

Diverse Beliefs, Survival, and the Market Price of Risk*

Timothy Cogley[†] Thomas J. Sargent[‡]

Revised: August 1, 2008

Abstract

We study prices and allocations in a complete-markets, pure-exchange economy in which there are two types of agents with different priors over infinite sequences of the aggregate endowment. Aggregate consumption growth evolves exogenously according to a two-state Markov process. The economy has two types of agents, one that learns about transition probabilities and another that knows them. We examine allocations, the market price of risk, and the rate at which asset prices converge to values that would be computed under the assumption that all agents know the transition probabilities.

KEY WORDS: Walrasian equilibrium, Bayes' Law, heterogeneity, market price of risk, survival.

1 Introduction

Cogley and Sargent (2008) showed that market prices of risk would be high in an economy with a risk-neutral Bayesian representative agent who learns about the parameters of a transition matrix for aggregate consumption growth starting with a pessimistic prior in 1933.¹ This paper studies the robustness of that finding to a perturbation that adds a small fraction of agents who know the parameters of the transition matrix. Traders participate in a Walrasian equilibrium in which they do

*We thank V.V. Chari, Michael Dotsey, Patrick Kehoe, Larry Jones, Aubhik Khan, Narayana Kocherlakota, Erzo Luttmer, Fabrizio Perri, Monika Piazzesi, Martin Schneider, and Andrew Scott for helpful discussions and comments on earlier drafts. We also thank an anonymous referee for useful suggestions. Sargent thanks the National Science Foundation for research support.

[†]University of California, Davis. Email: twcogley@ucdavis.edu

[‡]New York University and Hoover Institution. Email: ts43@nyu.edu

¹Friedman and Schwartz (1963) used a related evaporating pessimism hypothesis to explain the behavior of investors' portfolio choices during the two decades after the Great Depression.

not infer information from prices (see Grossman (1981)).^{2,3} Under what we assume to be the true data-generating mechanism, the survival-of-the-fittest force analyzed by Blume and Easley (2006) causes the more knowledgeable agents' influence on equilibrium prices to grow over time along with their wealth. We study how quickly that dissipates the effects of the initial pessimism of the less informed agents on equilibrium prices.⁴

To give the survival mechanism a large scope to moderate the effects of initial pessimism, we assume complete markets that allow agents to make trades of claims to wealth that are motivated *solely* by the different subjective probabilities they put on future states. That gives the agents many opportunities to place bets that over time stochastically increase the share of wealth of the traders who know transition probabilities.

We solve a Pareto problem to compute competitive equilibrium prices and allocations, thereby implicitly defining an initial allocation of wealth. We study how the market price of risk evolves as a function of the relative Pareto weight on the better-informed agent. Among other things, we want to know how large the Pareto weight on the better-informed agent has to be in order to eradicate the effects of the initial pessimism that Cogley and Sargent (2008) attributed to a representative agent.

Because we specify the data generating mechanism and agents' beliefs so that the truth is in the support of both agents' beliefs, the survival analysis of Blume and Easley (2006) lets both types of agents have positive wealths in the limit. However, the agents' ultimate shares of aggregate consumption are random variables that depend on the history of the growth rate of aggregate consumption. We calculate probability distributions of these shares for various horizons.⁵

2 The model

The aggregate endowment process is the same as in Cogley and Sargent (2008). The two types of agents have identical one-period utility functions but different priors that imply different beliefs about histories of aggregate growth rates. Agent 1 uses Bayes' law to learn about transition probabilities. Bayes' law induces a different probability measure over histories of aggregate consumption growth rates for agent 1

²We endow our two types of agents with different priors. For reasons discussed in subsections 2.1 and 2.2, if we had endowed the two types of agents with a common prior, our robustness exercise would be trivial: the presence of the more knowledgeable type 2 agent would cause the pessimism of the type 1 agent to evaporate completely at time 0.

³Kogan et al. (2006) study a setting in which most agents know the data generating mechanism and form forecasts accordingly but a low-wealth small measure of irrational agents still has important influences on prices.

⁴In the representative agent economy of Cogley and Sargent (2008), Bayes' law also dissipates pessimism.

⁵Beker and Espino (2008) study limiting portfolios and volume in a related economy.

than is believed by agent 2, who knows the true transition probabilities from the out-set. Agents take prices as given and trade history-contingent claims to consumption for all finite histories.

2.1 The endowment process, beliefs, and trading opportunities

The consumption good arrives exogenously and is nonstorable, so all current-period output is consumed immediately. Realizations for gross consumption growth follow a two-state Markov process with high and low growth states, denoted \bar{g}_h and \bar{g}_l , respectively. The Markov chain has a transition matrix F , where $F_{ij} = \text{Prob}[g_{t+1} = \bar{g}_j | g_t = \bar{g}_i]$. We let $g^t = [g_t, g_{t-1}, \dots, g_0]$ denote a history of aggregate growth rates. At time $t \geq 0$, agents of both types have the information set g^t .

The two types of agents begin with *different* priors in the form of joint densities over (F, g^∞) . The marginal distribution over F for the type 1 agent who learns is nontrivial and includes in its support the true F that actually governs the data. For the type 2 agent, the marginal distribution over F is concentrated on the true F .⁶

2.2 Walrasian versus rational expectations equilibrium

We study a complete markets economy with time 0 trading of a complete set of claims to consumption contingent on finite histories g^t for all $t \geq 0$. We study what Grossman (1981) calls a Walrasian equilibrium, 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 of future g 's. We do this by assuming that agents have different priors over g^∞ and common information sets, and not a common prior with different information sets.⁷

If we had made a different assumption about agents' beliefs, it would be have been appropriate to study an equilibrium in which traders do extract information from observed prices, what Grossman (1981) calls a rational expectations equilibrium.⁸ In particular, we could have started our two types of agents with a *common* prior over (F, g^∞, s) , where s is a signal that the type 2 agent but not the type 1 agent receives

⁶The term 'true F ' has no special status in the theory underlying the Walrasian competitive equilibrium prices and allocations. Two distinct probability densities over infinite sequences of consumption growth are simply ingredients of the two types of agents' preferences. We do not have to use the term 'true F ' until we simulate equilibrium allocations and prices under a particular probability distribution.

⁷Thus, we do not adhere to the "Harsanyi Doctrine" that asserts that differences in individuals' beliefs are to be attributed entirely to differences in information.

⁸Pearlman and Sargent (2005) study the rational expectations equilibrium for the private information environment of Townsend (1983) and show that firms can extract *all* of other firms' private information from prices.

Table 1: Maximum Likelihood Estimates of the Consumption Process (Cecchetti et al. (2000))

	F_{hh}	F_{ll}	μ_h	μ_l	σ_ε
Estimate	0.978	0.515	2.251	-6.785	3.127
Standard Error	0.019	0.264	0.328	1.885	0.241

at time 0. We could let the signal reveal the value of F and assume that while a type 1 consumer observes g^t and equilibrium prices at $t \geq 0$, a type 2 agent observes s as well as g^t and equilibrium prices. Under these assumptions about priors and information, rational-expectations equilibrium prices would immediately reveal s to the type 1 agents, and the two types of agents would have a common posterior over (g^∞, F) (see Milgrom and Stokey (1982)). In that case, the pessimism of the type 1 agent would completely evaporate at time 0.⁹

Since we are interested in continuing effects of gradual learning, we chose our Walrasian specification to keep learning alive.

2.3 Calibration

We calibrate F from estimates reported by Cecchetti et al. (2000) who estimated a hidden Markov model for aggregate consumption growth,

$$\Delta \ln C_t = \mu(s_t) + \varepsilon_t, \quad (1)$$

where s_t is an indicator variable that records whether consumption growth is high or low and ε_t is an identically and independently distributed normal random variable with mean 0 and variance σ_ε^2 . They estimated the model by maximum likelihood using data on annual per capita US consumption for the period 1890-1994. We reproduce their results in table 1.

The high-growth state is persistent, and the economy spends most of its time there. Contractions are more severe than typical post-WWII recessions, with a mean decline of 6.785 percent per annum. Moreover, because the low-growth state is moderately persistent, a run of contractions can occur with nonnegligible probability, producing something like the Great Contraction. For example, conditional on a contraction

⁹Sometimes differences in priors can be reformulated in terms of common priors but different information sets. Thus, we could have started our two types of agents with a common prior over (g^∞, F) , then given the type 1 agent the information set g^t at t and the type 2 agent both g^t and knowledge of particular tail events for the $\{g_t\}$ process in the form of limiting values of the empirical fractions of transitions from high to low growth and from low to high growth states. Knowing those tail events is tantamount to knowing F . In the text, we allow agents to exchange claims contingent only on finite histories g^t for all $t \geq 0$, so there is no market in which agents can bet on tail events.

having begun, the probability that it will last 3 more years is 14 percent, and if that were to occur, the cumulative fall in consumption would amount to 25 percent.

