Heterogeneous Beliefs and State Space Dimensionality*

Rishabh Kirpalani
University of Minnesota
kirp0005@umn.edu

July 4, 2013

Abstract

This paper studies the role of state space dimensionality in determining the survival of agents with heterogeneous beliefs. I consider a model in which Bayesian agents with heterogeneous priors learn about a probability distribution over a given state space. Agents are identical in all other aspects and are indexed by these priors. If the state space is finite, the set of agents who survive for almost all sample paths is a topologically large set. The reason for this is that strong consistency results hold in this setup and most agents' beliefs converge weakly to the truth. However if the state is countably infinite, the set of agents surviving almost surely is a topologically meager set. Most agents' beliefs do not converge to the truth and all such agents are eventually driven out. In this way, the paper sheds light on the effect of increasing the degree of uncertainty in models in which agents have heterogeneous beliefs.

*I thank V. V. Chari, Larry Jones and Chris Phelan for valuable comments. All remaining errors are mine alone.

1 Introduction

What are the Pareto efficient allocations in environments where Bayesian agents learn about a probability distribution over states of the world? More specifically, agents have priors over these distributions and learn by observing public signals. Given a true distribution and agents with heterogeneous priors who are identical in every other aspect, which agents are optimally driven out? This is related to the "market selection hypothesis" as discussed by Alchian (1950) and Friedman (1966) which posits that markets favor rational agents over irrational ones. The answer to these questions depends on the assumptions we choose to impose on the following key elements:

1. Preferences and Endowments
2. Market Structure
3. State Space

There has been a fair amount of work done in the trying to understand how assumptions on the first two features affect survival chances. In a seminal paper, Blume and Easley (2006) show that with separable preferences which satisfy inada conditions, complete markets and bounded aggregate endowments, Bayesians survive for almost all distributions in the support of their prior. This paper attempts to study the role of the state space dimensionality in determining the survival of agents with different beliefs. In the process, I show that conditions which guarantee the survival of agents when the state space is finite fail to ensure the same when the state space is countably infinite.

I assume that the agents are Bayesian and have heterogeneous priors on the space of distributions. The space of priors is endowed with the weak* topology and the associated notion of weak* star convergence is crucial to the analysis. One reason for taking this route is so that I can draw upon work done in the field of Bayesian asymptotics originating with Doob (1949). It is also worth noting that this notion of convergence is weaker than the concept used by Blackwell and Dubins (1962) whose "merging of opinions" theorem has been used widely in the survival literature. As a preview of the results, first consider the case when the set of possible states of the world lives is finite. Most models[1] that attempt to answer the types of questions posed in the previous paragraph make this assumption.[2] In this case, I prove that "most" agents survive for almost all sample paths. More precisely, any agent with a prior that puts positive probabilities on all neighborhoods around the true distribution survives almost surely and the set of such priors is a topologically large set. The reason for this is that in this case, most priors are consistent with the truth and almost surely, agents learn the true distribution. This compliments the results in Blume and Easley (2006) and in particular Theorem 4 of their paper which states that a Bayesian i with prior μ^i survives for μ^i almost all distributions. Theorem 1 of this paper provides a condition that ensures survival

[1] See for example Sandroni (2000) and Blume and Easley (2006)

[2] This can often be degenerate in the sense that agents learn about a particular state which occurs with probability one

given some value of θ. Additionally, the last proposition in section 4 shows a stronger "merging of opinions" result when the convergence concept is weak*.

Next, I relax the finite state space assumption and assume that the state of the world lives in a countably infinite set. The results are striking. If agents (identified by their priors) are identical in every other aspect, the set of agents who survive for almost all sample paths is topologically small. This is in sharp contrast to the finite dimensional case where a large set of agents with reasonable priors survive. In particular, conditions which guaranteed the survival in the finite dimensional case no longer do so. The reason is that Bayesian consistency theorems are quite strong in the case where agents learn about finite dimensional objects but quite weak in the infinite dimensional case. In the former, the beliefs of any agent who has the truth in the support of her prior eventually converges to the truth in a weak sense which ultimately ensures her survival. However in the latter, most priors are inconsistent. I prove that all such agents are optimally driven out of the market since they eventually have beliefs that are singular with respect to the truth. The "merging of opinions" fails miserably for an infinite state space and and as Freedman (1965) states "..for essentially any pair of Bayesians, each think the other is crazy." To prove these results I draw upon a wealth of work done in statistical theory. There is a literature originating with Freedman (1963) and more recently with Diaconis and Freedman (1986) that asks if Bayesian consistency results hold when the state space is infinite dimensional. They find that this is not the case unless special assumptions are imposed on priors.

