and so allocations are Pareto optimal, agents that have to estimate more parameters learn slower and are driven out of the market. However, they also provide examples showing that agents with incorrect beliefs can survive when markets are incomplete.¹

Similarly, Cogley and Sargent (2009) study aspects of the transition in a competitive complete-markets economy, including the rate at which less-well informed consumers are driven out, how wealth is transferred between well- and less-well-informed consumers, and how their financial positions affect subjective prices of risk. Their model is populated by two types of consumers, one that knows the true transition probabilities for an exogenous Markov endowment and another that learns about them. The less-well-informed consumers are pessimistic, initially over-estimating the probability of a contraction state. Because of their pessimism, Arrow securities paying off in that state are overpriced relative to rational-expectations valuations while those paying off in an expansion state are underpriced. The better-informed consumers regard these price gaps as attractive trading opportunities, buying low-priced expansion-state securities and selling high-priced contraction-state securities. They grow rich on average because the expansion state occurs more often and the contraction state less often than less-well-informed agents expect. But when a contraction occurs, better-informed consumers not only suffer a decline in their endowment but also are obliged to pay off on their contraction-state liabilities. Because their financial positions increase exposure to catastrophic risk, their subjective prices of risk are high during the learning transition. And since less-well-informed consumers buy 'contraction insurance,' their subjective prices of risk are low.

In the Cogley-Sargent model, the ability to trade a complete set of Arrow securities is central to the mechanism by which better-informed consumers grow rich at the expense of less-well-informed consumers. Complete markets give the agents many opportunities to make trades motivated solely by the different subjective probabilities they put on future outcomes. In this paper, we study what happens when some of those Arrow-security markets are closed. For a model very much like that of Cogley and Sargent, we find that the direction in which wealth is transferred is reversed when agents can trade only a risk-free bond. Less-well-informed consumers accumulate financial assets, and better-informed consumers are driven to debt limits. Moreover, both agents survive asymptotically, with better-informed agents rolling over their debt forever.

Precautionary motives play a central role. Because less-well-informed consumers

¹See also Kogan, et al. (2006), Becker and Espino (2011), Coury and Sciubba (2010), Cao (2011), and Tsyrennikov (2011).

are pessimistic, their precautionary motives are stronger than in the rational expectations version of the model. When markets are complete, they can guard against deep contractions by purchasing an Arrow security that pays off in that state. They lose wealth on average because those states occur less often than they expect. In a bond economy, they guard against deep contractions by buying risk-free bonds, thus accumulating wealth. The real-interest rate adjusts so that better-informed agents are content to sell risk-free bonds. This reverses the direction in which wealth is transferred during the learning transition.

2 The Model

Time is discrete and is indexed by $t \in \{0, 1, 2, \dots\}$. The set of possible states in each period is finite and is denoted by \mathcal{G} . In particular, \mathcal{G} is the set of all possible realizations of the aggregate income growth rate. The set of all sequences or histories of states is denoted by Σ . The partial history of the state through date t of is denoted by g^t . The set of all partial histories of length t is Σ^t . We will also make use of the “true” probability measure on Σ denoted by π^0 .

2.1 Preferences

There are two types of consumers, indexed by $i = 1, 2$. Agent i ranks different consumption plans $c = \{c(g^t) : \forall t, \forall g^t \in \Sigma^t\}_{t=0}^{\infty}$ using a time-separable, isoelastic welfare function:

$$U^i(c) = E^i \sum_{t=0}^{\infty} \beta^t u(c(g^t)), \quad \beta \in (0, 1), \quad (1)$$

where

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

We assume that the preference parameters β and γ are the same across types and that the consumers differ in how they form expectations. The expectation operator E^i signifies that each type forms predictions by averaging with respect to his own subjective probability distribution over future outcomes. The consumers choose consumption and savings plans to maximize expected utility subject to flow budget constraints and debt limits that are specified below.

2.2 The aggregate endowment and distribution of income

The two types receive constant shares of a non-storable aggregate endowment $y(g^t)$,

$$y^i(g^t) = \phi^i y(g^t), \quad i = 1, 2. \quad (3)$$

Growth in the aggregate endowment can take on one of three values $\{g_h, g_m, g_l\} \equiv \mathcal{G}$. The high-growth state represents an expansion, the medium-growth state is a mild contraction, and the low-growth state is a deep contraction or disaster. These outcomes depend on the realization of two independent random variables, s and d . The random variable s is a Markov-switching process with a transition matrix,

$$\Pi_s = \begin{bmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{bmatrix}. \quad (4)$$

The random variable d is an iid Bernoulli variate with success probability p_d . The mapping from (s, d) realizations to growth outcomes is as follows:

$$g = \begin{cases} g_h & \text{when } s = 1 \text{ and } d = 1 \text{ or } d = 2, \\ g_m & \text{when } s = 2 \text{ and } d = 1, \\ g_d & \text{when } s = 2 \text{ and } d = 2. \end{cases} \quad (5)$$

The high-growth state occurs when $s = 1$ independently of the outcome for d , a mild contraction occurs when $s = 2$ and $d = 1$, and a deep contraction occurs when $s = d = 2$. The resulting transition matrix for growth states is:

$$\Pi_g = \begin{bmatrix} p_{11} & (1 - p_{11})(1 - p_d) & (1 - p_{11})p_d \\ 1 - p_{22} & p_{22}(1 - p_d) & p_{22}p_d \\ 1 - p_{22} & p_{22}(1 - p_d) & p_{22}p_d \end{bmatrix}. \quad (6)$$

2.3 Information and beliefs

We chose the specification for the aggregate endowment in order to make the learning problem as simple as possible. In particular, because learning statistics become part of the state vector, we want to reduce the learning problem to a single unknown parameter. Toward that end, we assume that Π_s is known to both agents and that p_d is known only to agent 2. It follows that agent 2 knows the true transition matrix Π_g , while agent 1 does not. Agent 1 learns about p_d by applying Bayes's theorem.

Both agents observe realizations of the growth states g_t but not realizations of the underlying random variables (s_t, d_t) . Because s and d are independent and g_h can occur when d equals 1 or 2, entry into the high-growth state conveys no information

about p_d . Information about d is revealed only when the economy moves into a contraction and agents see whether it is mild or deep.²

We assume that less-well-informed, type-1 consumers have identical beta priors on p_d ,

$$f(p_d) = \mathcal{B}(n_0, m_0), \quad (7)$$

where $n_0 - 1$ is the prior number of disasters ($d = 2$) and $m_0 - 1$ is the prior number of non-disasters ($d = 1$). It follows that the prior mean for p_d is $\hat{p}_d = n_0/(n_0 + m_0)$.

Because d is an iid Bernoulli random variable, the likelihood function is proportional to

$$f(g^t|p_d) \propto p_d^{n_t} \cdot (1 - p_d)^{m_t}, \quad (8)$$

where g^t represents the observed history of growth states and n_t and m_t are the number of deep and mild contractions, respectively, counted through date t . These counters evolve according to

$$(n_{t+1}, m_{t+1}) = \begin{cases} (n_t, m_t) & \text{when } g_{t+1} = g_h, \\ (n_t, m_t + 1) & \text{when } g_{t+1} = g_m, \\ (n_t + 1, m_t) & \text{when } g_{t+1} = g_l. \end{cases} \quad (9)$$

Since the prior is beta and the likelihood function is binomial, the posterior is also a beta density,

$$f(p_d|g^t) = \mathcal{B}(n_0 + n_t, m_0 + m_t). \quad (10)$$

The posterior predictive density over a potential future trajectory g_t^f emanating from g^t is

$$f(g_t^f|g^t) = \int f(g_t^f|p_d, g^t) f(p_d|g^t) dp_d. \quad (11)$$

Type 1 consumers form expectations by averaging potential future sequences with weights assigned by $f(g_t^f|g^t)$. Their one-step ahead transition matrix is

$$\Pi_{gt}^1 = \begin{bmatrix} p_{11} & (1 - p_{11})(1 - \hat{p}_{dt}) & (1 - p_{11})\hat{p}_{dt} \\ 1 - p_{22} & p_{22}(1 - \hat{p}_{dt}) & p_{22}\hat{p}_{dt} \\ 1 - p_{22} & p_{22}(1 - \hat{p}_{dt}) & p_{22}\hat{p}_{dt} \end{bmatrix}, \quad (12)$$

where $\hat{p}_{dt} = (n_0 + n_t)/(n_0 + n_t + m_0 + m_t)$ is the posterior mean.³ The better-informed type-2 consumers form expectations using the true transition probabilities $f(g_t^f|p_d, g^t)$. Because our model satisfies the conditions of a Bayesian consistency theorem, differences in beliefs vanish eventually. However, learning will be slow because opportunities to learn arise only in contractions, which occur in 1 year out of 7 for our calibration. Hence differences in beliefs remain active for quite some time.

²The second and third columns of Π_g depend on p_d , but the first column does not.

³For recursive versions of the model, the one-step transition matrix is all we need.

Following Cogley and Sargent (2009), we study Walrasian equilibria in which traders take prices as given and do not infer information from prices. We put individuals in a setting in which the only information revealed by prices is subjective probabilities over future endowment paths. We short-circuit the problem of learning from prices by endowing agents with common information sets along with knowledge of each other’s prior. With this simplification, agents learn nothing from prices because there is nothing to learn.

3 A Complete-Markets Benchmark

As a point of departure, we study how the model behaves when markets are complete. We therefore temporarily assume that a full set of Arrow securities is traded, one for each of the aggregate growth states.⁴ Our preliminary objective is to establish a link with previous literature. We first establish that our model behaves in much the same way as those of Blume and Easley (2006) and Cogley and Sargent (2009). Then we address how matters differ when markets are incomplete.

Our model is a special case of one of Blume and Easley’s complete-markets frameworks. Because the beliefs of type-2 consumers are correct and more parsimonious than those of type-1 individuals, their results imply that

$$\limsup_{t \rightarrow \infty} \hat{c}_t^1 = 0, \tag{13}$$

where variables with hats represent shares of aggregate income, $\hat{x}(g^t) = x(g^t)/y(g^t)$. Type-1 consumers eventually come arbitrarily close to exhausting their borrowing capacity, after which time their consumption is arbitrarily close to zero.⁵ In that sense, type-2 agents are the only ones who survive in a competitive equilibrium.

Cogley and Sargent (2009) study aspects of the transition. To make contact with their model, we turn to a simulation.