We simplify the endowment process by suppressing the Gaussian innovation ε_t and assume instead that gross consumption growth follows a two-point process,

$$\begin{aligned} g_t &= \bar{g}_h \equiv 1 + \mu_h/100 & \text{if } s_t = 1, \\ &= \bar{g}_l \equiv 1 + \mu_l/100 & \text{if } s_t = 0. \end{aligned} \quad (2)$$

We retain the point estimates of μ_h and μ_l made by Cecchetti et al. (2000) as well as their estimates of the transition probabilities F_{hh} and F_{ll} .

We assume that this model represents the true process for consumption growth. We also assume that both agents know \bar{g}_h and \bar{g}_l and that a type 2 agent knows the true transition matrix F . A type 1 agent uses Bayes' law to learn about F .

2.4 Preferences over consumption plans

A *consumption plan* for agent i is a sequence of functions $C_{it}(g^t)$, $t \geq 0$, whose time t component maps a time t history g^t into agent i 's time t consumption. Two infinitely-lived consumers share the same isoelastic one-period utility function $u(C_{it}) = \frac{C_{it}(g^t)^{1-\alpha} - 1}{1-\alpha}$ and order consumption plans according to the expected utility functionals

$$U_i = \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \frac{C_{it}(g^t)^{1-\alpha} - 1}{1-\alpha} pr_i(g^t), \quad (3)$$

or

$$U_i = E_{i0} \sum_{t=0}^{\infty} \beta^t \frac{C_{it}^{1-\alpha} - 1}{1-\alpha}, \quad (4)$$

where $pr_i(g^t)$ is agent i 's subjective probability over g^t and E_{it} denotes conditional expectation with respect to $pr_i(g^t)$. The parameters α and β are the coefficient of relative risk aversion and subjective discount factor, respectively, and they are common across agents. In our simulations, we set $\alpha = 2$ and $\beta = 1.03^{-1}$.

2.5 The Pareto Problem

For some initial distribution of wealth, a competitive equilibrium allocation solves a Pareto problem:

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{\infty} \beta^t \sum_{g^t} \left\{ \lambda u[C_{1t}(g^t)] pr_1(g^t) + (1-\lambda) u[C_{2t}(g^t)] pr_2(g^t) \right. \\ & \left. + \mu(g^t) [C_t(g^t) - C_{1t}(g^t) - C_{2t}(g^t)] \right\}, \end{aligned}$$

where λ is the Pareto weight on the Bayesian consumer, g^t represents a history of states through date t , and $\mu(g^t)$ is the Lagrange multiplier on the aggregate resource

constraint at date t , history g^t . The Pareto planner distributes consumption so that the ratio of marginal utilities equals the ratio of Pareto weights,

$$\frac{u'(C_{2t}(g^t))pr_2(g^t)}{u'(C_{1t}(g^t))pr_1(g^t)} = \frac{\lambda}{1 - \lambda}. \quad (5)$$

With isoelastic utility, equation (5) can be solved to express consumption for agent 1 as a history-dependent share of the aggregate endowment,

$$C_{1t}(g^t) = \frac{\phi(g^t)}{1 + \phi(g^t)} C_t(g^t) \quad (6)$$

where

$$\phi(g^t) = \left[\frac{\lambda}{1 - \lambda} \frac{pr_1(g^t)}{pr_2(g^t)} \right]^{1/\alpha}. \quad (7)$$

The more informed type 2 agent gets the remainder,

$$C_{2t}(g^t) = \frac{1}{1 + \phi(g^t)} C_t(g^t). \quad (8)$$

With common beliefs (i.e., $pr_1(g^t) = pr_2(g^t)$), $\phi = [\lambda/(1 - \lambda)]^{1/\alpha}$, implying that the agents get constant shares of aggregate consumption. Diverse beliefs alter this common-beliefs allocation by shifting resources toward agent 1 in histories that he thinks are more likely than agent 2. In our model, agent 1 is initially pessimistic, overestimating the probability of a contraction state. Hence, along a sample path, agent 1 gets more consumption relative to the common-beliefs benchmark in the contraction state and less in the expansion state. Agent 2's allocation is twisted in the opposite direction. Since expansions are the norm, agent 2 frequently benefits from agent 1's pessimism. But in a contraction, the consumption of agent 2 falls by more than the aggregate endowment. Thus, the presence of a pessimistic agent alters equilibrium prices in ways that increase the consumption risk that the more informed agent chooses to bear in a competitive equilibrium.

2.6 Arrow security prices

As usual, we can support a Pareto optimal allocation with an appropriate initial distribution of wealth, sequential trading of one-period Arrow securities, and a set of 'natural' limits on the quantities of Arrow securities that can be issued in each history, date pair (for example, see Ljungqvist and Sargent (2004, ch. 8)). In the present context, it suffices to trade two Arrow securities each period, one that pays one unit of aggregate consumption when $g_{t+1} = \bar{g}_h$ and zero units otherwise, and the other paying off when $g_{t+1} = \bar{g}_l$. Their prices are denoted Q_{ht} and Q_{lt} , respectively. The gross return on the high-growth state asset is $1/Q_{ht}$ when $g_{t+1} = \bar{g}_h$ and 0 otherwise. Arrow securities prices are

$$Q_{ht} = \beta \frac{u'(C_{it+1}(\bar{g}_h, g^t))}{u'(C_{it})} pr_i(g_{t+1} = \bar{g}_h | g^t) \quad (9)$$

and

$$Q_{lt} = \beta \frac{u'(C_{it+1}(\bar{g}_l, g^t))}{u'(C_{it})} pr_i(g_{t+1} = \bar{g}_l | g^t). \quad (10)$$

2.7 Financial wealth

Given a division of the aggregate endowment into an amount $y_t^i(g^t)$ assigned to an agent of type i at time t and history g^t , we can use implied Arrow-Debreu prices to define a sequence of financial wealths that equal equilibrium quantities of one-period Arrow securities (see Ljungqvist and Sargent (2004, pp. 224-233)). For $t \geq 0$, let $q_\tau^t(g^\tau)$ be the time t Arrow-Debreu price for a claim to a unit of consumption at date τ after history g^τ and let $c_\tau^i(g^\tau)$ be the consumption of a type i agent at date τ after history g^τ . Then financial wealth at date t after history g^t is

$$W_t^i(g^t) = \sum_{\tau=t}^{\infty} \sum_{\hat{g}^\tau | \hat{g}^t = g^t} q_\tau^t(\hat{g}^\tau) [c_\tau^i(\hat{g}^\tau) - y_\tau^i(\hat{g}^\tau)]$$

where $\hat{g}^\tau | \hat{g}^t = g^t$ is a history g^τ , $\tau \geq t$ whose partial history up to t is g^t .

In principle, for a given assignment of individual endowments $\{y_t^i(g^t)\}$ for $i = 1, 2$, we can compute the financial wealths of our two types of consumers recursively, but for our model it is computationally demanding and we do not do so because we are principally interested in allocations and market prices of risk, which are easier to compute.

2.8 Alternative representations of the unique SDF

Because markets are complete, there is a unique stochastic discount factor with respect to the informed type 2 agent's probability measure. The stochastic discount factor has the alternative representations

$$SDF = m_{2t+1} = m_{1t+1} \frac{pr_1(g_{t+1}|g^t)}{pr_2(g_{t+1}|g^t)} \quad (11)$$

where $m_{it+1} = \beta(C_{it+1}(g^{t+1})/C_{it}(g^t))^{-\alpha}$ denotes the intertemporal marginal rate of substitution for consumer i .¹⁰

¹⁰Cogley and Sargent (2008) featured the learning wedge captured by the likelihood ratio $pr_1(g_{t+1}|g^t)/pr_2(g_{t+1}|g^t)$. In their economy, the representative agent is of type 1, there is no type 2 informed agent, but $pr_2(g_{t+1}|g^t)$ represents the true data generator. Bossaerts (2002, 2004) also uses a ratio of the likelihood of a Bayesian agent to the likelihood for a data-generating process to model prices of risk. Hansen (2007) and Hansen and Sargent (2007a) develop models of robust asset pricing that feature another probability ratio, $pr_{wc}(g_{t+1}|g^t)/pr_{dgp}(g_{t+1}|g^t)$, where the denominator represents probabilities under the true data-generating process and the numerator is a worst-case probability model.

Equality between the two representations of the common stochastic discount factor in (11) captures how agent-specific consumption growth rates must adjust to offset differences in subjective conditional probabilities. We can use these two representations of the stochastic discount factor to motivate alternative notions of the market price of risk.