The two main results of the paper can be understood in the following context. First, in the finite dimensional case, the theorem can be viewed as a strengthening of the Blume and Easley (2006) result as it provides a characterization of exactly which agents survive given a true distribution. Second, in the the infinite dimensional case it reaffirms Friedman's Market Selection Hypothesis by illustrating that in complex environments almost no agents survive since they are all unable to learn the truth and in fact have beliefs that are quite different from it.

The rest of the paper proceeds as follows. After a brief literature review, section 3 presents the description of the state space and the economic environment. In section 4, I prove results for the finite dimensional case and section 5 discusses the infinite dimensional case. Section 6 presents a discussion of the above results and section 7 concludes.

2 Related Literature

This paper is related and contributes to a recent literature on the the survival of agents with heterogeneous beliefs in general equilibrium models. Blume and Easley (2006) consider an environment with separable preferences, bounded aggregate endowment, complete markets and a finite state space and study the associated planner's problem. This paper builds on their approach by studying a model with the first three features while relaxing the fourth assumption. They show that in the presence of agents with rational expectations, an agent survives only if the truth is absolutely continuous with respect to her beliefs. If the agent is a Bayesian learner, she will survive for a prior probability one set of distributions. They also show that the market survival hypothesis may fail when markets are incomplete since opportunities for trade may restricted. In an earlier related paper, Sandroni (2000) studies a decentralized environment in which markets are dynamically completed using Lucas trees. He shows that agents with inaccurate predictions are driven out and the most prosperous agents are those with accurate beliefs. More recently, Kogan et al. (2009) and Borovička (2011) consider the role of preferences is determining the survival chances for an agent with incorrect beliefs. Kogan et al. (2009) show that for a general class of separable preferences and bounded aggregate endowment, the market selection hypothesis holds true. Borovička (2011) on the other hand studies a complete markets model in which agents have recursive preferences. A remarkable consequence of this is that agents with different beliefs that are agree with the truth on probability zero events can survive in the long run. However, an important assumption he imposes on the model is that agents are dogmatic about their beliefs and so do not learn.

As mentioned in the introduction, this paper draws from a vast literature in statistical theory concerning the asymptotic consistency of Bayesian estimates. Doob (1949) proved a very general theorem which states that the posterior estimates are consistent for a set of prior probability one. However, as indicated earlier, we would like to know whether beliefs are consistent *given* some true parameter. This was first addressed in a seminal paper by Freedman (1963) whose result I state and use in section 4. Freedman (1963), Freedman (1965) and Diaconis and Freedman (1986) discuss the failure of consistency theorems when agents learn about infinite dimensional objects. Additionally, Dubins and Freedman (1964) and Ghosh and Ramamoorthi (2003) present a nice treatment of spaces of priors and prove some interesting results which I use the paper.

Models with heterogeneous beliefs have proved to be quite useful in various applied problems. Some recent examples include Scheinkman and Xiong (2003) and Albagli et al. (2011) which try and use variants of these models to explain certain asset pricing phenomena. Scheinkman and Xiong (2003) construct a model in which agents overestimate the informative content of certain signals and show how this can generate price "bubbles" and high trading volumes both of which have been empirically documented. However most of the models in this literature do not allow agents to learn from all available information. One reason for this is tractability, as learning from endogenously generated signals such as prices can be quite difficult to describe, but at least in the context of a planners problem allowing the agents to learn may result in the "merging of opinions" as proved by Blackwell and Dubins (1962).

3 Mathematical Preliminaries and Economic Environment

Suppose that the true probability distribution on the states of the world N is given by $\theta \in \Theta$ where $\Theta = \mathcal{M}(N)$ where $\mathcal{M}(N)$ is the space of probability measures on $(N, \mathcal{B}_N)$, where N is a complete metric space and $\mathcal{B}_N$ be the corresponding Borel σ-algebra on the N. While there are multiple ways to topologize $\mathcal{M}(N)$, the most convenient and natural is the weak*-topology. Central to the results in this paper is the associated notion of weak*-convergence

Definition 1 *A sequence of measures ν_n converges weak* to ν or $\nu_n \xrightarrow{*} \nu$ if*

$$\int f \, d\nu_n \to \int f \, d\nu$$

for all $f \in \mathcal{C}^b(\Theta)$, the set of all continuous and bounded functions on Θ.