3.1 Asset markets, budget constraints, and debt limits

An Arrow security j bought in period t pays one unit of consumption in period $t + 1$ if the growth state $g_{t+1} = j$ is realized and zero units otherwise. The security that pays in state $g_{t+1} = j$ trades at price $Q(g_{t+1} = j|g^t)$.

⁴Payoffs cannot be contingent on the realization of the disaster-state random variable d because it is unobserved.

⁵Since the one-period utility function satisfies the Inada conditions, the natural borrowing limit never actually binds in equilibrium.

After history g^t , the flow budget constraint for agent i is

$$y^i(g^t) + a^i(g_t = k|g^{t-1}) \geq c^i(g^t) + \sum_{j \in \mathcal{G}} Q(g_{t+1} = j|g^t) a^i(g_{t+1} = j|g^t). \quad (14)$$

The variable $y^i(g^t)$ represents income, $c^i(g^t)$ is consumption, $a^i(g_{t+1} = j|g^t)$ denotes purchases at history g^t of the Arrow security paying off when $g_{t+1} = j$. When there is no risk of confusion, we will abbreviate the notation for Arrow prices and quantities as $Q_j(g^t)$ and $a_j(g^t)$, respectively.

Individuals can borrow by taking negative positions in Arrow securities subject to a debt limit. For the complete-markets version of our model, we adopt the natural borrowing limit, which constrains total borrowing to be no greater than the maximum that can be repaid with certainty,

$$\tilde{B}^i(g^t) = \sum_{\tau=t}^{\infty} \sum_{g^\tau \in \Sigma^\tau} Q(g^\tau|g^t) y^i(g^\tau), \quad (15)$$

where $Q(g^\tau|g^t) \equiv Q(g^\tau|g^{\tau-1}) \cdots Q(g^{t+1}|g^t)$. This borrowing limit equals the history- g^t value of the continuation of agent i 's endowment stream. This represents the amount of debt service consumer i could sustain if all future income were devoted to that purpose. Notice that the natural borrowing limit is in general history and individual dependent.

3.2 Complete-markets equilibrium

Because the endowment is perishable, the aggregate resource constraint is

$$y(g^t) \equiv y^1(g^t) + y^2(g^t) = c^1(g^t) + c^2(g^t), \quad \forall t, g^t \in \Sigma^t. \quad (16)$$

Since Arrow securities are in zero net supply, financial markets clear when

$$a_j^1(g^t) + a_j^2(g^t) = 0, \quad \forall t, g^t \in \Sigma^t, j \in \mathcal{G}. \quad (17)$$

Initial endowments of Arrow securities are given and satisfy this market-clearing condition.

We seek a recursive competitive equilibrium that satisfies the following conditions: (i) individuals formulate consumption, savings, and portfolio plans by maximizing subjective expected utility; (ii) beliefs are updated via Bayes theorem (iii) flow-budget constraints and borrowing limits are satisfied; (iv) consumption plans respect the aggregate-resource constraint; (v) and security prices adjust so that financial markets clear.

Appendix A describes a recursive formulation for the model and outlines an algorithm for computing its equilibrium. The algorithm exploits the fact that a competitive equilibrium solves a Pareto problem. We use a modified Negishi algorithm that replaces Pareto weights with an initial distribution of consumption. The solution to the Negishi problem tells us how to compute recursively consumption plans for each agent. Asset prices are then calculated from their subjective Euler equations. Then the consumption plans and price system are used to back out asset-trading plans. The last step delivers an initial wealth distribution that supports the computed prices and allocations as a competitive equilibrium. This relation can be inverted to compute a competitive equilibrium for any initial wealth distribution.

3.3 A numerical example

We use simulation methods to study the transition. The time period is one year, the discount factor $\beta = 1.04^{-1}$, and the coefficient of relative risk aversion is $\gamma = 2$. The endowment process is calibrated so that the high-growth state g_h represents an expansion, the medium-growth state g_m a mild recession, and the low-growth state g_l a deep contraction,

$$g_h = 1.03, \quad g_m = 0.99, \quad g_l = 0.90. \quad (18)$$

The true transition probabilities Π_g are calibrated so that the economy spends most of its time in the expansion state and visits the deep-contraction state rarely,

$$p_{11} = 0.917, \quad p_{22} = 0.50, \quad p_d = 0.10. \quad (19)$$

These numbers imply that an expansion has a median duration of 8 years, a mild recession has a median duration of 1 year, and that 1 in 10 contractions are deep. The implied one-step transition matrix is

$$\Pi_g = \begin{bmatrix} 0.917 & 0.0747 & 0.0083 \\ 0.50 & 0.45 & 0.05 \\ 0.50 & 0.45 & 0.05 \end{bmatrix}, \quad (20)$$

and the ergodic probabilities are

$$pr(g_h) = 0.8576, \quad pr(g_m) = 0.1281, \quad pr(g_l) = 0.0142. \quad (21)$$

The unconditional probability of a deep contraction is therefore in the same ballpark as the estimates of Barro and co-authors.⁶ Finally, we assume that each agent receives 50 percent of the aggregate endowment in each period: $\phi^i = 0.5, i = 1, 2$.⁷

⁶E.g., see Barro (2006), Barro and Ursua (2008), and Barro, Nakamura, Steinsson, and Ursua (2011).

⁷How ϕ^i is calibrated matters quantitatively but not qualitatively. Our main insights remain valid for economies with $\phi^1 = 0.2$ and 0.8 .

Following Cogley and Sargent (2009), we assume that less-well-informed type-1 consumers are initially pessimistic, over-estimating the probability of a deep contraction. In particular, we assume their prior is

$$p_d \sim \mathcal{B}(5, 5), \quad (22)$$

implying a prior mean $\hat{p}_{d0} = 0.50$. The implied prior transition and long-run probabilities are

$$\Pi_{g0}^1 = \begin{bmatrix} 0.917 & 0.0415 & 0.0415 \\ 0.50 & 0.25 & 0.25 \\ 0.50 & 0.25 & 0.25 \end{bmatrix}, \quad (23)$$

and

$$pr^1(g_h) = 0.8576, \quad pr^1(g_m) = 0.0712, \quad pr^1(g_l) = 0.0712, \quad (24)$$

respectively. Type 1 consumers therefore initially overestimate the likelihood of deep contractions and underestimate that of mild recessions. Appendix B demonstrates that type-1 consumers are only moderately pessimistic, in the sense that their priors would be statistically difficult to distinguish from those of type-2 consumers in samples 50 years long. Their pessimism dissipates as events unfold, but this happens slowly because opportunities to learn about p_d arise only in recessions or contractions.

We simulate 10,000 sample paths for g_t , each of length 200 years. Along each path, we compute consumption and financial wealth for each agent as well as equilibrium prices of Arrow securities. Figures 1-3 summarize the results.

We begin by examining prices for Arrow securities and comparing them with outcomes in the full-information rational-expectations version of the model. Because borrowing constraints never bind in equilibrium, Arrow-security prices can be expressed as

$$Q(g_{t+1}|g_t) = Q^{FI}(g_{t+1}|g_t) \left[\hat{c}^1(g_t) \left(\frac{\pi^1(g_{t+1}|g_t)}{\pi^2(g_{t+1}|g_t)} \right)^{1/\gamma} + (1 - \hat{c}^1(g_t)) \right]^\gamma, \quad (25)$$

where

$$Q^{FI}(g_{t+1}|g_t) = \beta g_{t+1}^{-\gamma} \pi^2(g_{t+1}|g_t), \quad (26)$$

represents security prices in the full-information, type-2 economy. When consumers agree on transition probabilities ($\pi^1(g_{t+1}|g_t) = \pi^2(g_{t+1}|g_t)$), prices in the heterogeneous-beliefs economy coincide with those in the full-information model. Conditional on the consumption share, $Q(g_{t+1}|g_t)$ is increasing in the probability ratio $\pi^1(\cdot)/\pi^2(\cdot)$, implying that Arrow securities are overpriced relative to full-information valuations when $\pi^1(\cdot) > \pi^2(\cdot)$ and underpriced when $\pi^1(\cdot) < \pi^2(\cdot)$.

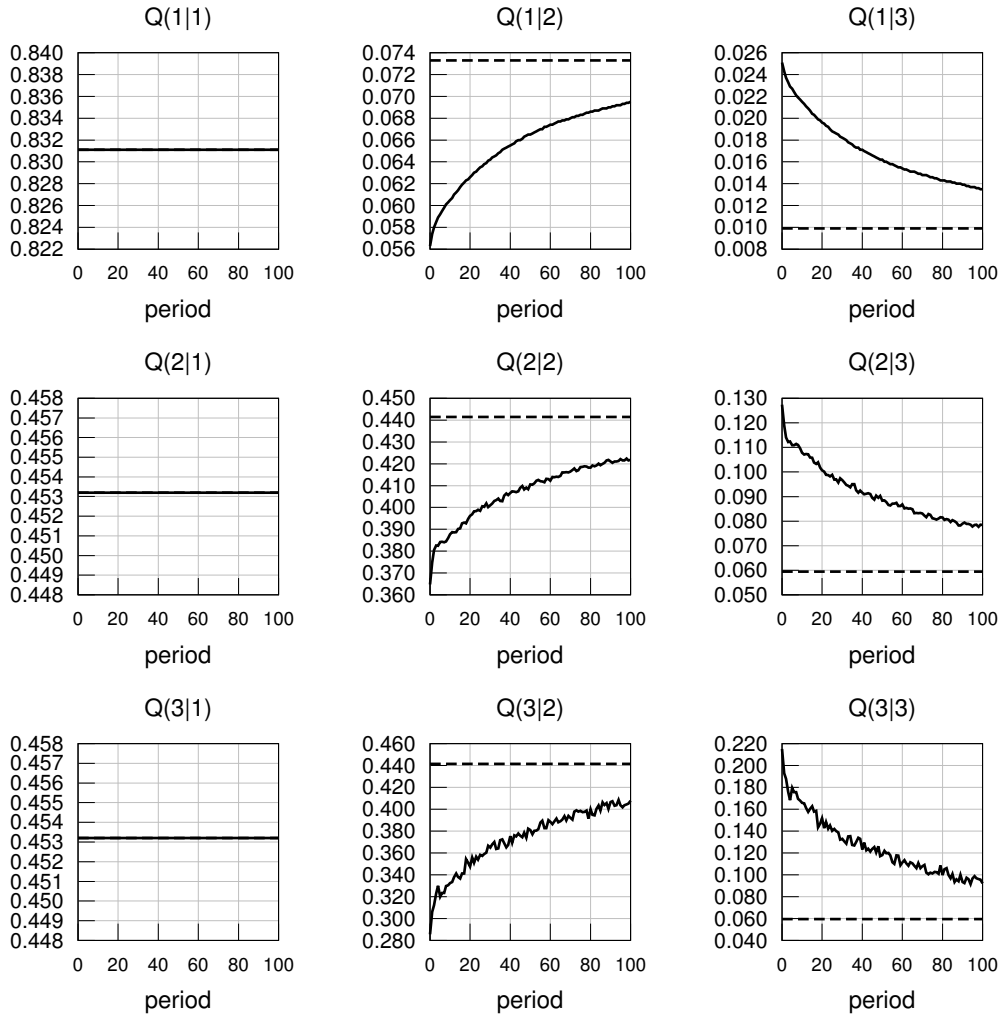


Figure 1: Average prices of Arrow securities

The solid lines in figure 1 depict average prices across sample paths in the diverse-beliefs economy, while dashed lines represent what prices would be in a homogenous-beliefs economy populated entirely by well-informed type 2 consumers. Because all consumers agree about transition probabilities into the high-growth state, the two sets of prices coincide for the security paying off in that state (see the first column of figure 1). Disagreements emerge about transition probabilities for entering mild and deep contractions, and gaps between the two sets of prices appear for securities paying off in those states. Relative to full-information rational-expectations values, the price of the security paying off in mild recessions is too low while that of the security paying off in deep contractions is too high. This follows from the fact that type-1 consumers overestimate the probability of deep contractions and underestimate that

of mild contractions.

