The Dynamics of Efficient Asset Trading with Heterogeneous Beliefs

Pablo F. BEKER*

Department of Economics
University of Warwick

Emilio ESPINO†

Department of Economics
Universidad Torcuato Di Tella

This Draft: November 7, 2007

Abstract

This paper analyzes the dynamic properties of portfolios that sustain Pareto optimal allocations. We consider an infinite horizon stochastic endowment economy where the actual process of the states of nature consists in i.i.d draws from a common probability distribution. The economy is populated by many Bayesian agents with heterogeneous prior beliefs over the stochastic process of the states of nature. Since Pareto optimal allocations are typically history dependent, we propose a method to provide a complete recursive characterization when agents know the likelihood function generating the data but have different beliefs about the probability distribution of these draws. Under these assumptions, we show that if every agent's belief contains the true probability distribution of the states of nature, then investors' equilibrium asset holdings converge with probability one and, consequently, any genuine asset trading vanishes with probability one. Finally, we provide examples in which asset trading does not vanish asymptotically because either (i) no agent has the true probability distribution of the states of nature in the support of her prior belief or (ii) agents disagree on the likelihood function that generates the data.

Keywords: heterogeneous beliefs, asset trading, dynamically complete markets.

*University of Warwick, Department of Economics, Warwick, Coventry CV4 7AL, UK. E-mail: Pablo.Beker@warwick.ac.uk.

†Universidad Torcuato Di Tella, Department of Economics, Saenz Valiente 1010 (C1428BIJ), Buenos Aires, Argentina. E-mail: eespino@utdt.edu

1 Introduction

Why do investors change their portfolio positions with the arrival of new information? For a long time, the conventional wisdom was that these changes in portfolios were mainly due to risk-sharing among agents with different attitudes towards risk. Recently, Judd *et al.* [15] topped the validity of this explanation. Indeed, in the context of a stationary Markovian economy, they show that each investor's equilibrium holdings of assets of any specific maturity is constant along time and across states after an initial trading stage. They conclude "that other factors considered in the literature, such as life-cycle factors, asymmetric information, heterogeneous beliefs, and incompleteness of the asset market, play a significant role in generating trade volume." Among these factors, differences of opinions due to heterogeneous beliefs have received special attention. For instance, Morris [18] suggests that "...trading volume presumably reflects a lack of consensus in the interpretation of the (publicly released) information." (p. 247) Indeed, recent work focusing on asset trading by Scheinkman and Xiong [24] and Hong and Stein [13] has emphatically put forward this idea. Hong and Stein [13], in particular, observe that "In conventional rational asset-pricing models with common priors ... the volume of trade is approximately pinned down by the unanticipated liquidity and portfolio rebalancing needs of investors. However, these motives would seem to be far too small to account for the tens of trillions of dollars of trade observed in the real world ... the bulk of volume must come from something else - for example, differences in prior beliefs that lead traders to disagree about the value of a stock even when they have access to the same information sets"...." (p. 111-112). They argue that an appropriate explanation of trading volume is one of the highest theoretical priorities in the study of asset markets and, then, they conclude that "...taken collectively, the disagreement models...represent the best horse on which to bet" (p. 126).

In this paper we assess the widespread idea that differences of opinion due to heterogeneous beliefs can generate persistent changes in portfolios. We consider an exchange economy where investors do not know the conditional probability of the states of nature and update their priors in a Bayesian fashion.[1] Our main contributions are the following. We first show that even though heterogeneous prior beliefs can indeed generate changes in the portfolios that sustain Pareto optimal allocations,

[1] To avoid any confusion, we use of the following terminology. By a prior, we refer to the subjective unconditional probability distribution over future states of nature. In the particular case where the prior can be characterized by a vector of parameters and a probability distribution over them, we call the latter the agent´s prior belief.

these changes vanish in the long run with probability one if investors know the likelihood function generating the data and the support of their prior beliefs contains the true probability distribution of the states of nature. Additionally, we characterize the limit portfolios and show that, even though agents learn the true probability of states of nature, these portfolios need not coincide with those of an otherwise identical economy with homogeneous priors. Afterwards, we show by means of examples that if one wants to argue that heterogeneity of priors can have enduring implications on the volume of trade then one needs to assume that either (i) no investor has the truth in the support of his prior belief or (ii) investors disagree about the likelihood function generating the data.