A useful fact to note is that under the weak*-topology, $\mathcal{M}(N)$ is a complete separable metric space and so we can define the Borel σ-algebra $\mathcal{B}$ on $\mathcal{M}(N)$ to be the smallest σ-algebra generated by all weak*-open sets .

Let $(\mathcal{X}, \mathcal{F})$ be the signal space, where the signals are publicly observable. For each $\theta \in \Theta$, we have an induced probability measure Q_θ on X and I assume that the map $\theta \to Q_\theta$ is 1-1 and Borel. Using this, we can construct a product measure Q_θ^∞ on the space $(\mathcal{X}^\infty, \mathcal{F}^\infty)$ that makes the individual observations i.i.d. We also have the t-fold product measure Q_θ^t defined on the space $(\mathcal{X}^t, \mathcal{F}^t)$. Let $\pi(\Theta)$ be the space of priors on Θ and endow it with the weak* topology. Like Θ, $\pi(\Theta)$ is also a complete separable metric space. A prior μ induces a probability measure (a joint distribution) on the product space $\Theta \times \mathcal{X}^\infty$, P_μ defined as

$$P_\mu(A \times B) = \int_A Q_\theta^\infty(B)\mu(d\theta). \tag{1}$$

where Q_θ^∞ is the infinite product measure on $\mathcal{X}^\infty$ which makes the observations independent and A, B are Borel in Θ and $\mathcal{X}^\infty$ respectively. For H borel in $\Theta \times \mathcal{X}^\infty$,

$$P_\mu(H) = \int_\Theta (\delta_\theta \times Q_\theta^\infty)(H)\mu(d\theta)$$

where δ_θ is point mass at θ.

Definition 2 *The Topological Carrier of μ denoted by $C(\mu)$ is the smallest compact subset of $\Theta = \mathcal{M}(N)$ of μ-measure 1 such that $\lambda \in C(\mu)$ if and only if $\lambda \in \Theta$ and every Θ-neighborhood λ has positive μ-measure.*

Definition 3 *When Q_θ is dominated by a sigma finite measure ν, the posterior given a prior μ*

and observations $X_1, \ldots, X_t$ is

$$\mu_t(A) = \frac{\int_A \prod_{j=1}^{t} q_\theta(X_j)\, \mu(d\theta)}{\int_\Theta \prod_{j=1}^{t} q_\theta(X_j)\, \mu(d\theta)} \tag{2}$$

for $A \in \mathcal{B}$ where $q_\theta = \frac{dQ_\theta}{d\nu}$ is the radon-nikodym derivative.

For the main results I will assume that N is a countable set. Notice that a posterior μ_t, in turn implies a posterior estimate on N,

$$\gamma_t(i) = \int_\Theta \theta(i)\mu_t(d\theta).$$

A key notion in Bayesian asymptotics is that of consistency.

Definition 4 *A pair (θ, μ) is consistent if $\mu_t \xrightarrow{*} \delta_\theta$ for Q_θ^∞ almost all $x \in \mathcal{X}^\infty$ where δ_θ is point mass at θ.*

Finally, to understand whether sets are topologically large or small, I will use the standard concepts of first category and residual.

Definition 5 *A subset of a complete separable metric space is of first category if it is a countable union of nowhere dense sets. A residual set has a compliment which is of first category*

Next, lets turn to the economic environment. Consider an endowment economy in which the time t aggregate endowment e_t, given by a deterministic $1-1$ map from X_t to a bounded subset on $\mathbb{R}_+$ such that e_t is $\mathcal{F}_t$ measurable. As mentioned earlier, the probability distribution on N is unknown to the agents and they learn about it over time by observing public signals X_t. Each agent in a set J is endowed with a prior over Θ, $\mu \in \pi(\Theta)$ and agents are indexed by these priors. In particular, I assume that given any prior, there exist agents in J having that prior and agent j has prior μ^j. I also assume that all agents have identical preferences represented by u which is strictly increasing, strictly concave and satisfies Inada conditions. Finally I assume that markets are complete which allows us to solve for the efficient allocation using the following social planner's problem

$$\max_{\{(c_t^j)_j\}_t} \int_J \alpha^j U^j \, dj$$

subject to for all t,

$$\int c_t^j \, dj \leq e_t \tag{3}$$

where α^j are the Pareto weights and for each j,

$$U^j = E_0^j \sum_{t=0}^{\infty} \beta^t \, u(c_t^j) = \sum_{t=0}^{\infty} \int_\Theta \int_{\mathcal{X}^\infty} \beta^t \, u(c_t^j) \, Q_\theta^t(dx) \, \mu_t^j(d\theta)$$

where E_0^j denotes the expectation operator for agent j. Note that the process c_t^j is restricted to be $\mathcal{F}_t$ measurable. We say that an agent survives if her consumption is not drives to zero. Formally,