2.9 Market prices of risk

Hansen and Jagannathan (1991) define the conditional market price of risk as the ratio of the conditional standard deviation of a stochastic discount factor to its conditional mean,

$$mpr_{it} = \frac{\sigma_{it}(m_{it+1})}{E_{it}(m_{it+1})}. \quad (12)$$

Here $E_{it}(\cdot)$ and $\sigma_{it}(\cdot)$ represent the conditional mean and standard deviation, respectively, evaluated with respect to consumer i 's predictive probabilities. To find unconditional prices of risk, we marginalize with respect to the growth state, using consumer i 's unconditional densities.

Because the two consumers have diverse beliefs, they also form different assessments about the price of risk. Since agent 2 knows the true transition probabilities, his perceived law of motion for aggregate consumption coincides with the actual law of motion. Hence we call mpr_{2t} the 'rational-expectations' price of risk. Similarly, because the Bayesian consumer uses his subjective predictive probabilities to make forecasts, we call mpr_{1t} the 'subjective' price of risk.

2.10 Bayesian learning and predictive probabilities

Agents observe the history of their own consumption, of aggregate consumption, and of Arrow security prices.¹¹ As in Cogley and Sargent (2008), we assume that the type 1 Bayesian consumer adopts a beta-binomial probability model for learning about F . A binomial likelihood is a natural representation for a two-state endowment process, and a beta density is the conjugate prior for a binomial likelihood. We inject additional pessimism by applying the T^2 risk-sensitivity operator of Hansen and Sargent (2007b). This operator distorts a benchmark beta prior by tilting probabilities towards the low-growth state.

¹¹They also know the sharing rules (6)-(8), and from this they can deduce the other type of agent's beliefs. Although the Bayesian consumer deduces the beliefs of the fully-informed agent, that does not mean that he recognizes they are the true probabilities. The Bayesian type 1 agent quickly discovers that the type 2 agent has a dogmatic prior. If agent 1 knew that agent 2 knew the truth, the sharing rule would reveal F , and agent 1 could also become fully informed. But we assume that from the perspective of agent 1, the beliefs of agent 2 are just someone else's prior. Therefore, they do not influence the type 1 agent's Bayesian updating. See the discussion in subsections 2.1 and 2.2.

Table 2: Prior Means for F_{hh} and F_{ll} .

	Beta		Worst-Case	
	F_{hh}	F_{ll}	F_{hh}	F_{ll}
1919 – 1933	0.915	0.805	0.869	0.966

Worst-case priors are calculated for $\alpha = 2$ and $\theta = 125$.

2.10.1 A type 1 agent's prior

As in Cogley and Sargent (2008), for the type 1 agent, we imagine a consumer who is about to emerge from the Great Contraction in 1933 and has a conventional beta prior

$$f(F_{hh}, F_{ll}) \propto f(F_{hh})f(F_{ll}), \quad (13)$$

where $f(F_{hh})$ and $f(F_{ll})$ are independent beta densities,

$$\begin{aligned} f(F_{hh}) &\propto F_{hh}^{n_0^{hh}-1}(1 - F_{hh})^{n_0^{hl}-1}, \\ f(F_{ll}) &\propto F_{ll}^{n_0^{ll}-1}(1 - F_{ll})^{n_0^{lh}-1}. \end{aligned} \quad (14)$$

The variable n_t^{ij} is a counter that records the number of transitions from state i to j through date t , and the parameters n_0^{ij} represent prior beliefs about the frequency of transitions. In the following subsections, we describe how we elicit a pessimistic prior for the type 1 agent.

2.10.2 Using a short sample

To elicit a pessimistic outlook, we calibrate $f(F_{hh}, F_{ll})$ by fitting to a short training sample covering the period 1919-1933 that oversamples contraction states. Because actual data on consumption growth are realizations of a continuous random variable, we fit a hidden Markov model to the actual data and then calibrate $f(F_{hh})$ and $f(F_{ll})$ so that they have the same mean and degrees of freedom.¹² The results are recorded in the middle column of table 2. Consumption growth was sharply negative during 1930-1933, and with a short training sample this experience would have made a Bayesian pessimistic about the onset and persistence of contractions. Thus, F_{hh} is lower and F_{ll} is higher than the estimates of Cecchetti et al. (2000).

2.10.3 Pessimistic twisting

In Cogley and Sargent (2008), we found that more pessimism is needed to attain good results for the equity premium and market price of risk. Accordingly, we multiply the benchmark beta prior by a nonnegative random variable $\zeta(F_{hh}, F_{ll}; \theta)$ that

¹²For details, see Cogley and Sargent (2008).

pessimistically distorts beliefs,

$$f_{wc}(F_{hh}, F_{ll}) \propto f(F_{hh}, F_{ll})\zeta(F_{hh}, F_{ll}; \theta). \quad (15)$$

To obtain the function $\zeta(F_{hh}, F_{ll}; \theta)$, we apply the \mathbb{T}^2 risk-sensitivity operator of Hansen and Sargent (2007b). This operator helps the consumer evaluate continuation values in a way that guards against misspecification of his prior. Application of the operator gives the indirect utility function for a problem in which the decision maker chooses a distortion to his benchmark prior $f(F_{hh}, F_{ll})$ in order to minimize the expectation of a continuation value function plus an entropy penalty. The penalty on entropy constrains the set of alternative priors against which the decision maker wants to guard, with the size of the set decreasing in a positive robustness parameter θ . The worst-case distortion to the prior is

$$\zeta(X_t, \theta) \propto \exp\left(\frac{-V(X_t)}{\theta}\right), \quad (16)$$

where $V(X_t)$ is the consumer's value function

$$V(X_t) = U(C_t) + \beta E_t[V(X_{t+1}|X_t)],$$

and the state X_t consists of statistics summarizing the observed history g^t along with the unobserved parameters F_{hh}, F_{ll} . Since we use \mathbb{T}^2 only to elicit a prior, we condition on the training sample for g_t , obtaining an initial distortion $\zeta(F_{hh}, F_{ll}; \theta)$. Notice that $\zeta(F_{hh}, F_{ll}; \theta) \rightarrow 1$ as $\theta \rightarrow +\infty$. Thus, in the limit as $\theta \rightarrow +\infty$ we recover the undistorted beta prior.

We set $\theta = 125$ to make $f_{wc}(F_{hh}, F_{ll})$ resemble one of the worst-case priors in our earlier paper. The results are recorded in the third column of table 2. Relative to the beta priors, the worst-case priors underestimate F_{hh} and exaggerate F_{ll} . Thus, the robust consumer initially believes that contractions occur more often and are longer when they do occur. Since long contractions have the character of Great Depressions, our robust consumer is initially wary of another crash. A Bayes factor actually favors the distorted prior over the benchmark beta prior, so the Bayesian consumer would not dismiss this as implausible. For further discussion of the worst-case priors, see Cogley and Sargent (2008).

Next, we approximate $f_{wc}(F_{hh}, F_{ll})$ by another product of beta densities,

$$p_{wc}(F_{hh}, F_{ll}) \propto p(F_{hh})p(F_{ll}), \quad (17)$$

where $p(F_{hh})$ and $p(F_{ll})$ are calibrated so that $p_{wc}(F_{hh}, F_{ll})$ has the same mean and degrees of freedom as $f_{wc}(F_{hh}, F_{ll})$. We do this partly for computation reasons, as it speeds our calculations quite a lot.

But there is also a substantive reason. The dependence induced by $\zeta(F_{hh}, F_{ll}; \theta)$ is an interesting feature in its own right, and in our earlier paper it contributed to

higher prices of risk. In that paper, α was calibrated at 0, and the worst-case prior induced positive correlation between F_{hh} and F_{ll} . Thus, after a transition from \bar{g}_h to \bar{g}_l , the mean of F_{hh} would increase, but so would the mean of F_{ll} . Similarly, after exiting a contraction, the mean of F_{ll} would decline but so would the mean of F_{hh} . In this way, every step in the direction of optimism was accompanied by a partially offsetting step toward pessimism. This caused pessimism to evaporate more slowly and contributed to high prices of risk.

With $\alpha = 2$, the effect seems to be different. In this case, (15) induces negative correlation between F_{hh} and F_{ll} . Thus, as F_{hh} rises, F_{ll} falls, and vice versa. Since inverse dependence makes pessimism evaporate more quickly, it causes the market price of risk to decline more rapidly. To counteract this effect, we endow the consumer with the independent prior (17). The two priors have the same mean and degrees of freedom, but (17) implies that F_{hh} and F_{ll} are updated separately.¹³

Figure 1 depicts the marginal priors for F_{hh} and F_{ll} . Solid lines portray the benchmark beta prior and dashed and dotted lines portray the two worst-case densities. Notice how the risk-sensitivity adjustment reshapes the benchmark priors by shifting probability mass toward lower values of F_{hh} and higher values of F_{ll} .