In order to purposely disentangle the role of heterogeneous priors in explaining why investors change their portfolios from those of the other candidates listed above, we proceed as follows. First, we analyze the evolution of portfolios that support a Pareto optimal allocation to discard the lack of some market to share risk as the driving force; i.e., markets are effectively complete in our model. Second, we assume that agents interpret public data differently to abstract from disagreement stemming from asymmetries in their information. Third, we consider a population of infinitely-lived agents to shut down the life-cycle factors motive. Finally, we assume that both the endowments as well as the assets returns are i.i.d. draws from a common probability distribution to isolate from the role of non-stationarities in fundamentals.

Our approach hinges on studying portfolios that support Pareto Optimal allocations. But solving directly for the portfolios is not always possible and, therefore, we follow an indirect strategy developed by Espino and Hintermaier [8]. We begin with a recursive characterization of the set of Pareto optimal allocations. The optimal plan for the planner's problem is history dependent whenever agents have heterogeneous priors. This is because optimality requires the ratio of marginal valuations of consumption of any two agents -which includes priors that could be subjectively held- to be constant along time. Consequently, at any date the ratio of marginal utilities at any future event must be proportional to the history dependent ratio of the agents' priors about that event, i.e. the likelihood ratio of the agents' priors. Since history dependence makes standard recursive methods unsuitable, we tackle this difficulty using a strategy similar to Lucas and Stokey's [17]. We obtain a recursive characterization of the set of Pareto optimal allocations in our stochastic framework under the assumption that investors know the likelihood function generating the data but have different prior beliefs about the probability of the states of nature.[2] The key

[2] Lucas and Stokey [17] characterize recursively optimal programs in a deterministic setting where

insight is that the planner does not need to know the partial history itself in order to continue the date zero optimal plan from date t onwards. In fact, it suffices that he knows the likelihood ratio of the agents' priors, the state of nature and the agents' prior beliefs over the probability of the states of nature. We argue that the sequential formulation of the planner's problem is equivalent to a recursive dynamic program where the planner allocates current feasible consumption and assigns next period attainable utility levels among agents.

Afterwards, we use the planner's policy functions to characterize recursively investors' financial wealth in any dynamically complete market equilibrium. This allows us to establish that the financial wealth distribution (and the corresponding portfolios) converges if and only if both the likelihood ratio as well as the investors' posterior beliefs over the unknown parameters converge.

When agents know the true likelihood function, the well-known consistency property of Bayesian learning implies that the agents' prior beliefs converge with probability one. To get a thorough understanding of the limiting behavior of portfolios, therefore, what remains to explain is the asymptotic behavior of likelihood ratios. When the support of the agents' prior beliefs over the parameters is a countable set containing the true probability distribution, the true probability distribution over paths is absolutely continuous with respect to the agents' priors and, therefore, the convergence of likelihood ratios follows from the well-known result in Blackwell and Dubins [1]. When the agents' prior beliefs have a positive and continuous density with support containing the true parameter, the hypothesis in Blackwell and Dubins [1] are not satisfied and so we apply a result in Phillips and Ploberger [21] to show that still the likelihood ratio of the agents' priors converges with probability one. We also show that equilibrium portfolios converge to those of a rational expectations equilibrium of an economy where the investors' relative wealth is determined by the densities of their prior beliefs evaluated at the true parameter. The important message here is that when investors are Bayesians who know the likelihood function generating the data and have the truth in the support of their prior beliefs, the heterogeneity of priors by itself can generate changes in portfolios but these changes necessarily vanish.