Prices eventually converge to full-information rational-expectations values, but convergence is slow, with substantial gaps remaining after 100 years. Type-1 consumers learn slowly because they update their posteriors for p_d only in recessions or contractions, which occur roughly 14 percent of the time. Many years must pass before a large sample of recessions/contractions is available.

Well-informed type-2 consumers regard these price gaps as attractive trading opportunities, for they believe the deep-contraction security will pay off less often and the mild-contraction security more often than type-1 consumers. They sell ‘overpriced’ deep-contraction securities and buy ‘underpriced’ mild-contraction securities, with type-1 consumers taking the opposite position. Type-1 consumers therefore pay out on their financial liabilities when a mild recession occurs, while type-2 consumers pay out in the event of a deep contraction.

Type-2 consumers profit on average from these trades because recession-state securities do in fact pay off more often and contraction-state securities less often than type-1 consumers expect. The consequences for consumption and wealth are shown in figure 2. The left column portrays cross-sample-path averages of consumption and financial wealth for the two agents, both normalized as a proportion of aggregate income.⁸ Consistent with the results of Blume and Easley (2006), the less-well-informed agent 1 loses wealth quickly. Her average debt is roughly 3 times aggregate income after 20 years and 4 times income after 50 years. As her liabilities accumulate, more and more of her income is devoted to debt service, and her consumption declines. Eventually she is driven to the vicinity of her borrowing limit and effectively exits the economy, although this usually does not occur within the first 100 years.

More formally, because borrowing constraints are slack in equilibrium, an analytical expression can be found for growth in the consumption share. After dividing agent 1’s Euler equation by that of agent 2 and re-arranging terms, we find

$$\frac{\hat{c}^1(g_{t+1})}{\hat{c}^1(g_t)} = \left[\hat{c}^1(g_t) + (1 - \hat{c}^1(g_t)) \left(\frac{\pi^2(g_{t+1}|g_t)}{\pi^1(g_{t+1}|g_t)} \right)^{1/\gamma} \right]^{-1}. \quad (27)$$

When consumers agree on transition probabilities ($\pi^1(\cdot) = \pi^2(\cdot)$), their consumption shares remain unchanged. Consumption shares therefore remain constant when the economy transitions into the expansion state. Because agent 1 is usually pessimistic about contractions, her consumption share declines when the economy transitions into a mild recession ($\pi^1(g_m|g_t) < \pi^2(g_m|g_t)$), and it increases when the economy transitions into a deep contraction ($\pi^1(g_l|g_t) > \pi^2(g_l|g_t)$). Because disasters are infrequent, agent 1’s consumption share trends down on average.

⁸Multiply by 2 to measure them as proportions of individual income.

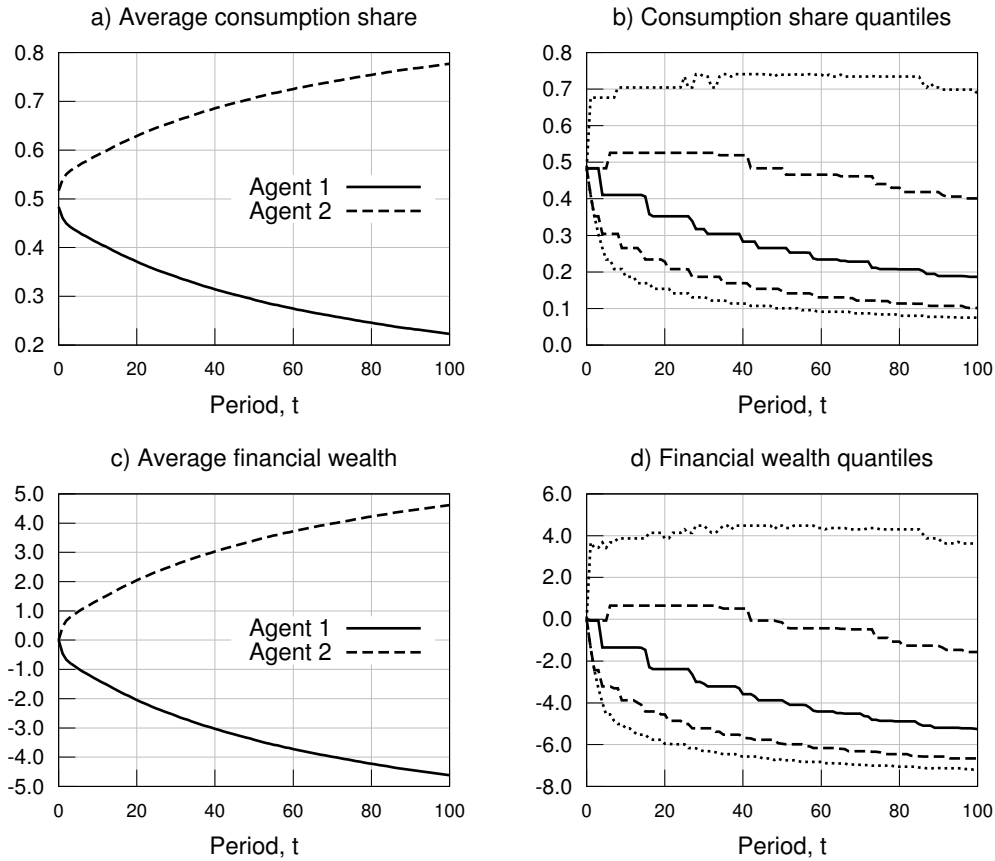


Figure 2: Wealth and consumption dynamics under complete markets

The right panel of figure 2 portray quantiles of the consumption and wealth distribution for agent 2. Percentiles are computed date-by-date from the respective cross-sample-path distributions. Consistent with the previous results, agent 2's median wealth drifts upward over time, as does his median share of consumption. Notice, however, that the marginal distributions are skewed to the left, and there is a small probability that agent 2's wealth and consumption shares decline over time. Agent 2 bets aggressively against deep contractions. With small probability, some sample paths emerge with many deep contractions, and agent 2's bets turn out badly in that case. Those paths account for the lower tails shown in the figure.

More detail is provided in 3, which illustrates four sample paths from the ensemble. The four paths correspond to the 1st, 10th, 90th, and 99th percentiles, respectively, of cross-sectional distribution for agent 2's consumption share in year 100. The top row illustrates paths on which agent 2's bets turn out well. No deep contractions and many mild ones occur on the sample path shown in the upper-left panel. Agent 2 never

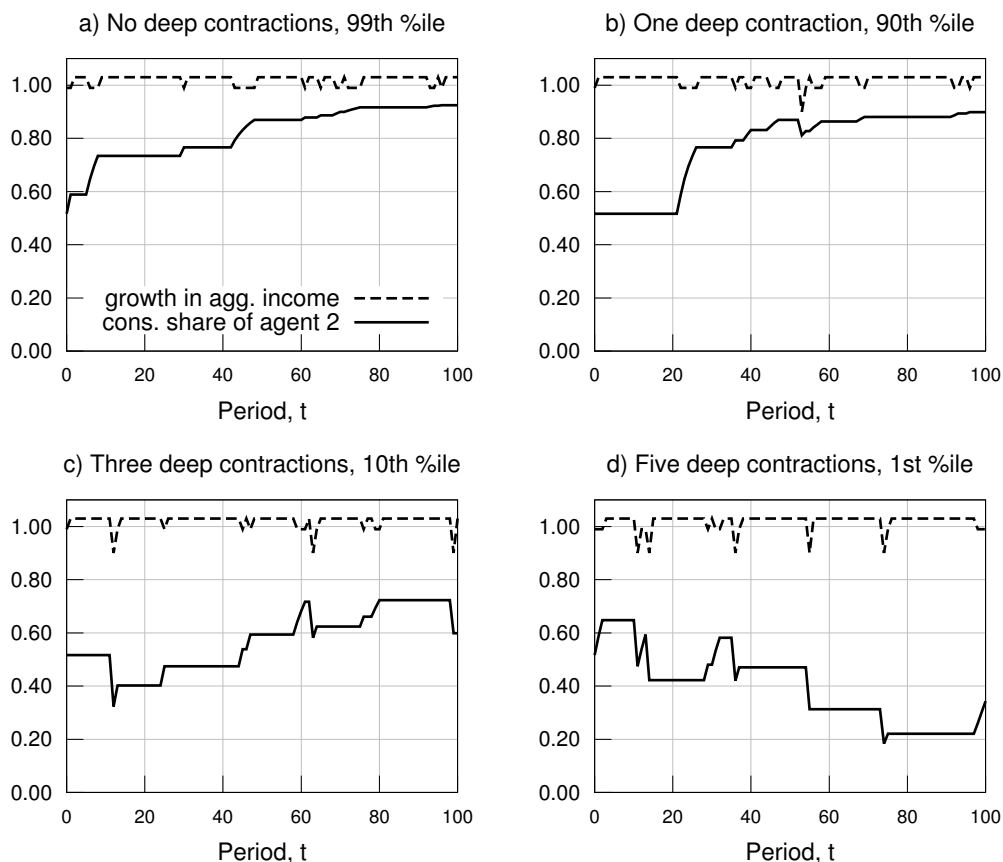


Figure 3: Selected sample paths under complete markets

has to pay out on contraction-state liabilities and frequently collects on recession-state assets. His consumption share ticks upward each time a mild recession occurs, reflecting the gain in his financial wealth. Similarly, the upper-right panel depicts a sample path with a single deep contraction and many mild ones. He again collects often on recession-state assets and has to pay out only once on contraction-state liabilities. Wealth and consumption tick up each time a mild recession occurs and decline slightly in the deep contraction.