2.10.4 The posterior on F_{hh} and F_{ll}

Next, we derive an expression for agent 1's posterior, $p_1(F_{hh}, F_{ll}|C^t, C_{1t}^t, Q^t)$. Given the history of aggregate and own consumption, the history of Arrow securities prices Q^t conveys no extra information. This follows from the fact that Arrow prices are deterministic functions of the other conditioning information (see equations (9) and (10)).¹⁴ Hence, the posterior simplifies to

$$p_1(F_{hh}, F_{ll}|C^t, C_{1t}^t, Q^t) = p_1(F_{hh}, F_{ll}|C^t, C_1^t). \quad (18)$$

Similarly, given knowledge of the sharing rule and history of aggregate consumption, C_{1t} is also a deterministic function of the other conditioning information (see equation (6)). Therefore, the history of own consumption is also redundant,

$$p_1(F_{hh}, F_{ll}|C^t, C_1^t, Q^t) = p_1(F_{hh}, F_{ll}|C^t). \quad (19)$$

Hence, Bayesian updating simplifies to learning about the transition probabilities in light of observations on aggregate consumption. Appendix A describes how this is done. Among other things, the appendix establishes that the prior p_{wc} is conjugate

¹³Because the worst-case prior depends on the utility function, the value of α affects the outcome of applying the T^2 operator. When $\alpha = 0$ as in Cogley and Sargent (2008), the agent cares only about mean consumption growth and not about smoothing consumption across states. When $\alpha = 2$, the agent prefers a smooth consumption process. The worst-case transition matrix for an $\alpha = 0$ consumer therefore differs from the worst-case transition matrix for an $\alpha = 2$ consumer.

¹⁴Bayesians remember their past beliefs as well.

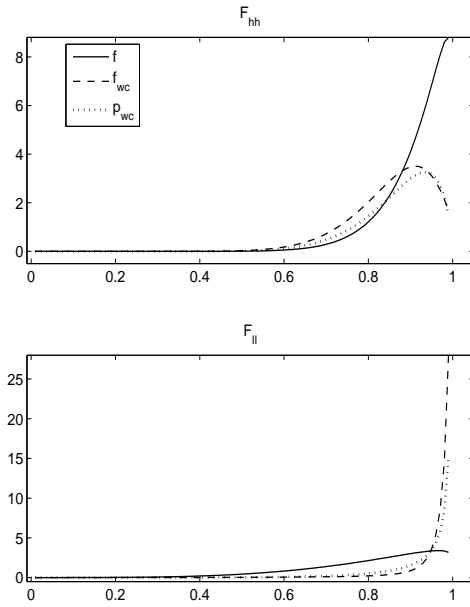


Figure 1: Beta and Robust Priors. Solid lines depict undistorted Beta priors, dashed lines portray worst-case priors for $\alpha = 2$ and $\theta_2 = 125$, and dotted lines illustrate the independent beta approximation to the worst-case prior.

to a binomial likelihood function. Hence the posterior is also a product of beta densities, and the vector of counters n_t constitute sufficient statistics. The Bayesian consumer enters each period with a prior of the form (17). After observing aggregate consumption growth, he updates the counters, incrementing by 1 the element n_{t+1}^{ij} that corresponds to the realizations of g_{t+1} and g_t :

$$\begin{aligned} n_{t+1}^{ij} &= n_t^{ij} + 1 && \text{if } g_{t+1} = \bar{g}_j \text{ and } g_t = \bar{g}_i, \\ n_{t+1}^{ij} &= n_t^{ij} && \text{otherwise.} \end{aligned} \tag{20}$$

Substituting the updated counters into (24) and normalizing delivers the new posterior, which then becomes his prior for the following period.

It is convenient to factor $pr_i(g^t)$ as

$$pr_i(g^t) = \prod_{s=1}^t pr_i(g_s | g^{s-1}), \tag{21}$$

where $pr_i(g_1 | g^0)$ is agent i 's prior distribution. To solve the Pareto problem and compute Arrow securities prices, we need the posterior predictive probabilities,

$$pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, n_t) = \iint pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, n_t, F_{hh}, F_{ll}) p(F_{hh}, F_{ll} | g_t = \bar{g}_i, n_t) dF_{hh} dF_{ll}. \tag{22}$$

In appendix A, we demonstrate that $pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, n_t) = \hat{F}_{ij}(t)$, the posterior mean of F_{ij} . Given our assumptions, the posterior mean reduces to $\hat{F}_{ij}(t) = n_t^{ij} / (n_t^{ih} + n_t^{il})$.

3 Two Benchmarks

Before studying the diverse-beliefs economy, we present results for two benchmarks, a rational-expectations economy populated only by a fully-informed consumer and a Bayesian economy in which the fully-informed consumer is absent. These benchmarks correspond to $\lambda = 0$ and $\lambda = 1$, respectively. Later we compare these polar cases to outcomes for intermediate values of λ .

3.1 A representative agent with full information

When $\lambda = 0$, the model reduces to a standard representative-agent economy in which the consumer knows the transition probabilities. Tables 3 and 4 report the Arrow security prices and market prices of risk that emerge from this model. Despite the possibility of a sharp decline in consumption, market prices of risk are quite small. The unconditional price of risk is just 0.032, an order of magnitude smaller than the lower bound of Hansen and Jagannathan (1991). The price of risk is higher in contractions than in expansions – 0.092 as opposed to 0.030 – but the contraction-state price of risk also falls well short of the Hansen and Jagannathan lower bound.

Table 3: Arrow Security Prices

	Q_{ht}	Q_{lt}
$g_t = \bar{g}_h$	0.908	0.025
$g_t = \bar{g}_l$	0.449	0.577

Table 4: Market Prices of Risk

	Expansion	Contraction	Unconditional
MPR	0.030	0.092	0.032

3.2 A representative Bayesian consumer

When $\lambda = 1$, the model reduces to a representative-agent economy in which the consumer uses Bayes' law to learn about the transition probabilities. Figures 2 and 3 summarize Arrow prices and market prices of risk averaged across 1000 sample paths.

The consumer's beliefs satisfy a Bayesian consistency theorem, so Arrow prices eventually converge to their $\lambda = 0$ full-information rational-expectations values. But this takes a long time. Along a sample path, the consumer is pessimistic about the onset of a contraction. Thus, conditional on being in an expansion, Q_h is lower and Q_l higher than their full-information values (see the first row of figure 2). The consumer is also pessimistic about the persistence of contractions. Therefore, conditional on being in a contraction, Q_h is again lower and Q_l higher than their limiting values (see the second row of the figure).

An outside observer who imputes knowledge of the transition probabilities to a representative consumer would say either that the consumer is making systematic pricing errors or that he is more risk averse than $\alpha = 2$.

Figure 3 depicts subjective and rational-expectations prices of risk. Because the consumer is mildly risk averse, subjective prices of risk are small, on the order of 0.03 to 0.09 (see the top panel of figure 3). These values are in the same ballpark as the full-information prices of risk reported above. The prices of risk needed to reconcile Arrow prices with rational expectations are substantially higher, however (see the bottom panel of figure 3). The unconditional rational-expectations price of risk is initially above 0.8 and declines gradually to 0.23 after 75 years ($1933 + 75 = 2008$). This is in the ballpark of Hansen and Jagannathan's (1991) lower bound. Going forward in time, the model predicts a further decline to around 0.15 after 200 years.

Table 5 summarizes the distribution of unconditional prices of risk in year 75. The subjective price of risk is always smaller than 0.1, but the rational expectations (RE) price of risk is greater than 0.2 on roughly half of the sample paths, and it exceeds

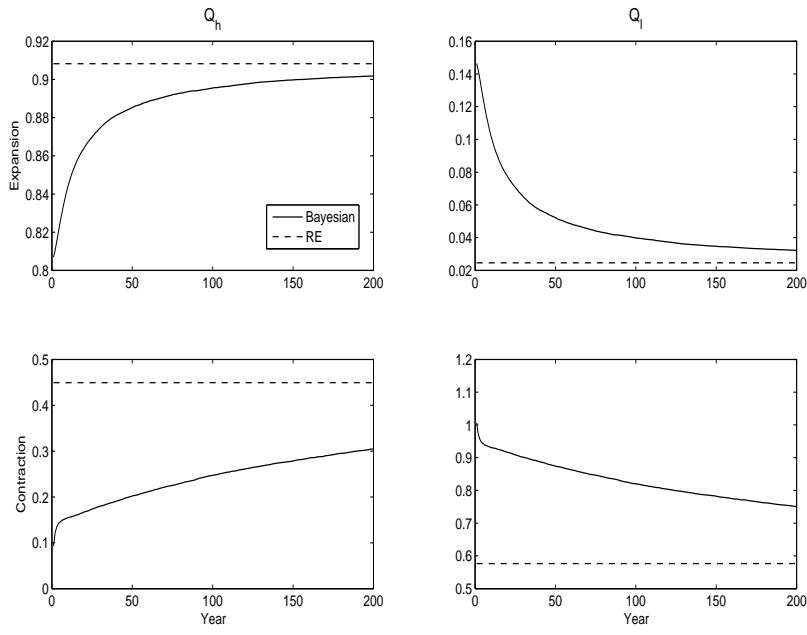


Figure 2: Arrow Security Prices. Solid lines portray Bayesian prices, and dashed lines depict full-information rational-expectations prices.