Later, we give two examples where agents are Bayesians but change portfolios infinitely often. In the first example, agents know the likelihood function generating the data but they do not have the truth in the support of their prior beliefs. For

recursive preferences induce the dependence upon histories.

simplicity we consider agents that are dogmatic in the sense that the support of their prior beliefs consists in only one point. We assume that their (degenerate) prior beliefs are such that the associated one-period-ahead conditional probabilities have identical entropy, a condition that ensures that the likelihood ratio of their priors fluctuates infinitely often between zero and infinity and, consequently, portfolios fluctuate infinitely often. The second example underscores the importance of assuming that every agent knows the true likelihood function for the portfolio to converge. To stretch the argument to the limit, we consider an example in which only one agent does not know the true likelihood function. This agent makes exact one-period-ahead forecasts infinitely often but it also makes mistakes infinitely often though rarely. We show that the likelihood ratio of these agents' priors fails to converge with probability one implying that the set of paths where the equilibrium portfolio converges has probability zero.

This paper relates to two branches of the literature on asset markets: models aiming to explain the dynamic consequences of belief heterogeneity on investors' behavior and models analyzing the market selection hypothesis. Harrison and Kreps [12] and Harris and Raviv [11] are the leading articles of the first branch and inspired subsequent work by Morris [19] and Kandel and Pearson [14], respectively. Those first-generation papers consider partial equilibrium models where a finite number of risk-neutral investors trade one unit of a risky asset subject to short-sale constraints. Investors do not know the value of some payoff relevant parameter but they observe a public signal and have heterogeneous, but degenerate, prior beliefs about the relationship between the signal and the unknown parameter. Belief heterogeneity implies that they value the asset differently in spite of having the same information. Since each investor is absolutely convinced their model is the correct one, their disagreement does not vanish as the data unfold.

Harrison and Kreps [12] consider the case where agents have different prior beliefs over the probability distribution of next period dividends and focuses on its asset pricing implications. They show that speculative behavior might arise, in the sense that the asset price might be strictly greater than every trader's fundamental valuation. This occurs, they argue, whenever the trader who holds the asset anticipates she will able to resell it in the future for strictly more than her short-term valuation. Harris and Raviv [11], on the other hand, concentrate on the relationship between trade volume and asset prices. Agents agree about the probability distribution of dividends but disagree on the likelihood of the signals. Risk neutrality ensures that the group with the higher valuation holds the asset and no further trade occurs as long as that group remains the one who values it the most. Trade occurs only when the two groups "switch sides."

The possibility that agents learn is addressed by Morris [19] who considers Harrison and Kreps' [12] model but assumes investors have non-degenerate prior beliefs about the probability distribution of dividends and characterizes the set of prior beliefs for which a speculative premium actually exists. He assumes the true process is i.i.d. and investors know the true likelihood function. Since they are Bayesian, they eventually learn the truth. Consequently, risk neutrality implies the price converges and the speculative premium vanishes in the long run. We underscore that even though in Morris [18] the speculative premium vanishes, asset trading does not because there is always a period in the future when the asset changes hands once again. His asymptotic results, however, are a direct consequence of the assumption that agents are risk-neutral. Indeed, under risk-neutrality the intertemporal marginal rates of substitution are independent of the equilibrium allocation and, therefore, they are linear in the agents' one-period-ahead conditional probabilities. This has two direct implications. On the one hand, when the individuals' one-period-ahead conditional probabilities *switch sides* perpetually, so do their intertemporal marginal rates of substitution and, therefore, new incentives for a change in the ownership of the asset arise infinitely often. On other hand, asset prices themselves are parameterized by the one-period-ahead conditional probabilities and, thus, they converge together. We argue that these forces do not operate in a setting where agents are risk-averse and allocations are Pareto Optimal. Indeed, the persistent switching in intertemporal rates of substitution in Morris [18] is not robust to the introduction of risk-aversion since in that case the agents' intertemporal marginal rates of substitution are always equalized in any efficient allocation. Portfolio changes might still occur persistently but this depends purely on the asymptotic behavior of the efficient allocation. Furthermore, as we emphasized above, the convergence of the one-period-ahead conditional probabilities by itself does not guarantee the convergence of allocations, asset prices and portfolios.