The bottom row illustrates sample paths on which agent 2's bets turn out less well. Fewer mild recession and two deep contractions occur on the sample path depicted in the lower-left panel. Gains in mild recessions still offset losses in deep contractions, but just barely. Agent 2's wealth and consumption are only slightly higher at the end of the sample than at the beginning. The lower-right panel portrays a rare sample path with many deep contractions and few mild ones. Agent 2's bets turn out badly in this case, and wealth and consumption share both decline over time.

These results establish that the complete-markets version of the model behaves in essentially the same way as that of Cogley and Sargent (2009). We turn next to our main question, which is how the evolution of the wealth and consumption distributions differ when markets are incomplete.

4 A Bond Economy

There are many ways that markets could be incomplete. In this section, we shut markets for Arrow securities and assume instead that only a risk-free bond is traded. Hence the flow-budget constraint for individual i after history g^t becomes

$$y^i(g^t) + b^i(g^{t-1}) \geq c^i(g^t) + q_b(g^t)b^i(g^t), \quad (28)$$

where $q_b(g^t)$ and $b^i(g^t)$ represent the price of the bond and the quantity held by consumer i , respectively. Individuals can borrow by taking a negative position in the bond subject to a borrowing limit,

$$B^i(g^t) \equiv B^i \cdot y(g^t). \quad (29)$$

We initially assume that a consumer's debt cannot exceed twice his annual income, a constraint that is quite a bit tighter than the natural debt limit considered in the previous section. Later we examine what happens when the debt limit is relaxed.

In all other respects, the model is identical to the one in sections 2 and 3. The only differences are the number of assets traded and the limit on borrowing.

4.1 A recursive formulation

We solve for a wealth-recursive Markov equilibrium in which the distribution of financial wealth is an endogenous state variable.⁹ Because the aggregate endowment process is non-stationary, we need to scale the variables appropriately. Since the one-period utility function is homothetic, we scale all variables by the current level of aggregate income, using the notation $\hat{x}(g^t)$ to represent $x(g^t)/y(g^t)$. Note that bond position at history g^t is chosen at g^{t-1} and so is scaled by $y(g^{t-1})$. Utilities are scaled by $[y(g^t)]^{1-\gamma}$.

Let $V_b^i(\hat{b}, \hat{w}, n, m, g)$ be the optimal value of individual i , normalized by $y^{1-\gamma}$. The state variables are the agent's bond position b , the distribution of financial assets

⁹See Kubler and Schmedders (2003) for a discussion of this equilibrium concept. Krusell and Smith (1998) employ the same equilibrium concept.

$\hat{w} = (\hat{b}^1, \hat{b}^2)$, the aggregate growth state g , and the counters (n, m) that summarize agent 1's beliefs. This value function satisfies the following Bellman equation,

$$V_b^i(\hat{b}, w, n, m, g) = \max_{\hat{c}, \hat{b}'} \left[u(\hat{c}) + \beta \sum_{j \in \mathcal{G}} V^i(\hat{b}', w', j, n', m') g_j^{1-\gamma} \pi^i(g_j | g) \right] \quad (30)$$

where the maximization is subject to the budget constraint,

$$\hat{c} + q_b \hat{b}' = \hat{b}/g + \phi^i, \quad (31)$$

and borrowing limit,

$$\hat{b}' \geq -B^i. \quad (32)$$

Agents take the evolution of the aggregate states as given,

$$\hat{w}' = \Omega(\hat{w}, n, m, g, j), \quad (33a)$$

$$(n', m') = \mathcal{L}(n, m, j), \quad (33b)$$

where the function $\mathcal{L}(\cdot)$ summarizes Bayes updating and $\Omega(\cdot)$ is the equilibrium wealth-transition map, yet to be determined. Markets clear when $\forall(\hat{w}, n, m, g)$,

$$\hat{b}'^1(\hat{w}, n, m, g) + \hat{b}'^2(\hat{w}, n, m, g) = 0, \quad (34a)$$

$$\hat{c}^1(\hat{w}, n, m, g) + \hat{c}^2(\hat{w}, n, m, g) = 1. \quad (34b)$$

Let z denote the combined exogenous aggregate state (n, m, g) . A wealth-recursive Markov equilibrium is a list of functions $\{(\rho_c^i(\hat{w}, z), \rho_b^i(\hat{w}, z))_{i \in \{1,2\}}, q_b(\hat{w}, z), \Omega(\hat{w}, z)\}$ such that (i) the functions (ρ_c^i, ρ_b^i) solve problem (30-39) given the price system q_b ; (ii) goods and financial markets clear; and (iii) the wealth-transition map $\Omega(\hat{w}, z)$ is consistent with individual decisions: $\Omega(\hat{w}, z) = (\rho_b^1(\hat{w}, z), \rho_b^2(\hat{w}, z))$.

When $\gamma \geq 1$, the argument in Duffie, et al. (1994) can be adapted to show that

$$B^i \leq \min_{g_{t+1}} [g_{t+1} \hat{y}^i(g^{t+1})]$$

is a sufficient condition for the existence of a wealth-recursive Markov equilibrium.¹⁰ This borrowing limit is extremely tight, however, because it guarantees that debt can always be repaid within a single period, implying that debt never exceeds an agent's minimum income. An equilibrium might or might not exist when the debt limit is more lax than this. In what follows, we successfully compute equilibria for weaker debt limits, but we anticipate that the equilibrium would eventually break down as the debt limit is progressively relaxed. It is important to bear existence in mind. The example in Kubler and Polemarchakis (2002) shows that approximate equilibria, which satisfy equilibrium conditions only approximately, may exist even when an exact equilibrium does not.

¹⁰The assumption that $\gamma \geq 1$ is needed to insure that consumption is bounded away from zero.

4.2 Computing the equilibrium

We solve for the equilibrium policy and price functions using a projection method. For a given state $z = (n, m, g)$, we use cubic splines to approximate the bond-price map $q_b(\hat{b}, z)$, the decision rules $\rho_c^1(\hat{b}, z), \rho_c^2(\hat{b}, z), \rho_b^1(\hat{b}, z), \rho_b^2(\hat{b}, z)$, and the Lagrange multipliers on borrowing constraints, $\mu^1(\hat{b}, z), \mu^2(\hat{b}, z)$. The splines are defined on $[-B, B]$. Our grid has $500 \times 3 \times 20 \times 220 = 6,600,000$ nodes. The state space for the counter pair (n, m) is such that the probability of reaching the boundary after 500 periods is less than 0.0001.

Appendix C describes a number of tests for assessing the accuracy of the approximation. There we demonstrate that the maximum approximation error amounts to less than \$22 per \$10,000 of consumption, or 0.14% of total income. For comparison, the statistical discrepancy in the U.S. NIPA between 1929 and 2010 averaged at 0.54% of total income.

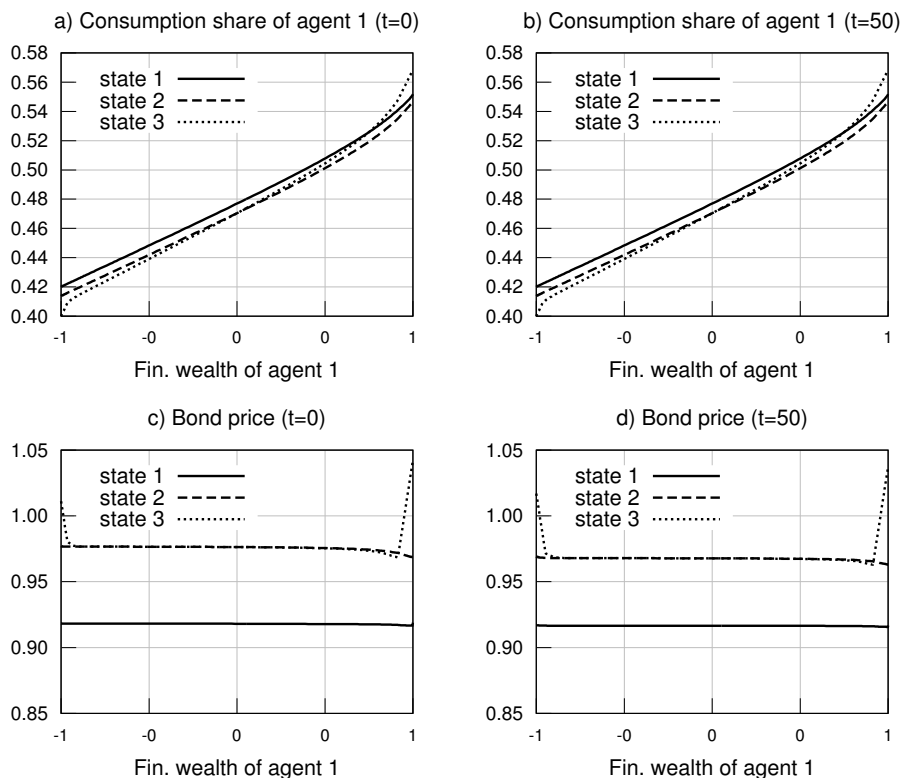


Figure 4: Decision rules and equilibrium price functions

Figure 4 illustrates decision rules for agent 1 as well as the equilibrium map for bond prices. The time-zero plot is based on the prior counters $(n, m) = (5, 5)$, while that for $t = 50$ conditions on the median value of the counters in that period, $(n, m) =$

(6, 11). A number of salient features emerge. Agent 1's consumption share depends on financial wealth, but it varies little across growth states, being only slightly lower in contractions than in expansions. In addition, consumption policies hardly change as time passes and agent 1 learns. Third, bond prices are higher in contractions and lower in expansions, but away from borrowing constraints they are insensitive to the distribution of financial wealth.

4.3 Simulation results

To compare the bond economy with the complete-markets model, we simulate consumption and savings outcomes and calculate equilibrium bond prices for the same endowment paths as in section 3. We initialize the financial wealth distribution at $(\hat{b}^1, \hat{b}^2) = (0, 0)$. Figures 5-7 portray the results.