Definition 6 *Agent j survives on path $x \in \mathcal{X}^\infty$ if $\limsup c_t^j(x) > 0$. Given some fundamental θ^*, an agent θ^*-survives if he survives for $Q_{\theta^*}^\infty$ almost all x.*

4 The Finite Dimensional Case

Assume that N is a finite set. Recall that agents are indexed by their priors on Θ.

Theorem 1 *Any agent with prior μ such that the truth $\theta^* \in C(\mu)$ survives $Q_{\theta^*}^\infty$ a.s. The set of agents who θ^*-survive is residual (i.e. whose complement is a countable union of nowhere dense sets).*

The intuition for this result is that the condition $\theta^* \in C(\mu)$ ensures that the agent's prior is consistent with the truth. This was a result first proved in Freedman (1963). It is a stronger notion than the Doob (1949) (see Theorem 4) theorem which says that the prior is consistent at all parameters except for a μ measure zero set. Notice that Doob's theorem does not tell us exactly for which parameter values the prior is consistent. On the other hand, given some true distribution, Freedman's theorem allows us to identify agents whose beliefs merge weakly with the truth. The proof of the above theorem involves using the first order conditions of the planner's problem to show that these agents survive and that this set is topologically large. Before proving this theorem let's formally state the relevant result from Freedman (1963).

Theorem 2 (Freedman (1963)) *Suppose $\theta \in \Theta$ and $\{i \mid \theta(i) > 0\}$ is finite. Let μ be a prob. on $\mathcal{B}$. Then (θ, μ) is consistent iff θ is in the topological carrier of μ.*

Proof of Theorem 1. Using Theorem 2, we know that if $\theta^* \in C(\mu)$, (θ^*, μ) is consistent. Therefore for all continuous bounded functions f on Θ, $\int f\, d\mu_t \to f(\theta^*)$. Notice that for any finite t, the map $\theta \to \int \beta^t u(c_t) dQ_\theta^t$ is continuous, since u is assumed to be continuous. One can also show that the map $\theta \to \int f\, dQ_\theta^\infty$ is continuous for all continuous bounded functions on $\mathcal{X}^\infty$.[3] Therefore for all agents with $\theta^* \in C(\mu)$ and t large,

$$\int_\Theta \int_{\mathcal{X}^\infty} \beta^t\, u(c_t(x))\, dQ_\theta^t(x)\, d\mu_t(\theta) \approx \int_{\mathcal{X}^\infty} \beta^t\, u(c_t(x))\, dQ_{\theta^*}^t(x)$$

with

$$\lim_{t\to\infty} \left| \int_\Theta \int_{\mathcal{X}^\infty} \beta^t\, u(c_t(x))\, dQ_\theta^t(x)\, d\mu_t(\theta) - \int_{\mathcal{X}^\infty} \beta^t\, u(c_t(x))\, dQ_{\theta^*}^t(x) \right| = 0$$

[3] See Appendix A in Diaconis and Freedman (1986).

Let a be an agent whose prior has the above property and b be an agent with correct beliefs, i.e. $\mu^b = \delta_{\theta^*}$.[4] The first order conditions of the planning problem imply

$$\frac{u'(c_t^a(x))}{u'(c_t^b(x))} = \frac{\alpha^b}{\alpha^a} \frac{q_{\theta^*}^t(x)}{\int_\Theta q_\theta^t(x) d\mu_t^a(\theta)} \tag{4}$$

Since $\mu_t^a \xrightarrow{*} \delta_\theta$, $\int_\Theta q_\theta^t d\mu_t(\theta) \to q_{\theta^*}^t(x)$ and so,

$$\frac{u'(c_t^a)}{u'(c_t^b)} \to \frac{\alpha^a}{\alpha^b}$$

for $Q_{\theta^*}^\infty$ almost all x. Since the utility functions satisfy inada conditions, as long as the Pareto weights on the agents are positive, both agents survive $Q_{\theta^*}^\infty$ a.s.