The aforementioned work assumes a capital market imperfection (i.e., short-sale constraints) to argue that belief heterogeneity can have a fundamental effect on asset prices and the volume of trade. But this source of heterogeneity may matter even if they do not give rise to a speculative premium and even in the absence of any market imperfection. As a notable exception, Cogley and Sargent [5] focus on the effect on asset prices due solely to prior belief heterogeneity under the assumption that agents know the true likelihood function. They consider a Lucas [16] tree model with a risk-neutral representative agent with a pessimistic but non-degenerate prior belief over the growth rate of dividends. Even though learning eventually erases pessimism, pessimism contributes a volatile multiplicative component to the stochastic discount factor that an econometrician assuming correct priors would attribute to implausible degrees of risk aversion.[3] Thus, their work is close in spirit to ours in that they use a general equilibrium model without any additional market imperfection. Since they study a representative agent framework, however, they are silent about the implications for trading volume.

This literature has not disentangled yet the asset trading implications stemming purely from differences in priors and, more importantly, it is still an open question what the limiting behavior of asset trading is when agents eventually learn the true one-period-ahead conditional probability.

The second branch of the literature related to our paper analyses the market selection hypothesis and is exemplified by the work of Sandroni [22] and Blume and Easley [3]. Sandroni [22] shows that, controlling for discount factors, if some trader's prior merge with the true distribution then she survives and any other trader survives if and only if her prior merges with the true distribution as well.[4] He also considers some cases in which no agent's prior merges with the truth. He shows that the entropy of priors determines survival and, therefore, an agent who persistently makes wrong predictions vanishes in the presence of a learner. To see the scope of Sandroni's results, recall that an agent's prior merges with the true distribution if and only if the true distribution is absolutely continuous with respect to that agent's prior. This is a strong restriction on priors that is not satisfied, for instance, if the true process is i.i.d., the agent knows this fact but her prior beliefs over the probability of the states of nature have continuous and positive density. In that case, since the entropy of every agent's prior is the same, one cannot apply Sandroni's results relating survival with the entropy of priors either. This is precisely the case that Blume and Easley [3] consider. In a setting similar to ours, they show that the evolution of the agents' consumption in any efficient allocation depends only on the discount rates and on the likelihood ratio of their priors. They prove that among Bayesian learners who know the true likelihood function generating the data, have prior beliefs over the parameter with positive and continuous density on a set containing the truth, only those with the lowest dimensional support can have positive consumption in the long

[3] Their model can generate substantial and declining values for the market prices of risk and the equity premium and, additionally, can predict high and declining Sharpe ratios and forecastable excess stock returns.

[4] An agent is said to survive if her consumption does not converge to zero.

run. Technically speaking, Blume and Easley's notion of convergence is *in probability* and they establish their asymptotic result for *almost all parameters* in the support of the agent's prior belief. Although we do not focus on survival, one side contribution of this paper is to make Blume and Easley's results more robust because we show that every Bayesian agent with a prior belief with the lowest dimensional support actually survives *with probability one* (not just in probability), not only for *almost every parameter* in the support of her prior belief but actually *for all parameters* in the support of her prior belief.[5] Our treatment of priors is very general in that we consider a family that includes priors for which the one-period-ahead conditional probability converges to the truth regardless of whether the agents' priors merge with the truth or whether traders know the true likelihood function. In addition, it includes cases in which the entropy of all agents is the same but some agents do not learn the true one-period-ahead conditional probability.

Our results on the dynamics of portfolios that support a Pareto Optimal allocation are a novel contribution to the literature because neither Sandroni [22] nor Blume and Easley [3] analyze portfolio dynamics. However, one might hastily conjecture that their results on the asymptotic behavior of consumption when agents have different priors would map easily into properties on the asymptotic behavior of the supporting portfolios. On the contrary, this mapping can actually be rather intricate. To grasp the difficulty, consider the simplest case in which investors have homogeneous priors. In that case it has been known for a long time that Pareto optimal individual consumption in a stationary Markovian economy is a time homogeneous process with finite support (see section 20 in Duffie [6] for example). Nonetheless, it has been surprisingly difficult to establish how these properties translate into properties of the portfolio in a dynamically complete markets equilibrium (see Judd et al. [15]).