The left column of figure 5 depicts ensemble averages for consumption and financial wealth, which are again expressed as proportions of aggregate income. In contrast to the complete-markets model, the less-well-informed agent 1 accumulates financial wealth in the bond economy, and better-informed agent 2 goes into debt. On average, agent 2's debt reaches approximately 50 percent of aggregate income (100 percent of individual income) after 20 years and approaches his borrowing limit of -1 (200 percent of individual income) after 60 years. Beyond that point, agent 2 rolls over his debt forever.

The underlying economics is primarily about precautionary saving and its effect on the equilibrium real-interest rate. Because agent 1 is pessimistic about deep contractions, his precautionary motives are stronger than they would be in the full-information rational-expectations version of this model. When markets are complete, type-1 consumers guard against deep contractions by purchasing Arrow securities which pay off in that state. They lose wealth on average because those states occur less often than they expect. Because that market is closed in the bond economy, they guard against deep contractions by buying risk-free bonds, thus accumulating wealth.¹¹

As shown in figure 6, agent 1's attempts to save more initially drive up the equilibrium bond price. Agent 2 is happy to borrow at this low real-interest rate in order to transfer consumption from the future to the present. The consumption profile for

¹¹To verify our intuition about precautionary saving, we also examined an economy in which type-1 agents are initially optimistic. As expected, this reverses the direction in which wealth is transferred. Because type-1 consumers are less concerned about deep contractions, their precautionary motives are weaker than under rational expectations. Hence they value bonds less than type-2 consumers and decumulate wealth along the transition path. Nevertheless, their financial wealth and consumption shares decline slowly. For instance, when the prior is $B(1, 19)$, the consumption share for type-1 consumers drops only 2% in a hundred years.

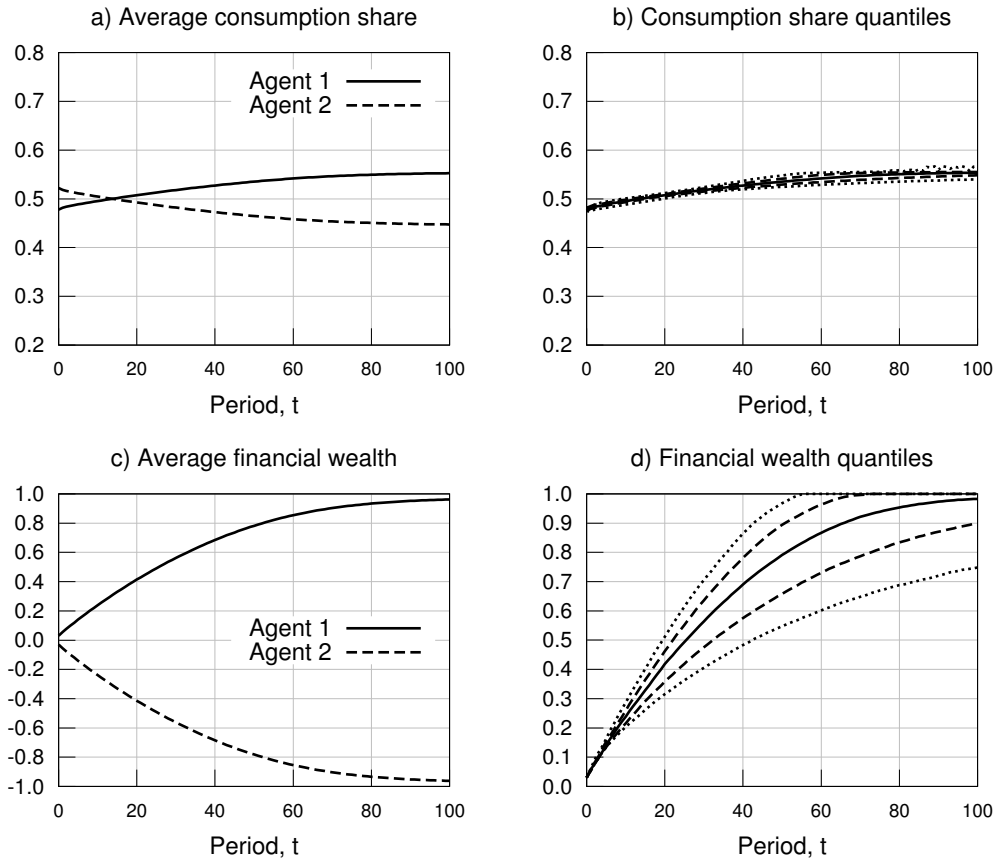


Figure 5: Wealth and consumption when markets are incomplete

agent 2 therefore slopes downward, while that of agent 1 slopes upward.¹² As agent 1's pessimism dissipates, he engages in less precautionary saving, and the bond price falls. Beliefs eventually converge, and the two agents settle into a homogeneous-beliefs equilibrium with a distribution of wealth favoring agent 1.

Thus, less-well-informed type-1 agents not only survive, they prosper. Indeed, the direction in which wealth is transferred is reversed. This provides yet another example in which market incompleteness disarms Blume and Easley's survival hypothesis. For instance, Tsyrennikov (2011) studies a model in which agents are allowed to renege on their obligations, as in Alvarez and Jermann (2000). Because agents face endogenous solvency constraints, financial markets are incomplete, yet all agents survive.¹³ Yet

¹²Appendix D provides analytical results along these lines for a two-period economy.

¹³Becker and Espino (2011) and Cao (2011) analyze closely-related environments and also demonstrate that all agents can survive when markets are incomplete. Becker and Espino focus on explaining asset-pricing anomalies. Cao analyzes an environment with collateral constraints *a la* Geanakoplos and Zame (2002).

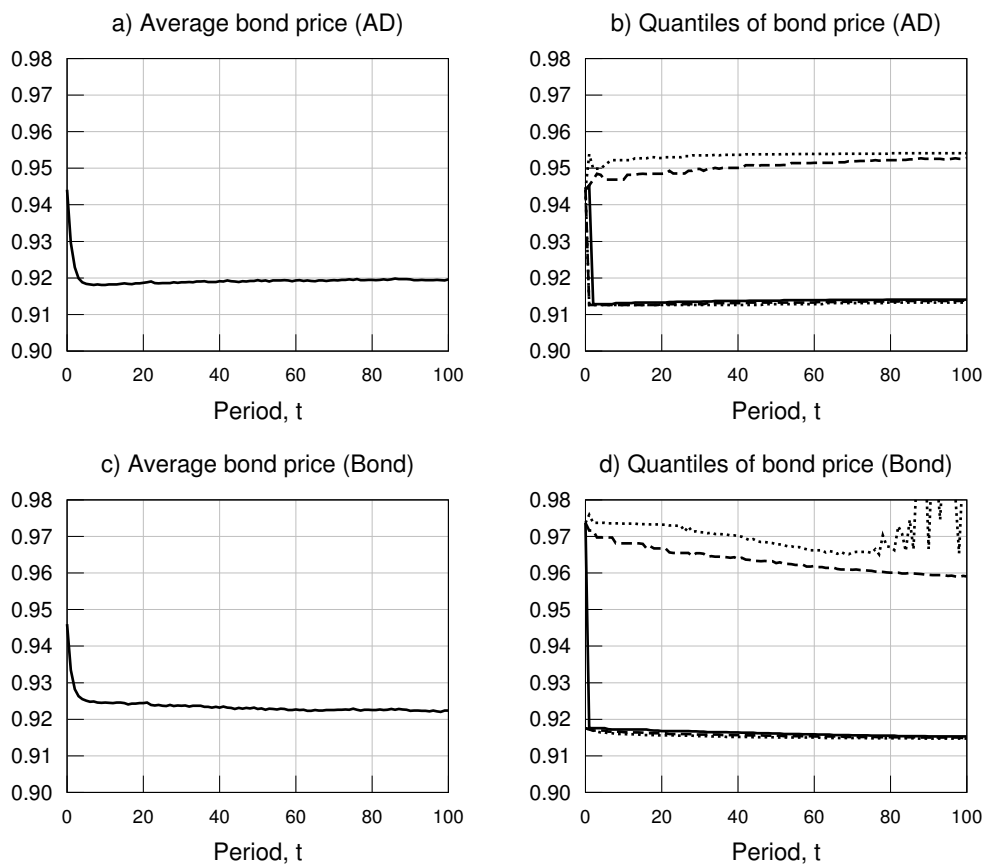


Figure 6: Bond price dynamics

another route to survival is that of Coury and Sciuabba (2010), who demonstrate that impatient traders with incorrect beliefs can survive and that their incorrect beliefs impact prices. We rule out default and assume that agents have identical discount factors, and we focus on the survival effects of financial-asset spanning.

The right column of figure 5 portrays quantiles of the respective cross-sample-path distributions for agent 1. The distributions of consumption and wealth are more tightly concentrated than under complete markets, and the distribution of consumption shares stays close to that of income shares. This reflects the fact that policy functions for consumption shares are less sensitive to recessions and contractions than in the complete-markets model.