Finally, we have to show that the set of agents surviving is residual. Consider the set A of priors containing $\mu \in \pi(\Theta)$ such that μ assigns positive mass to every nonempty open set of Θ. This set is a subset of the set B of μ such that $\theta^* \in C(\mu)$. Result 3.13 from Dubins and Freedman (1964)[5] establishes that A is residual in $\pi(\Theta)$. Therefore, since $A \subset B$, B must be residual as well.

As a final point, a valid question one might raise is whether an agent with prior μ such that $\theta^* \in C(\mu)$ can die out in finite time. This is not true when N is finite. To see this, notice that the posterior μ_t induces a distribution on N given by

$$\gamma_t(i) = \int_\Theta \theta(i)\, \mu_t(d\theta)$$

Suppose that $\gamma_t(i) = 0$ for some i and t. Then it must be that $\mu_t(\tilde{\Theta}) = 0$ where $\tilde{\Theta} = \{\theta : \theta(i) > 0\}$. But this must mean that $\mu(\tilde{\Theta}) = 0$ but since $\tilde{\Theta}$ is a neighborhood of θ^*, this contradicts the fact that $\theta^* \in C(\mu)$. ■

To understand the relationship between the above result and the "merging of opinions" theorem from Blackwell and Dubins (1962) it is useful to illustrate the fact that under the assumptions of Theorem 1, agents' predictive distributions eventually merge. To formally define this, we use a notion of weak merging from Ghosh and Ramamoorthi (2003). Given a prior μ define for any measurable set C of $\mathcal{X}^\infty$,

$$\lambda_{\mu,t}(C) \equiv \int_\Theta Q_\theta^\infty(C)\, \mu_t(d\theta) \tag{5}$$

to be the predictive distribution of $X_{t+1}, X_{t+2}, \ldots$ given $X_1, X_2, \ldots, X_t$.

Definition 7 *Given any two priors μ^a and μ^b, we say that their predictive distributions merge weakly with respect to Q_θ^∞ if*

$$\left| \int_{\mathcal{X}^\infty} f(x)\, \lambda_{\mu^a,t}(dx) - \int_{\mathcal{X}^\infty} f(x)\, \lambda_{\mu^b,t}(dx) \right| \to 0$$

[4] We know that there exists such an agent by assumption.

[5] See Theorem 5

Q_θ^∞ *almost surely, for all* $f \in \mathcal{C}^b(\mathcal{X}^\infty)$.

In other words, agents' predictive distributions merge weakly if they eventually make the same predictions about the distribution of future observations.

Proposition 1 *Fix some* $\theta_0 \in \Theta$. *Consider any two priors on* Θ, μ^a *and* μ^b *such that* $\theta \in C(\mu^i)$, $i \in \{a, b\}$. *Then* $\lambda_{a,t}$ *and* $\lambda_{b,t}$ *merge weakly as* $t \to \infty$,

Proof. As shown in Diaconis and Freedman (1986), the map $\theta \to \int f(x) dQ_\theta^\infty$ is continuous for all $f \in \mathcal{C}^b(\mathcal{X}^\infty)$. Next, from Theorem 2 we know that μ^a and μ^b are consistent at θ. Given (5), we have that

$$\int_{\mathcal{X}^\infty} f(x)\, \lambda_{\mu^i,t}(dx) = \int_\Theta \int_{\mathcal{X}^\infty} f(x) Q_\theta^\infty \mu_t^i(d\theta)$$

for $i \in \{a, b\}$. Given the continuity of the integral map, the inside integral is a continuous function of θ and the definition of weak* convergence and consistency imply that the above converges to $\int_{\mathcal{X}^\infty} f(x) dQ_{\theta_0}^\infty$ for all $i \in \{a, b\}$. This satisfies the condition stated in 7. ∎

5 The Infinite Dimensional Case

For this section assume that N is a countably infinite set. Recall that space of distributions on N is Θ and the space of priors is $\pi(\Theta)$. Define $\Theta_0 = \{\theta : \theta \in \Theta, \exists i \in I \text{ s.t. } \theta(i) = 0\}$. We know[6] that Θ_0 is a dense F_σ set of first category. Let $\Theta_1 = \Theta - \Theta_0$ and Q be a countably dense subset of Θ_1.