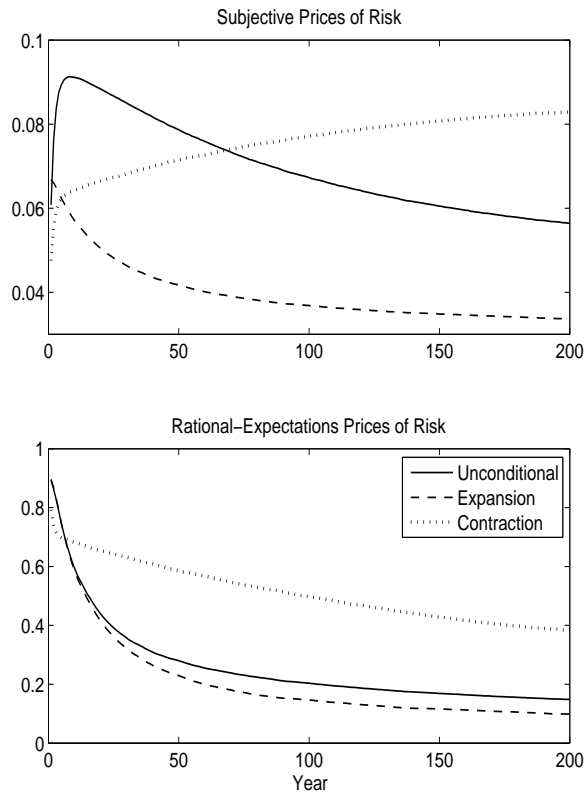


Figure 3: Market Prices of Risk in a Bayesian economy. Dashed and dotted lines depict prices of risk in expansions and contractions, respectively, while the solid line portrays the unconditional price of risk.

Table 5: Cumulative Distribution of Unconditional MPR in year 75

$pr(mpr > x)$	$x = 0.1$	0.15	0.2	0.25
Subjective MPR	0	0	0	0
RE MPR	1.0	1.0	0.495	0.279

0.25 on more than a quarter of the paths.

As in Cogley and Sargent (2008), the high rational-expectations price of risk reflects the change of measure that reconciles Bayesian asset prices with the true transition probabilities. That change of measure introduces a highly volatile learning wedge into a rational-expectations Euler equation, disconnecting the rational-expectations pricing kernel from the consumer’s subjective IMRS. The learning wedge makes the rational-expectations pricing kernel highly volatile, even though the consumer’s subjective IMRS is not. That explains why a high rational-expectations price of risk can coexist with a mild degree of risk aversion.

4 The Diverse Beliefs Economy

4.1 How consumption is distributed

Figure 4 illustrates the share of consumption for agent 2 averaged across 1000 sample paths. The respective values of λ are calibrated to deliver initial mean shares of 1, 5, 10, and 50 percent.¹⁵

On average, the well-informed agent’s consumption share increases over time. The rate of growth is higher the lower is his initial mean share. For instance, when his average share is initially 10 percent or less, his share of consumption increases by a factor of 3 or 4 over the first 100 years and by a factor of 4 to 6 over 200. But when average consumption shares are initially even, agent 2’s share increases by a factor of only 1.7 over 200 years.

Figure 5 depicts histograms for consumption shares in various years, with the share allocated to the better-informed agent on the x -axis and the proportion of sample paths on the y -axis. As time passes, the histograms shift to the right, illustrating how the better-informed agent’s consumption share increases with high probability.

¹⁵Although this is a convenient way to calibrate λ , one should keep in mind that agent 2’s initial consumption share understates his initial share of wealth. In this economy, for both types of consumer, the time 0 *financial* wealths defined in subsection 2.7 are zero. *Total* wealth, which includes the present value of the consumer’s endowment, equals the present value of future consumption, evaluated with Arrow-Debreu history-contingent prices. Since the better-informed agent’s consumption share increases over time, it follows that his initial share of wealth exceeds his initial share of consumption. Thus, for example, when λ is calibrated so that the two agents initially share consumption equally, we allocate more than half of initial aggregate wealth to the well-informed agent.




Figure 4: Average Consumption Share of the Fully-Informed Agent for Various λ .

The histograms also spread out over time and acquire long lower tails. This means that there are some sample paths on which his consumption share fails to increase and a few paths where it actually declines. Thus, although the fully-informed consumer frequently does quite well relative to the Bayesian consumer, he does not always prosper. As t increases, the histograms converge to a non-degenerate ergodic distribution,¹⁶ but that takes a long time, and we did not run the simulation long enough to learn what it looks like.

Figure 6 displays particular sample paths that show how various events alter the consumption shares. This figure refers to the simulation in which the fully-informed agent has an initial consumption share of 5 percent; the figures for initial shares of 1 and 10 percent are similar.¹⁷ Solid lines portray aggregate consumption growth, and dashed and dotted lines depict consumption shares for the Bayesian and fully-informed consumers, respectively. The message of this figure is that a consumer's share increases when his cumulative forecasting record is superior to that of his counterpart. The full set of sample paths contains a variety of experiences. The figure portrays just a few of them in order to inform intuition.

The upper left panel illustrates a path along which no contractions occur. At first, this favors the fully-informed consumer. The Bayesian consumer is initially pessimistic about the onset of a contraction and worries too much about contractions that do not occur. Eventually, their roles reverse. As good outcomes recur, \hat{F}_{hh} rises and eventually surpasses F_{hh} . At that point, expansions are more likely under the Bayesian predictive density, and his forecasting record makes a comeback relative to

¹⁶This follows from Blume and Easley (2006) and the fact that the truth lies in the support of both agent's beliefs.

¹⁷Matters are slightly different when the initial shares are 0.5.

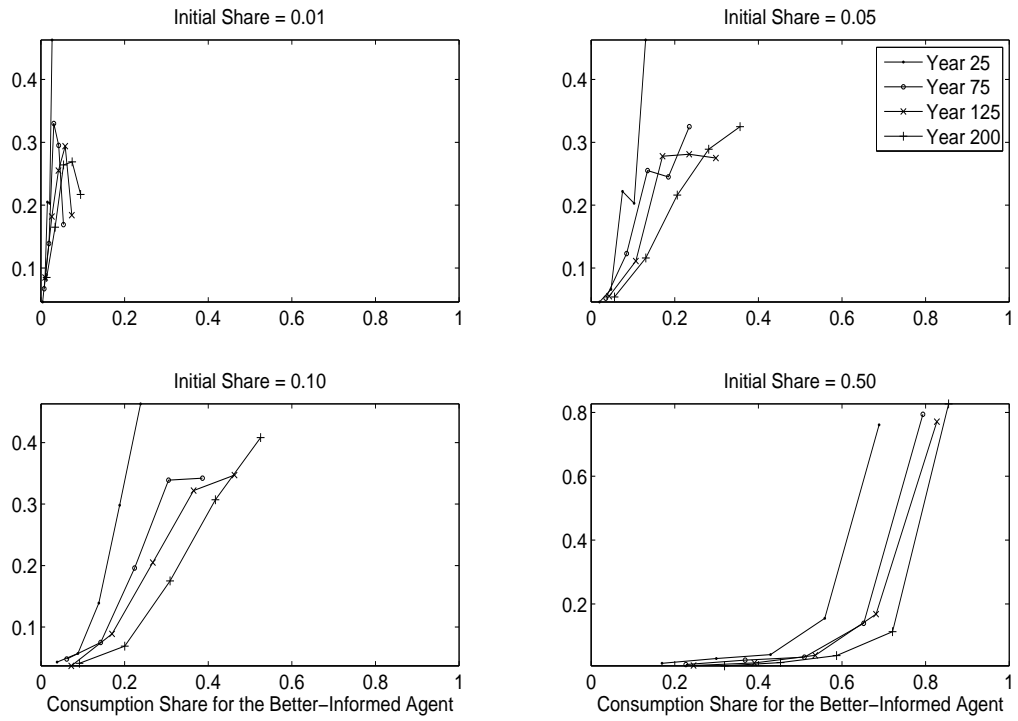


Figure 5: Histograms of Consumption Shares for the Better-Informed Agent.