This paper is organized as follows. In Section 2 we describe the model. In section 3 we present a simple example that illustrate the main ideas in this paper. The recursive characterization of Pareto optimal allocations is in section 4. Section 5 characterizes the asymptotic behavior of the agents present discounted value of excess demand. Finally, sections 6 and 7 discuss when the agents' portfolio converge and when it does not. Conclusions are in section 8. Proofs are gathered in the Appendix.

[5] This distinction is economically relevant because both in Blume and Easley's [3] setting as well as in ours the data (and agents' ultimate fate) may be produced by a probability measure with parameters that may lie in a zero measure set of the agents' support.

2 The Model

We consider an infinite horizon pure exchange economy with one good. In this section we establish the basic notation and describe the main assumptions.

2.1 The Environment

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The set of possible states of nature at date $t \geq 1$ is $S_t \equiv \{1, ..., K\}$. The state of nature at date zero is known and denoted by $s_0 \in \{1, ..., K\}$. We define the set of partial histories up to date t as $S^t = \{s_0\} \times \left(\times_{k=1}^{t} S_k\right)$ with typical element $s^t = (s_0, ..., s_t)$. $S^\infty \equiv \{s_0\} \times (\times_{k=1}^{\infty} S_k)$ is the set of infinite sequences of the state of nature and $s = (s_0, s_1, s_2, \cdots)$, called a path, is a typical element.

For every partial history s^t, $t \geq 0$, a *cylinder* with base on s^t is the set $C(s^t) \equiv \{s \in S^\infty : s = (s^t, s_{t+1}, \cdots)\}$ of all paths whose t initial elements coincide with s^t. Let $\mathcal{F}_t$ be the σ-algebra that consists of all finite unions of the sets $C(s^t)$. The σ-algebras $\mathcal{F}_t$ define a filtration on S^∞ denoted $\{\mathcal{F}_t\}_{t=0}^{\infty}$ where $\mathcal{F}_0 \subset \ldots \subset \mathcal{F}_t \subset \ldots \subset \mathcal{F}$ where $\mathcal{F}_0 \equiv \{\emptyset, S^\infty\}$ is the trivial $\sigma-$algebra and $\mathcal{F}$ is the σ-algebra generated by the algebra $\bigcup_{t=0}^{\infty} \mathcal{F}_t$.

For any probability measure $\Pi : \mathcal{F} \to [0,1]$ on $(S^\infty, \mathcal{F})$, $\Pi_{s^t} : \mathcal{F} \to [0,1]$ denotes its posterior distribution after observing s^t.[6] Let $\Pi_t(s)$ be the probability of the finite history s^t, i.e. the $\mathcal{F}_t-$measurable function defined by $\Pi_t(s) \equiv \Pi(C(s^t))$ and $\Pi_0 \equiv 1$. Let $\pi\left(s_t \,\middle|\, s^{t-1}\right) \equiv \frac{\Pi_t(s)}{\Pi_{t-1}(s)}$ denote the one-period-ahead conditional probability of state s_t. Finally, for any random variable $x : S^\infty \to \Re$, $E^{\Pi}(x)$ denotes its mathematical expectation with respect to Π.

Let Δ^{K-1} be the $K-1$ dimensional unit simplex in $\Re^K$, $\mathcal{B}\left(\Delta^{K-1}\right)$ be its Borel sets and $\mathcal{P}(\Delta^{K-1})$ be the set of probability measures on $\left(\Delta^{K-1}, \mathcal{B}\left(\Delta^{K-1}\right)\right)$. Consider a set of probability measures on $(S^\infty, \mathcal{F})$ parameterized by $\theta \in \Delta^{K-1}$, with typical element Π^θ, with the additional property that the mapping $\theta \mapsto \Pi^\theta(B)$ is $\mathcal{B}\left(\Delta^{K-1}\right)-$measurable for each $B \in \mathcal{F}$. This set includes the subset of probability measures on $(S^\infty, \mathcal{F})$ uniquely induced by i.i.d. draws from a common distribution $\theta : 2^K \to [0,1]$, where $\theta(\xi) > 0$ for all $\xi \in \{1, ..., K\}$, with typical element P^θ. We make the following assumption.

A.0 The true stochastic process of states of nature is P^{θ^*} for some θ^*.