Theorem 3 *Fix some fundamental* $\theta^* \in \Theta_1$. *Then the set of agents who* θ^**-survive is of first category (i.e. is a countable union of nowhere dense sets).*

The first part of the proof closely follows the the argument used in Freedman (1965). The idea is to show that given some distribution θ, and *any* open subset of Θ, most priors are likely to assign probability 1 to it after observing sequences from $\mathcal{X}^\infty$ drawn according to Q_θ^∞. To prove that all agents with such priors eventually do not survive, first consider those who assign zero probability to events which the truth assigns a positive probability in finite time. Clearly with positive probability these agents do not survive. Finally for those agents whose beliefs have the property that the truth is absolutely continuous with respect to them for finite t but do not converge to truth eventually disagree as they assign zero probability to a tail event of probability 1 implying that they do not survive asymptotically. This is due to a classic result from Kakutani (1948). Before the main proof, I prove a series of useful lemmas.

Lemma 1 *Given a pair* (θ, μ), *suppose that*

$$\limsup_{t\to\infty} \int_{\mathcal{X}^\infty} \mu_t(U) dQ_\theta^\infty = 1 \tag{6}$$

simultaneously for all nonempty open subsets U *of* Θ, (θ, μ) *is not consistent*

[6]See for example Remark 1 in Freedman (1965). An F_σ set is a countable union of closed sets.

Proof. The proof of this is clear. If the above condition holds we cannot have $\mu_t \overset{*}{\to} \delta_\theta$. ■

Lemma 2 *For some $q \in Q$, let V^q be the set of $\mu \in \pi(\Theta)$ such that*

1. *μ assigns positive mass to q*
2. *The set $\{\theta : \mu(\theta) > 0\}$ is finite and $\{\theta : \mu(\theta) > 0\} - q \subset \Theta_0$*

Then V^q is dense and $\lim_{t\to\infty} \mu_t = \delta_q$ Q_θ^∞ a.s. for any $\theta \in \Theta_1$ and $\mu \in V^q$.

Proof. Let $\Theta_2 \equiv Q \cup \Theta_0$. Then since both Q, Θ_0 are dense, Θ_2 is dense as well. Then clearly for all $\mu \in V^q$, $\mu(\Theta_2) = 1$. But then it is easy to see that V^q is dense as we can approximate any $\mu \in \pi(\Theta)$ with a finite linear combination of point masses.
For the second part, consider some $\mu \in V^q$. Every other point in its support besides q has the property that it assigns zero probability to at least one $i \in N$. But then for any $\theta \in \Theta_1$ and Q_θ^∞ almost all sequences, such an i will show up leading to agent eventually assigning a posterior mass of zero to all points other than q. Since the support only contains a finite number of elements, posterior mass will shift towards q which proves the claim. ■

Lemma 3 *Define $\hat{\pi}(\Theta)$ to be the set of priors μ such that $\mu(\Theta_0) < 1$. Then $\hat{\pi}(\Theta)$ is residual in $\pi(\Theta)$.*

Proof. Using a result from Diaconis and Freedman (1986),[7] we know that the set of priors π_2 such that $\mu(\Theta_0) = 0$ for all $\mu \in \pi_2$ is residual. Since $\pi_2 \subset \hat{\pi}(\Theta)$, the result follows. ■
Define the set R^{θ^*} to be the set priors μ such that (6) holds Then a key part of proving the theorem is to show that R^{θ^*} is residual in $\pi(\Theta)$.

Lemma 4 *Fix some $\theta^* \in \Theta_1$. Then the set R^{θ^*} is residual in $\pi(\Theta)$.*

Proof. Let $\{V_k^q\}_{k=1}^\infty$ be a sequence of open subsets of Θ with $V_k^q \supset \text{ cl } V_{k+1}^q$ and $\bigcap_k V_k^q = q$. Define q_k to be a continuous function from Θ to $[0,1]$ equaling 1 on V_{k+1}^q and 0 on V_k^q. From 2 and the definition of weak* convergence we have,

$$\lim_{t\to\infty} \int_{\mathcal{X}^\infty} \int_{\Theta} q_k \, d\mu_t \, dQ_{\theta^*}^\infty = 1 \text{ for } \mu \in V^q \tag{7}$$

which implies that the map $\mu \to \int_{\mathcal{X}^\infty} \int_{\Theta} q_k \, d\mu_t \, dQ_{\theta^*}^\infty$ is continuous on $\hat{\pi}(\Theta)$. Next define the set

$$R_{qkjt}^{\theta^*} = \{\mu : \mu \in \hat{\pi}(\Theta), \int_{\mathcal{X}^\infty} \int_{\Theta} q_k \, d\mu_t \, dQ_{\theta^*}^\infty \leq 1 - j^{-1}\}.$$