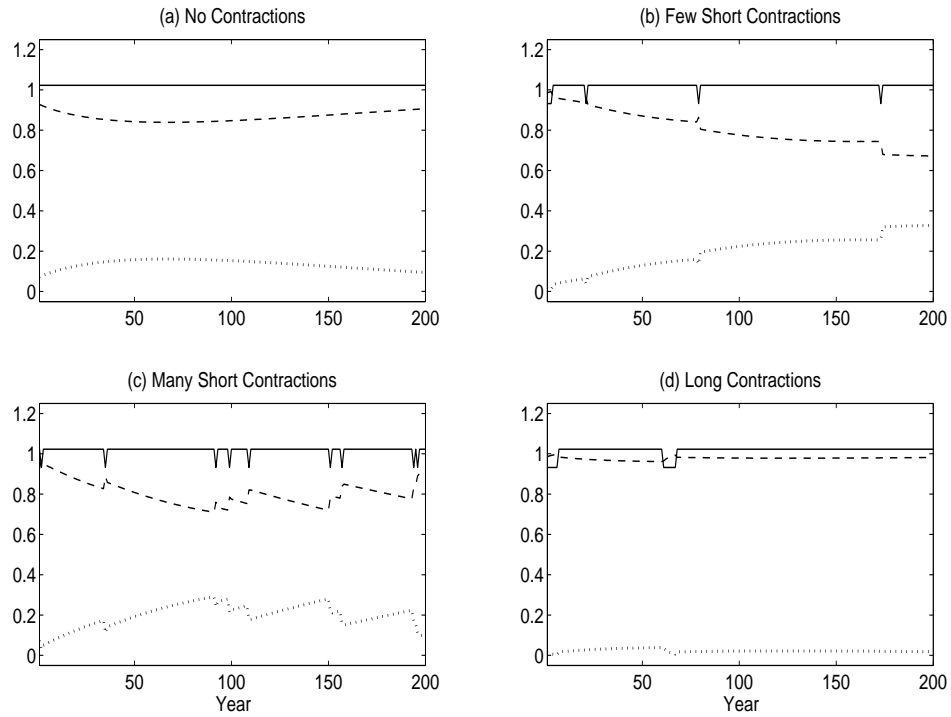


Figure 6: Particular Sample Paths, Initial Mean Share = 0.05. Solid lines portray aggregate consumption growth, and dashed and dotted lines depict consumption shares for the Bayesian and fully-informed consumers, respectively.

that of the fully-informed agent. That shifts consumption back toward agent 1.

The upper right panel depicts a sample path on which there are just a few short contractions. Because most transitions are from \bar{g}_h to \bar{g}_l , \hat{F}_{hh} converges fairly quickly to the neighborhood of F_{hh} , and since there is little disagreement about this transition probability, the onset of a contraction has only a slight effect on their consumption shares. But because no long contractions occur on this sample path, there is substantial disagreement about the persistence of contractions, the Bayesian remaining much more pessimistic. A quick end to a contraction therefore favors the fully-informed agent, improving his cumulative forecasting record relative to that of the Bayesian consumer. That shifts consumption toward the fully-informed agent.

The sample path depicted in the lower left panel has many short contractions. The occurrence of *many* contractions favors the Bayesian consumer. He is pessimistic about \hat{F}_{hl} and predicts more \bar{g}_h to \bar{g}_l transitions than the fully-informed consumer. The realization of many such transitions therefore improves his relative forecasting record, shifting consumption in his favor at the onset of a contraction. On the other hand, the fully-informed agent attaches a higher probability to high-growth states, and he rallies when the economy transits into an expansion.

Finally, the bottom right panel illustrates a sample path with a pair of long contractions. According to the fully-informed agent, these are rare events, so his relative forecasting record suffers when they are realized. Hence consumption shifts toward the Bayesian consumer when long contractions occur. Notice also how long the penalty persists. Because of the history dependence in the Pareto allocation, the share of the fully-informed agent remains low for many, many years after the contraction ends.

Relative to what happens in the limit, along the transition path the Bayesian consumer buys extra contraction insurance from the fully-informed agent. That insurance pays off in the event of frequent and/or long contractions, in which case the Bayesian's share increases at the expense of the fully-informed agent. So during frequent and/or long contractions, Bayesian consumption falls by less than aggregate consumption and that of the fully-informed agent falls by more. The payoff for the fully-informed agent is a higher consumption share in expansions. In those states, his consumption increases by more than aggregate consumption.

For initial consumption shares of 10 percent or less, the risk-sharing arrangement damps consumption volatility for the Bayesian agent and amplifies it for the fully-informed agent. Figure 7 illustrates consumption for the four sample paths shown above. In each panel, the top line portrays aggregate consumption on a log scale, and the bottom line depicts consumption for the fully-informed consumer. In the aggregate, contractions are small bumps, in some cases visually hard to discern. For the fully-informed consumer, contractions are proportionally much larger, often resulting in a very substantial decline in consumption.

Table 6 reports the standard deviation of consumption growth in the aggregate and for the two agents. Averaging across all sample paths, the standard deviation of

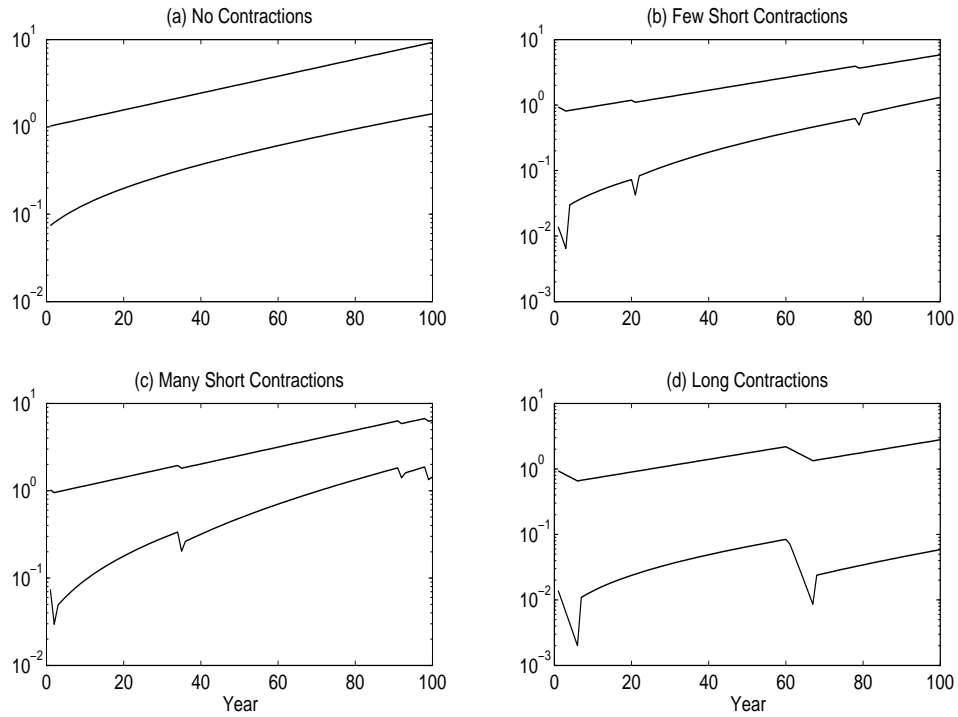


Figure 7: Particular Sample Paths, Initial Mean Share = 0.05. In each panel, the top line portrays aggregate consumption, and the bottom line depicts consumption for the fully-informed consumer.

Table 6: Standard Deviation of Consumption Growth

Initial Share	$\Delta \ln C$	$\Delta \ln C_1$	$\Delta \ln C_2$
0.01	0.0196	0.0170	0.1785
0.05	0.0196	0.0155	0.1679
0.10	0.0196	0.0177	0.1565
0.50	0.0196	0.0477	0.0911

aggregate consumption growth is 1.96 percent per annum. When agent 2 has a small initial consumption share, Bayesian consumption is a bit smoother than the aggregate. But for the fully-informed agent consumption growth is almost 10 times more volatile. Thus, the better-informed agent bears a disproportionate share of aggregate consumption risk. In effect, the risk-sharing arrangement amplifies consumption catastrophes for the fully-informed consumer relative to the aggregate data-generating process. This affects his perception of the market price of risk.

Matters are slightly different when the fully-informed trader has an initial consumption share of 0.5 (see figure 8 and the last row of table 6). The risk-sharing arrangement still amplifies consumption volatility for the well-informed consumer, but it also increases consumption volatility for the Bayesian consumer. Figure 8 portrays consumption shares for this calibration along the same sample paths shown in figure 6. Consumption covaries negatively across consumers, so there is still a lot of risk sharing. But relative to the previous calibrations, the consumption of the Bayesian consumer 1 increases a lot more in favorable states of the world and falls more in unfavorable states. This amplifies the Bayesian consumer's perception of the price of risk.

4.2 Arrow security prices

Figure 9 illustrates Arrow prices in the diverse-beliefs economies and compares them with those in polar Bayesian ($\lambda = 1$) and RE ($\lambda = 0$) economies. As before, the figures depict mean prices averaged across 1000 sample paths. The top and bottom row depict prices conditional on being in an expansion and contraction, respectively.