Figure 7 makes the same point, depicting outcomes for the bond-economy on the endowment paths shown in figure 3. Whereas recessions and contractions induce big jumps in figure 3, they result in little ticks in figure 7. This occurs because agent 2 can no longer bet against deep contractions or in favor of mild recessions because no

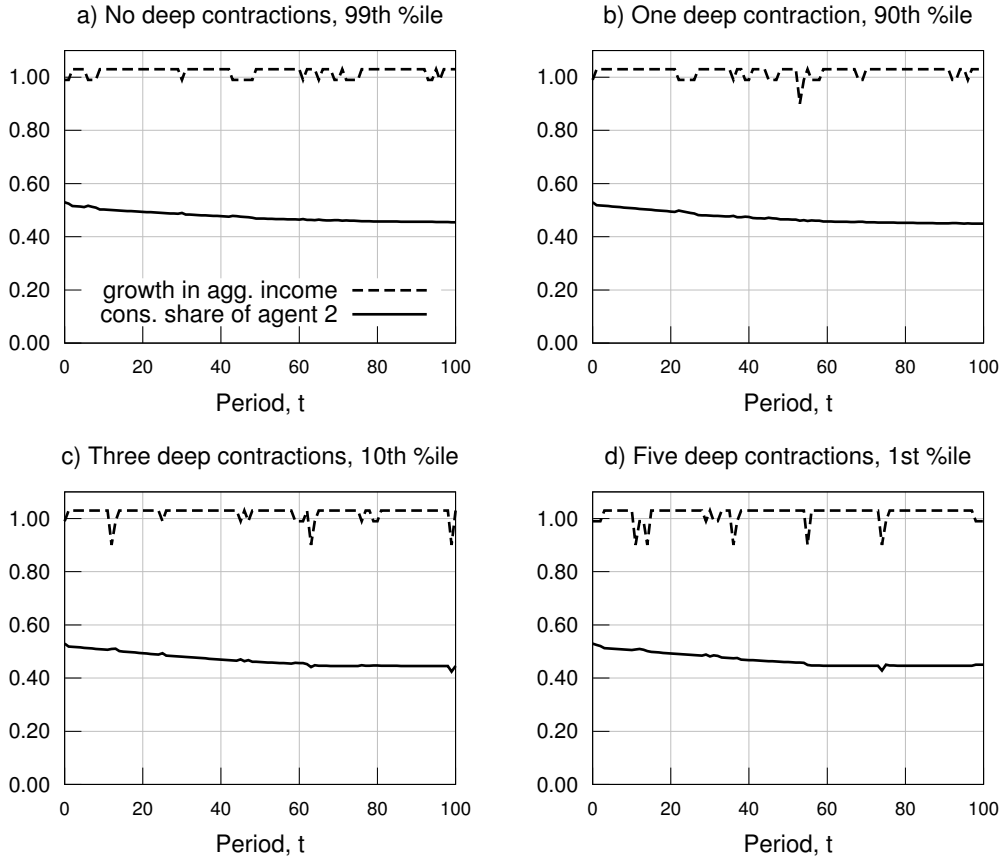


Figure 7: Selected sample paths when markets are incomplete

markets exist in which to place those bets. Shutting Arrow-security markets therefore deactivates the main mechanism through which the survival hypothesis operates in section 3. As a consequence, consumption shares in the bond economy never stray far from income shares, even on endowment paths in the tails of the complete-markets economy.

Although precautionary motives are important for understanding the bond economy, they arise here for a different reason than in standard models of incomplete markets (see Heathcote, Storesletten and Violante (2009)). To understand why, rewrite agent 1's welfare function as

$$\sum_{t=0}^{\infty} \sum_{g^t} \beta^t \frac{\pi^1(g^t)}{\pi^2(g^t)} \pi^2(g^t) u(c_t^1(g^t)) = E^2 \left[\sum_{t=0}^{\infty} \frac{\pi^1(g^t)}{\pi^2(g^t)} u(c_t^1(g^t)) \right]. \quad (35)$$

It follows that our model is equivalent to one in which agents have the same beliefs (those of agent 2) but in which agent 1 is hit by a preference shock of the form $\pi^1(g^t)/\pi^2(g^t)$.

If beliefs were homogenous, these “preference shocks” would be absent, agents would have no reason to trade, and consumption would always equal income. In this case, income variation would enter Euler equations only through marginal utility, and the degree of risk aversion would matter quantitatively for precautionary savings.




Figure 8: Consumption and wealth dynamics with different degree of risk aversion

Figure 8 degree of risk aversion also matters when precautionary motives arise from heterogeneity in beliefs. The figure portrays average financial wealth for agent 1 for two economies, one in which the coefficient of relative risk aversion $\gamma = 2$ and another in which $\gamma = 4$. Consumption and financial wealth dynamics change little when risk-aversion increases. Importantly, unlike the case with idiosyncratic income risk, financial wealth decreases at a slower pace when risk-aversion is high. This happens because higher risk aversion dampens the increase in the bond price, making it less attractive for type-2 agents to borrow.

4.4 Relaxing the debt limit

According to conventional wisdom, market incompleteness matters more because debt limits are typically tighter than the natural debt limit and less because of absence of spanning. In our model, however, absence of spanning seems crucial because this is what prevents agent 2 from exploiting the pessimism of agent 1. In this section, we try to disentangle the effects of debt limits v. absence of spanning by progressively relaxing the debt limit. If that is what really mattered, outcomes in the bond economy would move toward those under complete markets when the debt limit is relaxed.

Figure 9 plots the path for agent 1’s average financial wealth in the economies with debt limits of $B = 1$ and $B = 4$, respectively. Relaxing the debt limit moves the bond economy farther away from the complete-markets allocation. Type-1 consumers




Figure 9: Consumption and wealth dynamics under different borrowing restrictions

accumulate wealth even faster as the adverse effect of a binding borrowing limit is not felt during the first 100 years. Absence of spanning therefore seems to be the central force accounting for differences between complete-market outcomes and those in the bond economy. Differences in the debt limit matter quantitatively, but less so qualitatively.

4.5 A dogmatic-beliefs economy

Finally, to highlight the role of learning, we examine a heterogeneous-beliefs economy in which the pessimistic agent never learns. Since the pessimism of type-1 consumers now remains constant, they accumulate wealth faster relative to consumers who update their beliefs. Figure 10 confirms this intuition.




Figure 10: Consumption and wealth dynamics under different learning strategies

4.6 Welfare

Our model puts two Pareto-optimal allocations on the table. The complete-markets equilibrium is an ex-ante Pareto-optimum, i.e. with respect to subjective beliefs starting from date zero. On the other hand, the endowment is an ex-post Pareto-optimum, i.e. with respect to the asymptotic beliefs of the two agents. Closing Arrow-security markets and introducing a risk-free bond moves the allocation away from the ex-ante and toward the ex-post Pareto optimum.

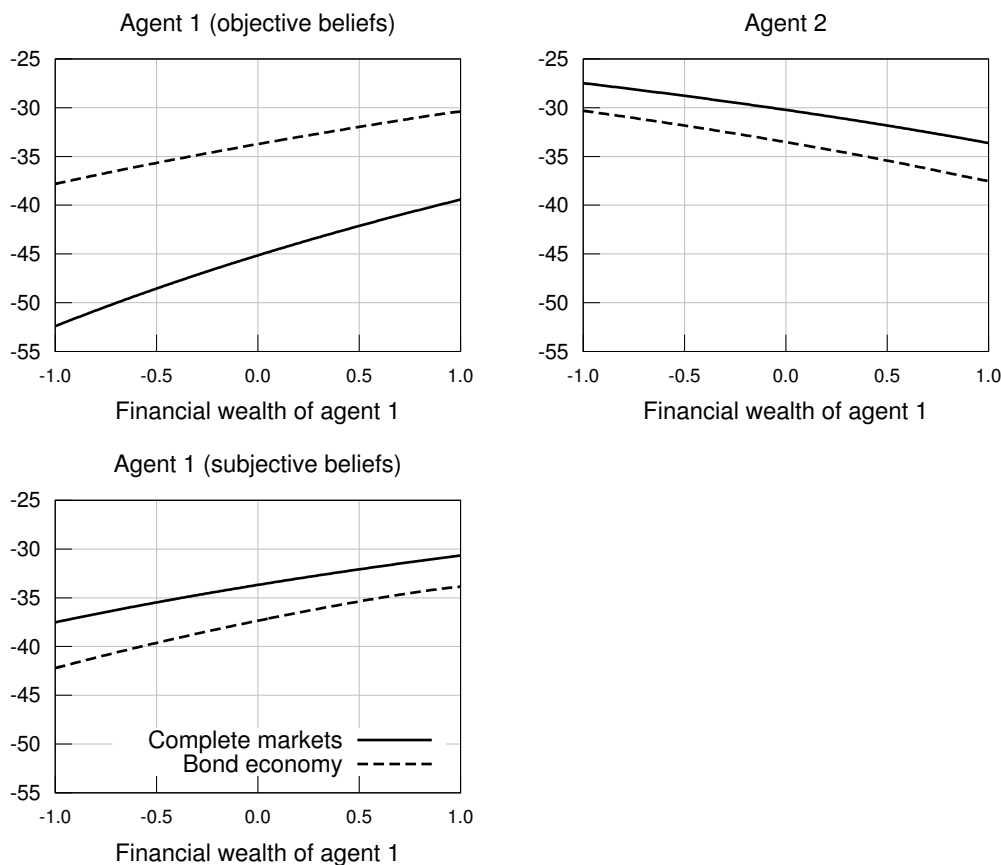


Figure 11: Welfare functions under complete and incomplete markets

The welfare effects are substantial. Figure 11 portrays expected utility as a function of agent 1's initial wealth. For agent 2, initial and asymptotic beliefs are the same, implying that welfare is the same whether evaluated from an ex-ante or ex-post perspective. In either case, closing Arrow-security markets reduces agent 2's welfare. For instance, for an initial wealth distribution of $(0, 0)$, closing Arrow-security markets reduces agent 2's welfare by an amount equivalent to a permanent 9.8% decrease in consumption. From an ex-ante point of view, agent 1 also believes that closing

Arrow-security markets will reduce her expected utility, but it increases when evaluated with respect to the true probabilities. Again focusing on an initial wealth distribution of $(0, 0)$, her ex-ante expected utility falls by an amount equivalent to a 9.8% decrease in consumption, but it increases by an amount equivalent to a 33% increase in consumption when evaluated with respect to true probabilities. Although neither agent would vote to close markets at date zero, agent 1 would regret her choice after the fact.

5 Concluding remarks

In ongoing research, we study a number of economies that occupy the middle ground between the polar cases contrasted here. In each economy, a single Arrow security trades along with a risk-free bond. That allows consumers to synthesize a portfolio of Arrow securities across the remaining two states, leaving them one asset short of a complete market. We study how wealth dynamics and survival depend on which market is closed. When the additional asset is an expansion-state Arrow security, we conjecture that the middle-ground economy will resemble the bond economy. But when the additional asset is one of the contraction-state securities, we suspect that the results will resemble those under complete markets. In this way we hope further to elucidate how the forces of the survival hypothesis depend on market structure.

6 References

- Alvarez, F. and U.J. Jermann, 2000, Efficiency, Equilibrium, and Asset Pricing with Risk of Default, *Econometrica* 68, 775-797.
- Barro, R.J., 2006, Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics* 121, 823-866.
- Barro, R.J. and J.F. Ursua, 2008, Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity*, 255-335.
- Barro, R.J., E. Nakamura, J. Steinsson, and J.F. Ursua, 2011, Crises and Recoveries in an Empirical Model of Consumption Disasters, unpublished manuscript.
- Becker, P.F. and E. Espino, 2011, A General Equilibrium Explanation for Financial Market Anomalies: Belief Heterogeneity Under Limited Enforceability, unpublished manuscript.
- 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.

Cao, D., 2011, Belief Heterogeneity, Wealth Distribution, and Asset Prices, unpublished manuscript.