[6] Formally, $\Pi_{s^t}(A) \equiv \frac{\Pi\left(A_{s^t}\right)}{\Pi(C(s^t))}$ for every $A \in \mathcal{F}$, where $A_{s^t} \equiv \left\{s \in S^\infty : s = \left(s^t, s'\right),\ s' \in A\right\}$.

2.2 The Economy

There is a single perishable consumption good every period. The economy is populated by I (types of) infinitely-lived agents where $i \in \mathcal{I} = \{1, ..., I\}$ denotes an agent's name. A consumption plan is a sequence of functions $\{c_t\}_{t=0}^{\infty}$ such that $c_t : S^{\infty} \rightarrow \mathbb{R}_+$ is $\mathcal{F}_t-$measurable for all t and $\sup_{(t,s)} c_t(s) < \infty$. The agent's consumption set, denoted by $\mathbf{C}$, is the set of all consumption plans.

2.2.1 Preferences

We assume that agents' preferences satisfy Savage's [23] axioms and, therefore, they have a subjective expected utility representation. This representation provides a prior P_i over paths and, as it is well-known, it also implies that agents are Bayesians (i.e., they update their prior using Bayes' rule as information arrives). But, most importantly, it does not otherwise restrict agent's priors in any particular way.[7]

We denote by P_i the probability measure on $(S^{\infty}, \mathcal{F})$ representing agent i's prior and we make the standard assumptions that the utility function is time separable and the discount factor is the same for all agents. That is, her preferences are represented by

$$U_i^{P_i}(c_i) = E^{P_i}\left(\sum_{t=0}^{\infty} \beta^t \ u_i(c_{i,t})\right),$$

where $\beta \in (0, 1)$ and $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuously differentiable, strictly increasing, strictly concave and $\lim_{x \rightarrow 0} u_i'(x) = +\infty$ for all i.

One particular family of priors is that where the agent believes that the true process of states of nature belongs to a parametric family of probability measures, $\{\Pi^{\theta}\}$, but the agent does not know the parameter $\theta \in \Delta^{K-1}$. That is,

$$\Pi(A) = \int_{\Delta^{K-1}} \Pi^{\theta}(A)\, \mu(d\theta) \qquad \text{for every } A \in \mathcal{F}, \tag{1}$$

where $\mu \in \mathcal{P}(\Delta^{K-1})$ is the prior belief over the unknown parameters. The hypothesis of rationality can be further strengthened to require that the agent is a Bayesian who knows that the true process generating the data is i.i.d. but does not know the true probability of the states of nature. Accordingly, we say that an agent with prior Π knows the likelihood function (of the stochastic process) generating the data if $A.1$ holds.[8]

A.1 $\Pi^{\theta} = P^{\theta}$ for every $\theta \in \Delta^{K-1}$.

[7] See Blume and Easley [2] for a complete discussion on the implications of Savage's axioms.

[8] The celebrated De Finetti theorem states that this is equivalent to the prior being exchangeable.

We want to emphasize that even though under A.1 agents agree that the states of nature are generated by i.i.d. draws from a common distribution θ, they might still disagree about θ itself. The following assumption imposes more structure on the subjective distribution of θ and it will be discussed further below.

A.2 μ has density f with respect to Lebesgue that is continuous at θ^* with $f(\theta^*) > 0$.

Another interesting specification of prior beliefs is a point mass probability measure on θ defined as $\delta_\theta : \mathcal{F} \to [0,1]$ where

$$\delta_\theta(B) \equiv \begin{cases} 1 & \text{if } \theta \in B \\ 0 & \text{otherwise.} \end{cases}$$

When priors belong to the class represented by (1), Bayes' rule implies that prior beliefs evolve according to

$$\mu_{i,s^t}(d\theta) = \frac{\pi^\theta(s_t \left| s^{t-1}\right.)\ \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}} \pi^\theta(s_t \left| s^{t-1}\right.)\ \mu_{i,s^{t-1}}(d\theta)}, \tag{2}$$

where $\mu_{i,0} \in \mathcal{P}(\Delta^{K-1})$ is given at date 0 and $\pi^\theta(s_t \left| s^{t-1}\right.) \equiv \frac{\Pi^\theta(C(s^t))}{\Pi^\theta(C(s^{t-1}))}$. Observe that under A.1, $\pi^\theta(s_t \left| s^{t-1}\right.) = \theta(s_t)$.