The presence of a fully-informed type 2 consumer moves prices toward their rational-expectations values. But when his initial consumption share is small, his effect on prices is also small. Thus, when the initial mean share is 0.01, prices in the diverse-beliefs economy are almost identical to those in the Bayesian economy. In this case, the two lines are visually hard to distinguish. The differences are a bit larger when his initial mean share is 0.05 or 0.10, but prices remain closer to Bayesian outcomes than to RE values. His influence is greater, however, when his initial mean consumption share is 0.5, and in that case convergence to the rational-expectations

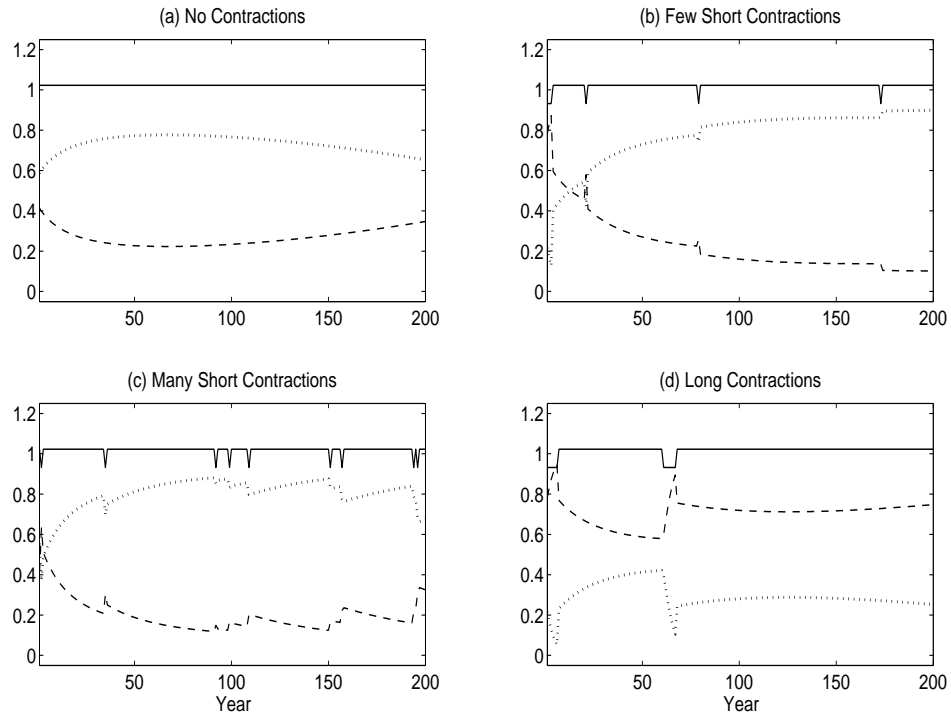


Figure 8: Particular Sample Paths, Initial Mean Share = 0.5. Solid lines portray aggregate consumption growth, and dashed and dotted lines depict consumption shares for the Bayesian and fully-informed consumers, respectively.

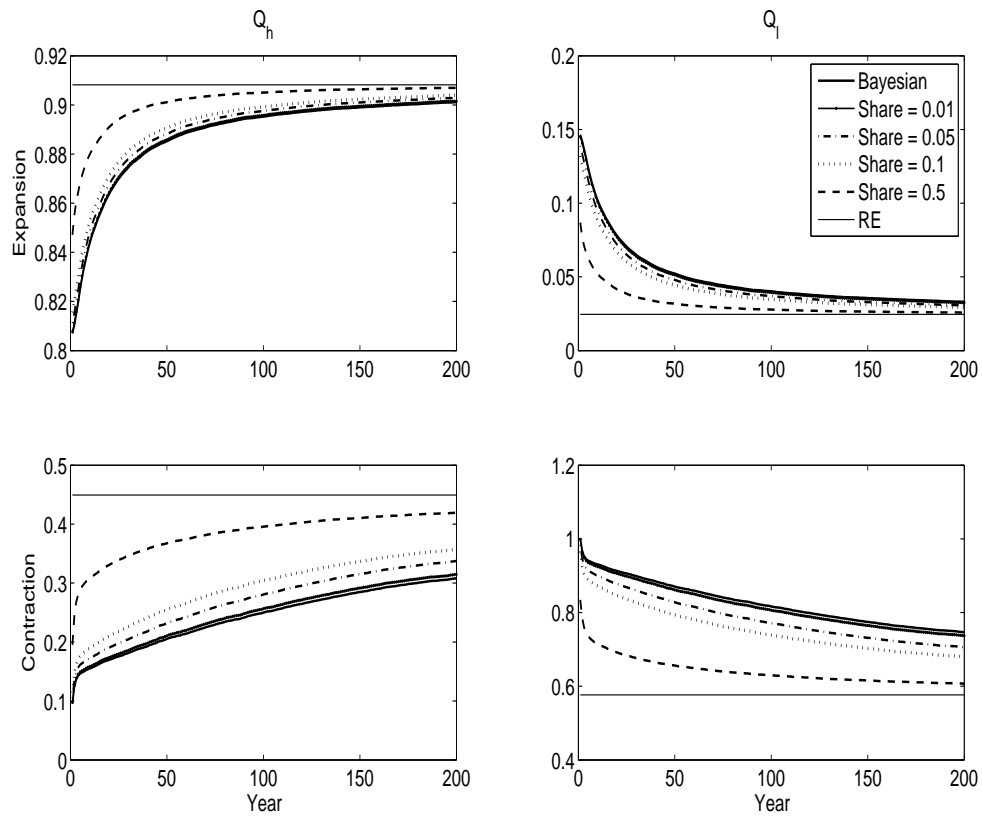


Figure 9: Arrow Security Prices. Top: conditional on expansion; bottom: conditional on contraction.

benchmark is much more rapid.

Because the fully-informed agent bears a disproportionate share of aggregate consumption risk, the price gaps shown in figure 2 no longer present such attractive opportunities to trade. For instance, consider the Arrow security paying off in the low-growth state. At first glance, figure 2 might seem to suggest that the fully-informed trader could profit by selling to the Bayesian consumer because the benchmark Bayesian price exceeds the RE price. That impression is misplaced, however, because the fully-informed trader bears much more downside consumption risk in the diverse-beliefs economy than in the representative-agent, rational-expectations economy and requires a higher selling price to compensate for the extra risk.

4.3 Market prices of risk

Figure 10 illustrates market prices in the diverse-beliefs economies and compares them with prices of risk in the purely Bayesian economy. The left-hand column portrays the subjective price of risk for the Bayesian consumer, and the right-hand column depicts rational-expectations prices of risk for the fully-informed trader.¹⁸

In all but one of the calibrations, the subjective price of risk is quite small. When the well-informed trader has a small initial consumption share, the Bayesian consumer can buy contraction insurance and therefore bears less consumption risk than in the pure Bayesian economy. Hence subjective prices of risk are lower in the mixed economy. The one exception occurs when the agents initially share consumption equally. As we have seen, in that case both agents bear more consumption risk, and that elevates the Bayesian consumer's subjective price of risk.

RE prices of risk are also smaller than in the Bayesian economy, but in most cases only slightly. For the fully-informed type 2 consumer, prices of risk remain high because he is exposed to more consumption risk than in the aggregate. For the Bayesian consumer, the price of risk is high because of the learning wedge that appears after expressing his Euler equation in terms of the true probabilities.

Table 7 summarizes the distribution of the rational expectations market price of risk in year 75. When the fully-informed agent has a small initial consumption share, the distribution is much like that in the Bayesian economy. The distribution shifts to the left as the fully-informed trader becomes more important, but high prices of risk emerge on a substantial fraction of sample paths even when his initial consumption share is 0.1. Hence, the introduction of a small measure of fully-informed consumers does not reverse the results of Cogley and Sargent (2008). A large measure of fully-informed agents is needed to overturn those results.

¹⁸In the purely Bayesian economy, the rational-expectations price of risk involves a pricing kernel which reconciles Bayesian asset prices with the true transition probabilities. In the diverse-beliefs economy, this coincides with the fully-informed consumer's price of risk.