Cogley, T. and T.J Sargent, 2009, Diverse Beliefs, Survival, and the Market Price of Risk, *Economic Journal* 119, 354-376.

Coury, T. and E. Sciuabba, 2010, Belief Heterogeneity and Survival in Incomplete Markets, *Economic Theory*, forthcoming.

Duffie, D., J. Geanakoplos, A. Mas-Colell, A. McLennan, 1994, Stationary Markov Equilibria, *Econometrica* 62, 745-781.

Friedman, M., 1953, *Essays in Positive Economics*, University of Chicago Press.

Geanakoplos, J. and W.R. Zame, 2002, Collateral and the Enforcement of Intertemporal Contracts, unpublished manuscript.

Heathcote, J., K. Storesletten, and G. Violante, 2009, Quantitative Macroeconomics with Heterogeneous Households, *Annual Review of Economics*, Vol. 1, 319-354.

Kogan, L. S.A. Ross, J. Wang, and M.W. Westerfield, The Price Impact and Survival of Irrational Traders, 2006, *Journal of Finance* LXI, 195-228.

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

Kubler, Felix, And H. M. Polemarchakis, 2002, Markov Equilibria In Overlapping Generations, Brown University mimeo.

Kubler, F. and K. Schmedders, 2003, Stationary Equilibria in Asset-Pricing Models with Incomplete Markets and Collateral, *Econometrica*, 71(6), 1767-1793.

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

A A recursive formulation for the complete-markets model

Because the aggregate endowment is non-stationary, we begin by scaling the variables appropriately. Since the one-period utility function is homothetic, we scale all variables by the current level of aggregate income. We use $\hat{x}(g^t)$ to denote $x(g^t)/y(g^t)$. The bond position at history g^t is chosen at g^{t-1} and so is scaled by $y(g^{t-1})$. Utilities are scaled by $y(g^t)^{1-\gamma}$.

We solve for a wealth-recursive Markov equilibrium. In a wealth-recursive Markov equilibrium, the distribution of financial wealth $\hat{w} = (\hat{a}^1, \hat{a}^2)$ is an endogenous state variable. The other state variables for individual i are his own financial assets \hat{a}^i , the

current aggregate growth state g , and the counters m, n that are sufficient statistics for summarizing agent 1's beliefs.

Let $V_{cm}^i(\hat{a}, \hat{w}, n, m, g)$ be the optimal value of individual i , normalized by $y^{1-\gamma}$. This value function must satisfy the following Bellman equation,

$$V_{cm}^i(\hat{a}, \hat{w}, n, m, g) = \max_{\hat{c}, \hat{a}'_j} \left[u(\hat{c}) + \beta \sum_{j \in \mathcal{G}} V_{cm}^i(\hat{a}'_j, \hat{w}', g_j, n', m') g_j^{1-\gamma} \pi^i(g_j|g) \right], \quad (36)$$

where expectations are taken with respect to agent i 's predictive distribution, $\pi^i(g_j|g)$. The maximization is subject to the budget constraint,

$$\hat{c} + \sum_{j \in \mathcal{G}} Q_j(\hat{w}, n, m, g) g_j \hat{a}'_j = \hat{a} + \phi^i, \quad (37)$$

and the natural borrowing limit,

$$\hat{a}'_j \geq -\hat{B}^i(\hat{w}', n', m', g_j). \quad (38)$$

The borrowing limit can be computed recursively by iterating on

$$\hat{B}^i(\hat{w}, n, m, g) = \phi^i + \sum_{j \in \mathcal{G}} Q_j(\hat{w}, n, m, g) \hat{B}^i(\hat{w}', g_j, n', m').$$

Agents take the evolution of aggregate states as given, including those for aggregate financial wealth and the counters,

$$\hat{w}' = \Omega(\hat{w}, n, m, g, g'), \quad (39a)$$

$$(n', m') = \mathcal{L}(n, m, g'). \quad (39b)$$

The function \mathcal{L} is given in equation (9), and the function Ω is the part of the equilibrium map, which is yet to be computed.

Markets clear when

$$\hat{a}_j^{1'}(\hat{w}, n, m, g) + \hat{a}_j^{2'}(\hat{w}, n, m, g) = 0, \quad \forall j \in \mathcal{G}, \quad (40a)$$

$$\hat{c}^1(\hat{w}, n, m, g) + \hat{c}^2(\hat{w}, n, m, g) = 1, \quad (40b)$$

for all \hat{w}, n, m, g .

B How pessimistic is the prior?

To assess whether the assumed degree of pessimism is plausible, we consider whether consumers could easily distinguish it statistically from the beliefs of type-2

consumers based on a finite data sample. Let M_i denote the prior probability model of types $i = 1, 2$. Following Hansen and Sargent (2008), we measure a consumer's ability to distinguish models by the detection-error probability. Assuming prior weights of $1/2$ on each of the models, the detection-error probability is

$$dep = 0.5[prob(\mathcal{L}_2(X) > \mathcal{L}_1(X)|M_1) + prob(\mathcal{L}_1(X) > \mathcal{L}_2(X)|M_2)] \in [0, 0.5],$$

where $\mathcal{L}_i(X)$ is the likelihood of sample X evaluated using model M_i . The two models are easy to distinguish when the detection-error probability is close to zero. Using 100,000 simulated samples of length $T = 50$ years, we obtain $dep = 12.8\%$, a figure comparable to those used for calibrating uncertainty aversion in models of robustness.

C Testing the accuracy of the approximation for the bond economy

The solution consists of the consumption $\rho_c^i(\hat{b}, n, m, s)$ and the bond investment $\rho_b^i(\hat{b}, n, m, s)$ policy functions, the Lagrange multipliers associated with borrowing limits $\rho_\mu^i(\hat{b}, n, m, s)$ 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 $e_\rho = 10^{-5}$. Decreasing this threshold improves accuracy only marginally.¹⁴

We verify the computed solution on a grid that is 5 times denser than the one used to compute policy functions. Verification procedure consists of computing the following errors:

$$\begin{aligned} e_1(\hat{b}, n, m, s) &= 1 - \frac{1}{\rho_c^1(\hat{b}, n, m, s)} \left[\frac{q_b(\hat{b}, n, m, s)}{E[(\rho_c^1(\hat{b}, n, m, s)g(s))^{-\gamma}] + \rho_\mu^1(\hat{b}, n, m, s)} \right]^{1/\gamma}, \\ e_2(\hat{b}, n, m, s) &= 1 - \frac{1}{\rho_c^2(\hat{b}, n, m, s)} \left[\frac{q_b(\hat{b}, n, m, s)}{E[(\rho_c^2(\hat{b}, n, m, s)g(s))^{-\gamma}] + \rho_\mu^2(\hat{b}, n, m, s)} \right]^{1/\gamma}, \\ e_3(\hat{b}, n, m, s) &= 1 - \frac{1}{\rho_c^1(\hat{b}, n, m, s)} [0.5 + \hat{b} - q_b(\hat{b}, n, m, s)\rho_b^1(\hat{b}, n, m, s)]. \end{aligned}$$

The errors answer the following question: “what fraction should be added/subtracted from an agent's consumption so that the respective equilibrium condition holds exactly?” The first two equations are the consumption Euler equations for agent 1 and 2 respectively. The third equation is the budget constraint of agent 1. Feasibility constraint is imposed on the solution; so, the budget constraint of agent 2 hold exactly.

¹⁴This means that approximation errors begin to dominate numerical errors.

The “free parameter” in the numerical analysis is the number of grid points for the bond position. Starting from 50, we increase the number of grid points for the bond position until a sufficient level of accuracy is achieved. With 500 grid points the errors are smaller than 0.22% of the average consumption.¹⁵ For comparison the statistical discrepancy in the U.S. NIA between 1929 and 2010 averaged at 0.54% of the *total income*.¹⁶ Errors are larger for higher levels of risk-aversion but stay below 0.35% of individual consumption as long as $\gamma \leq 5$ (holding the number of grid points at 500).

Figure 12 plots *maximal* errors (over all possible pairs of counters) for each (b^1, s) pair for $\gamma = 2$ and 500 grid points. These errors are largest when $g = 3$ which is expected.

We also conduct another test of solution accuracy. The error in equilibrium conditions for the solution with 100 grid points is 1.13% of consumption; so, about 5 times that of the 500-grid-point solution. We ask: “how different are simulated paths for solutions with 100 and 500 grid points?” The test path that we chose is a deterministic repetition of the sequence 1, 1, 1, 1, 1, 2, 2, 2, 2, 3. In this path states 2 and 3 are over-represented relative to the true measure. It turns out (see figure 13) that, even though differences exist, paths do not diverge and the difference stays bounded.

D A two-period example

Consider the following two-period setup. In period 0, aggregate income is 1. In period 1, it is $e_l > 0$ with probability π_l and $e_h > e_l$ with complementary probability. Each agent receives half of aggregate income. Agent 2 has correct beliefs, and agent 1 is pessimistic. He assigns probability $\pi_l + d, d \in [0, 1 - \pi_l)$ to the low endowment state.

Agents rank a consumption stream $(c_0^i, c_{1l}^i, c_{1h}^i)$ according to:

$$u(c_0^i) + \pi_l^i u(c_{1l}^i) + \pi_h^i u(c_{1h}^i),$$

where u is a strictly increasing and strictly concave function.

Financial markets trade only a risk-free bond. Clearly, when $d = 0$, competitive equilibrium is autarkic. How does competitive equilibrium depend on d ?

Let us analyze the bond-demand function of agent i for a given price q_b . Her consumption-Euler equation is

$$q_b u'(1/2 - q_b b^i) = \sum_{s \in \{l, h\}} \pi_s^i u'(e_s/2 + b^i)$$

¹⁵This amounts to 22\$ for every 10,000\$ of consumption.

¹⁶Model errors when normalized by the total income are less than 0.14%.