R^{θ^*} is closed by the above continuity result and so $\bigcap_{t=m}^\infty R_{qkjt}^{\theta^*}$ is closed in $\hat{\pi}(\Theta)$. From 2 and (7) we also have that $\bigcap_{t=m}^\infty R_{qkjt}^{\theta^*}$ has an empty interior. Since $\pi(\Theta)$ is a complete metric space, the

[7] See Theorem 6

Baire Category theorem implies that it is a Baire space. On a Baire space, a countable union of closed sets with empty interior is empty.[8] Therefore $\bigcup_{q\in Q}\bigcup_{k=1}^{\infty}\bigcup_{j=1}^{\infty}\bigcup_{m=1}^{\infty}\bigcap_{t=m}^{\infty} R_{qkjt}^{\theta^*}$ also has empty interior. Finally, since

$$\hat{\pi}(\Theta) - R^{\theta^*} \subset \bigcup_{q\in Q}\bigcup_{k=1}^{\infty}\bigcup_{j=1}^{\infty}\bigcup_{m=1}^{\infty}\bigcap_{t=m}^{\infty} R_{qkjt}^{\theta^*},$$

$\hat{\pi}(\Theta) - R^{\theta^*}$ is of first category which proves the lemma. ■

Proof of Theorem 3. From 4 we know that the the set of agents (indexed by their priors) who eventually learn the truth (in a weak sense) is of first category. To prove the theorem we need to show that agents not learning the truth do not survive. To do this define

$$\gamma_t(i) = \int_{\Theta} \theta(i)\, \mu_t(d\theta)$$

to be the posterior induced on N by μ_t. The set of agents who have priors that are inconsistent with θ^* can be divided into two types: agents for whom $\gamma_t(i) = 0$ for some t and some i and those for whom this is not true. Note that those agents with posteriors which have the former property do not survive with positive probability since the truth assigns positive probability to each i. Consider agents for whom $\theta^* << \gamma_t^i$ for all t. Clearly for any finite t the induced product measures on $\mathcal{X}^t$, Q_γ^t and $Q_{\theta^*}^t$ are also absolutely continuous so we need consider the induced measures on $\mathcal{X}^\infty$. Let $\xi_t = \frac{dQ_{\theta^*}^t}{dQ_\gamma^t}$ be the Radon-Nikodym derivative of the measures induced on $\mathcal{X}^t$ by γ_t and θ^*. To prove the result we need to understand the relationship between the product measures Q_γ^∞ and $Q_{\theta^*}^\infty$. First notice that since $E^t\xi_t = \int \xi_t\, dQ_\gamma^t = 1$, $E^t\sqrt{\xi_t} < 1$.[9] Then it must be that

$$\prod_t E^t\sqrt{\xi_t} = 0.$$

For the above not to be true, a necessary condition is that $\lim_t E^t\sqrt{\xi_t} \to 1$ which would be true only if $\xi \to 1$ which in turn would be true only if $\gamma_t \to \theta^*$ which we have ruled out. Then using Kakutani's Dichotomy theorem (Kakutani (1948))[10] we have that $Q_{\theta^*}^\infty \perp Q_\gamma^\infty$. Therefore there exists an event to which the truth assigns probability 1 and the agent (with posterior γ_t) assigns zero mass. Therefore in the presence of an agent with true beliefs,[11] this agent's consumption is driven to zero $Q_{\theta^*}^\infty$ a.s. ■

[8]See Munkres (2000)

[9]By Jensen's inequality, $(E^t\sqrt{\xi_t})^2 < E^t(\xi_t) = 1$ which implies that $E^t\sqrt{\xi_t} < 1$.