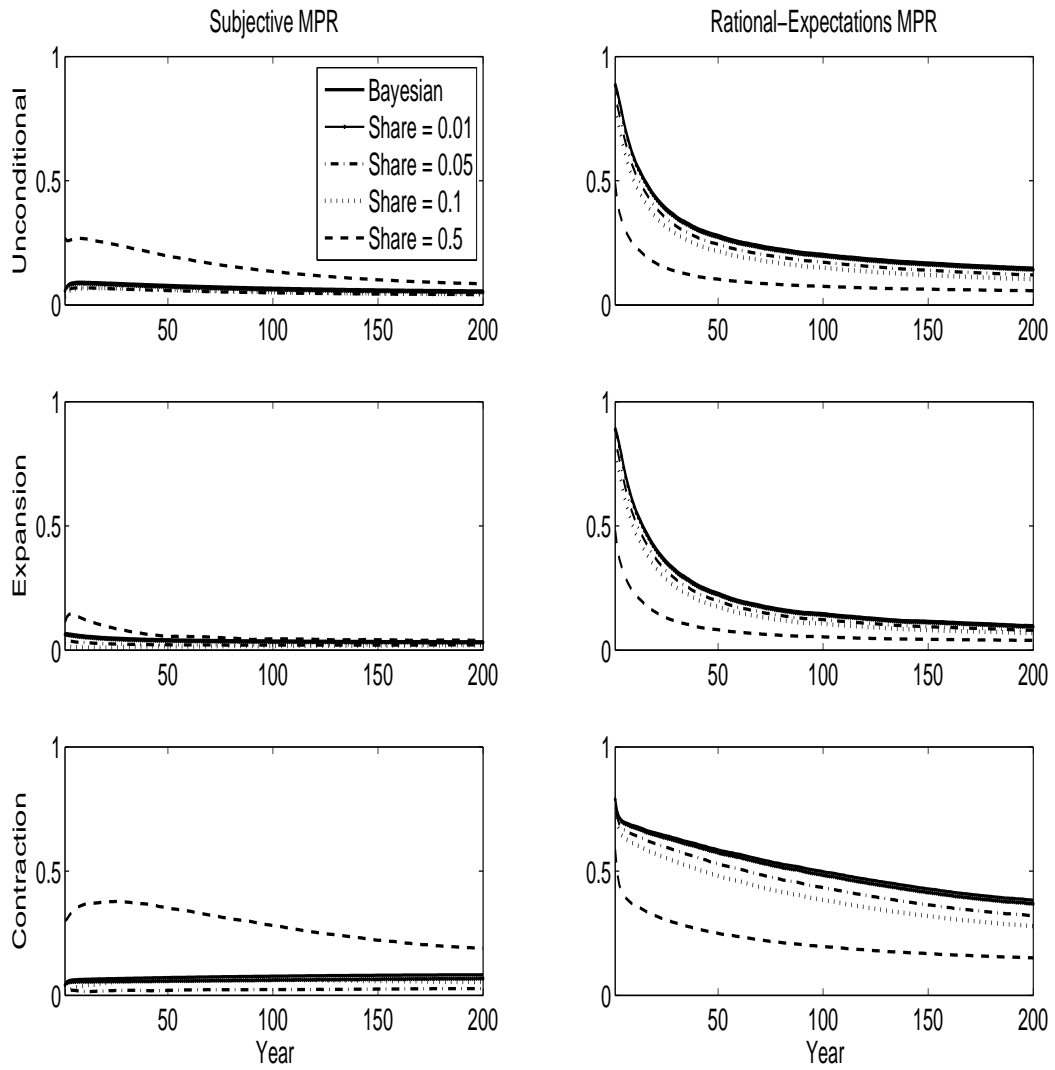


Figure 10: Market Prices of Risk. Top: unconditional; middle, conditional on expansion; bottom, conditional on contraction.

Table 7: Distribution of the Unconditional RE-MPR in Year 75

Initial Share	Pr(mpr > 0.1)	Pr(mpr > 0.15)	Pr(mpr > 0.2)	Pr(mpr > 0.25)
0	1.0	1.0	0.495	0.279
0.01	1.0	0.913	0.493	0.252
0.05	1.0	0.726	0.339	0.185
0.10	1.0	0.492	0.237	0.114
0.50	0.184	0.060	0.022	0.012

5 Concluding Remarks

Survival arguments depend sensitively on market completeness as well as the distribution of risks and agents' attitudes toward risk.¹⁹ Complete markets give the agents in this paper many opportunities to make trades that are motivated by the different subjective probabilities they put on future outcomes. Thus, by assuming complete markets, we have given the survival argument ample scope to eradicate the effects of initial pessimism on equilibrium prices. It would be interesting to compute how market incompleteness would change that.

A Updating

The likelihood function for a batch of data, $g^t = \{g_s\}_{s=1}^t$, is proportional to the product of binomial densities,

$$f(g^t | F_{hh}, F_{ll}) \propto F_{hh}^{(n_t^{hh} - n_0^{hh})} (1 - F_{hh})^{(n_t^{hl} - n_0^{hl})} F_{ll}^{(n_t^{ll} - n_0^{ll})} (1 - F_{ll})^{(n_t^{lh} - n_0^{lh})}, \quad (23)$$

where $(n_t^{ij} - n_0^{ij})$ is the number of transitions from state i to j observed in the sample.²⁰ Multiplying the likelihood by the prior (17) delivers the posterior kernel,

$$\begin{aligned} k(F_{hh}, F_{ll} | g^t) &= F_{hh}^{n_t^{hh} - 1} (1 - F_{hh})^{n_t^{hl} - 1} F_{ll}^{n_t^{ll} - 1} (1 - F_{ll})^{n_t^{lh} - 1}, \\ &\propto p(F_{hh} | g^t) p(F_{ll} | g^t), \end{aligned} \quad (24)$$

where

$$\begin{aligned} p(F_{hh} | g^t) &= \text{beta}(n_t^{hh}, n_t^{hl}), \\ p(F_{ll} | g^t) &= \text{beta}(n_t^{ll}, n_t^{lh}). \end{aligned} \quad (25)$$

Hence, the prior and likelihood form a conjugate pair. The posterior is also a product of independent beta densities, and the counters are sufficient statistics.

¹⁹This is a message of Kogan et al. (2006).

²⁰According to this notation, n_t^{ij} represents the sum of prior plus observed counters.

The posterior predictive probabilities are

$$\begin{aligned}
pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, n_t) &= \iint pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, n_t, F_{hh}, F_{ll}) p(F_{hh}, F_{ll} | g_t = \bar{g}_i, n_t) dF_{hh} dF_{ll}, \\
&= \iint pr_1(g_{t+1} = \bar{g}_j | g_t = \bar{g}_i, F_{hh}, F_{ll}) p(F_{hh}, F_{ll} | g_t = \bar{g}_i, n_t) dF_{hh} dF_{ll}, \\
&= \iint F_{ij} p(F_{hh}, F_{ll} | g_t = \bar{g}_i, n_t) dF_{hh} dF_{ll}, \\
&= \hat{F}_{ij}(t),
\end{aligned}$$

where $\hat{F}_{ij}(t)$ is the posterior mean of F_{ij} . After integrating, one can show that $\hat{F}_{ij}(t) = n_t^{ij} / (n_t^{ih} + n_t^{il})$.

References

- Beker, Pablo F. and Emilio Espino. 2008. The Dynamics of Efficient Asset Trading with Heterogeneous Beliefs. Department of Economics University of Warwick Department of Economics Universidad Torcuato Di Tella.
- Blume, Lawrence and David Easley. 2006. If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* 74 (4):929–966.
- Bossaerts, Peter. 2002. *The Paradox of Asset Pricing*. Princeton, New Jersey: Princeton University Press.
- . 2004. Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory. *Review of Economic Studies* 71:6386.
- Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark. 2000. Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True? *American Economic Review* 90 (4):787–805.
- Cogley, Timothy and Thomas J. Sargent. 2008. The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression? *Journal of Monetary Economics* 55:454–478.
- Friedman, M. and A.J. Schwartz. 1963. *A Monetary History of the United States, 1857-1960*. Princeton, New Jersey: Princeton University Press.
- Grossman, Sanford J. 1981. An Introduction to the Theory of Rational Expectations Under Asymmetric Information. *Review of Economic Studies* 48 (4):541–559.
- Hansen, Lars Peter. 2007. Beliefs, Doubts, and Learning: Valuing Macroeconomic Risk. *American Economic Review* 92:1–30.

- Hansen, Lars Peter and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99:225–262.
- Hansen, Lars Peter and Thomas J. Sargent. 2007a. Fragile beliefs and the price of model uncertainty. Unpublished.
- . 2007b. Recursive Robust Estimation and Control without Commitment. *Journal of Economic Theory* 136:1–27.
- Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield. 2006. The Price Impact and Survival of Irrational Traders. *The Journal Of Finance* LXI (1):195–228.
- Ljungqvist, Lars and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory, Second Edition*. Cambridge, Mass.: MIT Press.
- Milgrom, Paul and Nancy Stokey. 1982. *Journal of Economic Theory* 26:17–27.
- Pearlman, Joseph G. and Thomas J. Sargent. 2005. Knowing the Forecasts of Others. *Review of Economic Dynamics* 8 (2):480–497.
- Townsend, Robert M. 1983. Forecasting the Forecasts of Others. *Journal of Political Economy* 91 (4):546–88.