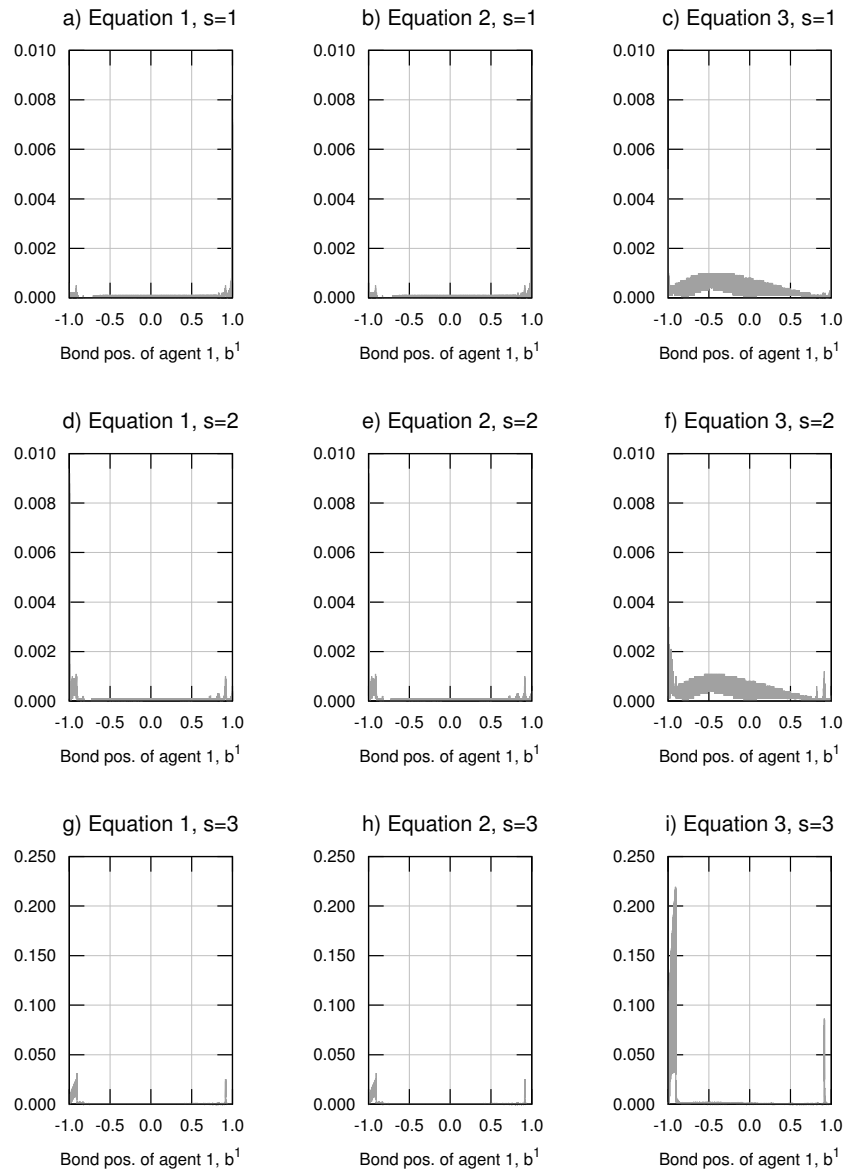


Figure 12: Bond economy: off the grid errors

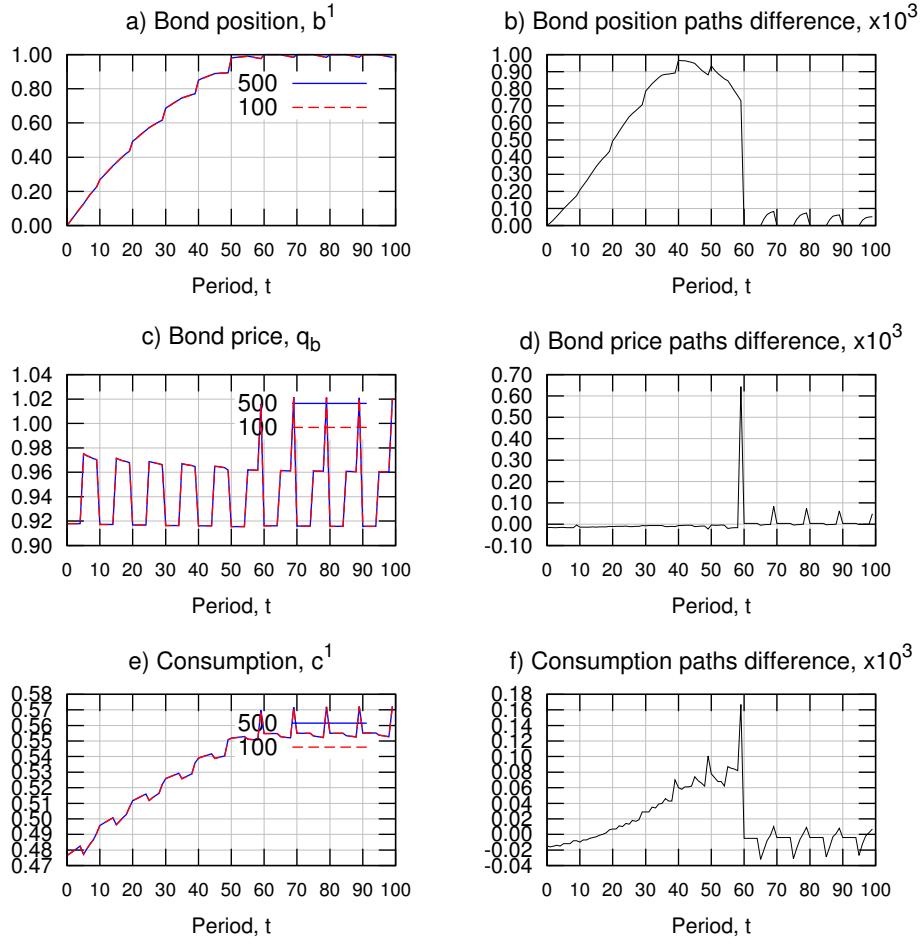


Figure 13: Comparing solutions with different degree of accuracy

For a fixed q_b how does b^1 change with d ? It is easy to show that $\partial b^1(d, q_b)/\partial d > 0$. Since $b^2(q_b) = b^1(0, q_b)$ we must have $b^1(d, q_b) > b^2(q_b), \forall q_b, d > 0$. Then the bond-market clearing condition implies: $b^1(d, q_b^*) = -b^2(q_b^*) > 0$, where q_b^* is the equilibrium bond price. This, in turn, implies that $c_0^1 < c_0^2$ and $c_{1s}^1 > c_{1s}^2, s \in \{l, h\}$. Finally, using the consumption Euler-equation again, we can show that the bond price must be higher than in the homogenous beliefs case. Using strict concavity of u we get:

$$q_b = \frac{E^2[u'(e_s/2 + b^2)]}{u'(1/2 - q_b b^2)} > \frac{E^2[u'(e_s/2)]}{u'(1/2)}.$$

The term on the right hand side equals the equilibrium bond price when beliefs are homogeneous. This proves the following claim.

Claim. In the two-period economy, relative to the competitive equilibrium with homogeneous beliefs, the pessimistic agent holds a larger position in bonds and consumes

less in period 0 and more in period 1. In addition, the equilibrium bond price is higher.

E Model solution

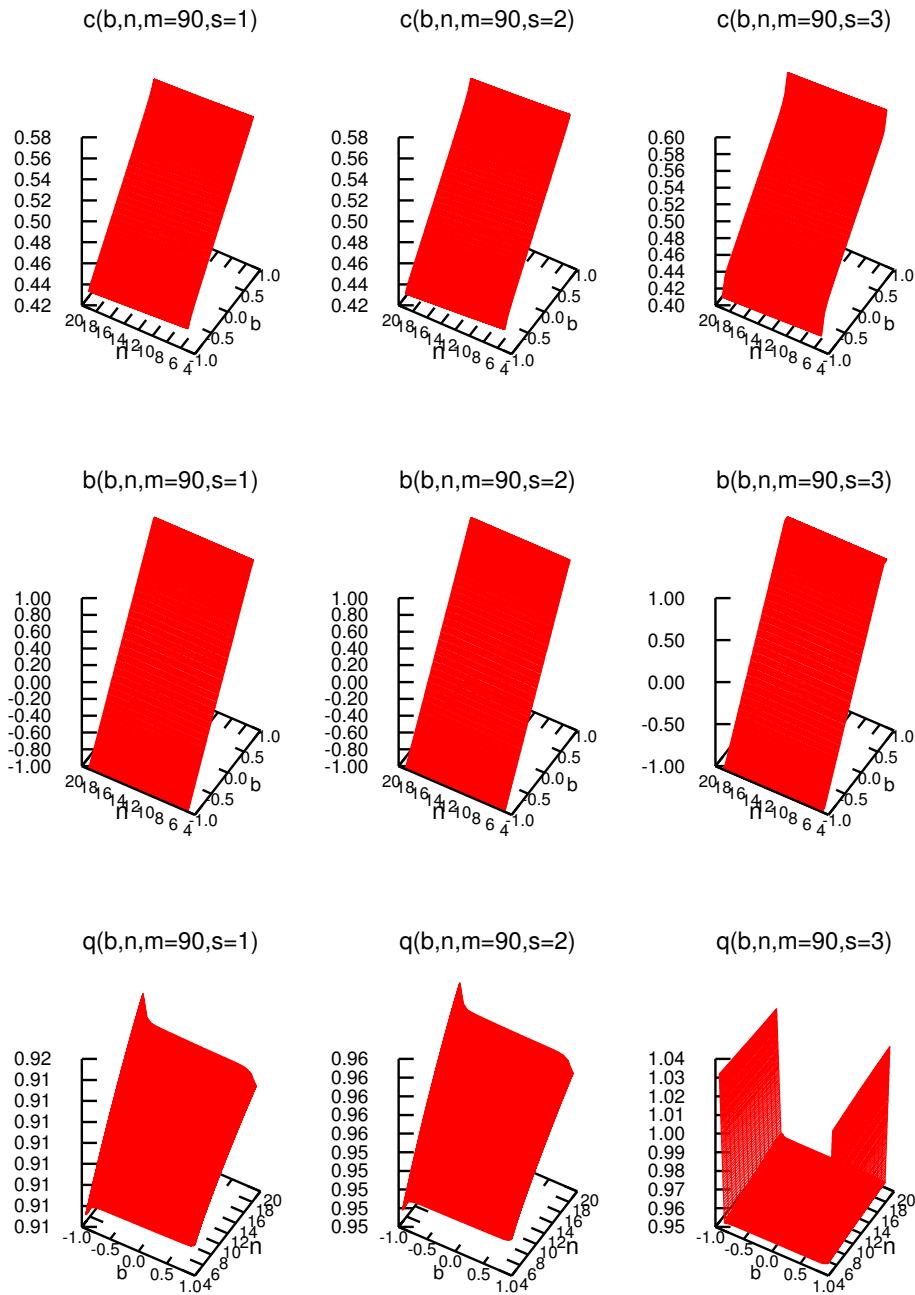


Figure 14: Decision rules and equilibrium price functions