[10]See Theorem 7

[11]We know that there exists at least one since by assumption all priors are represented and in particular one that puts point mass on the truth

6 Discussion

The results in the previous two sections warrant some discussion. First, the assumptions laid out are fairly standard in the literature. I have assumed that agents have identical preferences so that they only differ in their priors which helps isolate the effect of enlarging the state space, but this can easily be relaxed as long as all agents have separable preferences which satisfy inada conditions. The other assumption which is somewhat restrictive is that the distribution over states in i.i.d over time. This is unlike the case in Blume and Easley (2006) where agents learn about distributions over *sequences* whose finite dimensional distributions may not independent. The reason for allowing agents to learn about a distribution over a set of states rather than sequences is to isolate the role of the state space. One could allow this distribution to be time varying without any change to the main results. In particular, notice that Theorem 7 doesn't make any assumptions about independence and so the survival analysis would also hold in a more general setup.

7 Conclusion

This paper studies the role of increasing the state space on the survival of agents with heterogeneous beliefs in a general equilibrium model. In summary, the main argument illustrates how allowing the state space to be countably infinite has dramatic effects on the survival chances of agents. If all possible priors types were represented by agents who are identical in all other aspects, a topologically large set of agents survive if the state space is finite dimensional. This is because most priors are consistent and so agents eventually learn the true distribution. However, in the infinite dimensional case, the set of agents surviving almost surely is a topologically meager set. The reason for this is the inconsistency of most priors in this environment which implies that agents' beliefs will assign zero probability to a tail event which occurs with probability one.

A Appendix

Theorem 4 (Doob (1949)) *$\mu_t \xrightarrow{*} \delta_\theta$ for μ almost all θ.*

Theorem 5 (Dubins and Freedman (1964)) *The set π_1 of $\mu \in \pi(\Theta)$ assigning positive probability to all nonempty open subsets of Θ is a dense G_δ and so residual.*

Theorem 6 (Dubins and Freedman (1964)) *The set π_2 of $\mu \in \pi(\Theta)$ such that $\mu(\Theta_0) = 0$ is a dense G_δ and so residual.*

Theorem 7 (Kakutani (1948)) *Let $(\Omega_n, \mathcal{F}_n)$ be a sequence of measurable spaces and P_n and Q_n are probability measures on $(\Omega_n, \mathcal{F}_n)$ such that $Q_n << P_n$ for all n. Then either the product measure $\otimes Q_n$ is absolutely continuous with respect to $\otimes P_n$ or the two are mutually singular. The former occurs if and only if*

$$\prod_n \int \left(\frac{dQ_n}{dP_n}\right)^{1/2} dP_n > 0.$$

References

Albagli, E., C. Hellwig, and A. Tsyvinski (2011): "A Theory of Asset Pricing Based on Heterogeneous Information," Tech. rep., National Bureau of Economic Research.

Alchian, A. A. (1950): "Uncertainty, evolution, and economic theory," *The Journal of Political Economy*, 58, 211–221.

Blackwell, D. and L. Dubins (1962): "Merging of opinions with increasing information," *The Annals of Mathematical Statistics*, 33, 882–886.

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, 929–966.

Borovička, J. (2011): "Survival and long-run dynamics with heterogeneous beliefs under recursive preferences," .

Diaconis, P. and D. Freedman (1986): "On the consistency of Bayes estimates," *The Annals of Statistics*, 1–26.

Doob, J. L. (1949): "Application of the theory of martingales," *Le calcul des probabilites et ses applications*, 23–27.

Dubins, L. and D. Freedman (1964): "Measurable sets of measures." *Pacific Journal of Mathematics*, 14, 1211–1222.

Freedman, D. A. (1963): "On the asymptotic behavior of Bayes' estimates in the discrete case," *The Annals of Mathematical Statistics*, 34, 1386–1403.

——— (1965): "On the asymptotic behavior of Bayes estimates in the discrete case II," *The Annals of Mathematical Statistics*, 454–456.

Friedman, M. (1966): *Essays in positive economics*, vol. 231, University of Chicago Press.

Ghosh, J. K. and R. Ramamoorthi (2003): *Bayesian nonparametrics*, Springer.

Kakutani, S. (1948): "On equivalence of infinite product measures," *Ann. of Math*, 49, 214–224.

Kogan, L., S. Ross, J. Wang, and M. M. Westerfield (2009): "Market selection," Tech. rep., National Bureau of Economic Research.

Munkres, J. (2000): *Topology*, Prentice Hall, Incorporated.

Sandroni, A. (2000): "Do markets favor agents able to make accurate predictions?" *Econometrica*, 68, 1303–1341.

Scheinkman, J. and W. Xiong (2003): "Overconfidence, short-sale constraints, and bubbles," *Journal of Political Economy*, 111, 1183–1219.