Lemma 1 *Suppose agent i's prior satisfies (1). Then,*

$$P_{i,s^t}(B) = \int_{\Delta^{K-1}} \Pi^\theta_{s^t}(B) \frac{\pi^\theta(s_t \left| s^{t-1}\right.)\ \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}} \pi^\theta(s_t \left| s^{t-1}\right.)\ \mu_{i,s^{t-1}}(d\theta)}. \tag{3}$$

It is well-known that Bayesian learning is consistent for any prior satisfying A.1. However, this property applies to more general specifications of priors (for instance, those satisfying (1), see Schwartz [26, Theorems 3.2 and 3.3]), and since our example 4 in Section 7.2 does not satisfy A.1 but it does satisfy (1), we state the consistency result in the following Lemma to make precise its scope.

Lemma 2 *Suppose that for $\mu_{i,0}-$ almost all $\theta \in \Delta^{K-1}$ the probability measures Π^θ on $(S^\infty, \mathcal{F})$ are mutually singular. Then $\left\{\mu_{i,s^t}\right\}_{t=0}^\infty$ converges weakly to $\mu^\infty = \delta_\theta$ for $\Pi^\theta-$almost all $s \in S^\infty$, for $\mu_{i,0}-$ almost all $\theta \in \Delta^{K-1}$.*

2.2.2 Endowments

Agent i's endowment at date t is $y_{i,t}(s) = y_i(s_t)$ for all s and the aggregate endowment is $y(s_t) = y_t(s) \equiv \sum_{i=1}^I y_{i,t}(s) \leq \overline{y} < \infty$. An allocation $\{c_i\}_{i=1}^I \in \mathbf{C}^I$ is *feasible* if $c_i \in \mathbf{C}$ for all i and $\sum_{i=1}^I c_{i,t}(s) \leq y_t(s)$ for all $s \in S^\infty$. Let Y^∞ denote the set of feasible allocations.

3 Heterogeneous Priors and Portfolios: Examples

The main purpose of this section is to illustrate our main results using simple examples of dynamically complete markets equilibria.

Suppose there are two states, $A.0$ holds with $\theta^*(1) = \frac{1}{2}$, two agents, $u(c) = \ln c$ and $y_i(\xi) = \lambda_i y(\xi) > 0$ for all $\xi \in \{1,2\}$ where $\lambda_1 + \lambda_2 = 1$. Agents can trade a full set of Arrow securities. Arrow security ξ' pays 1 unit of the consumption good if $s_{t+1} = \xi'$ and 0 otherwise. The price of Arrow security $\xi' \in \{1,2\}$ and agent i's holdings at date t on path s are denoted by $m_t^{\xi'}(s)$ and $a_{i,t}^{\xi'}(s)$, respectively.

In Appendix A we show that equilibrium consumption and portfolios are

$$\begin{aligned} c_{i,t}(s) &= \left(\lambda_i + \lambda_j \tfrac{P_{j,t}(s)}{P_{i,t}(s)}\right)^{-1} \lambda_i \; y_t(s), \\ a_{i,t}^{\xi'}(s) &= \frac{1}{1-\beta}\, y(\xi')\; \lambda_i \left(\left(\lambda_i + \lambda_j \tfrac{P_{j,t}(s)}{P_{i,t}(s)} \tfrac{p_j(\xi' | s^t)}{p_i(\xi' | s^t)}\right)^{-1} - 1\right), \quad \xi' \in \{1,2\}. \end{aligned} \tag{4}$$

Observe that the evolution of individual portfolios is completely determined by the evolution of the likelihood ratio, $\frac{P_{j,t}(s)}{P_{i,t}(s)}$, and the ratio of the one-period-ahead conditional probabilities, $\frac{p_j(\xi' | s^t)}{p_i(\xi' | s^t)}$. Portfolios converge if and only if the product of these two ratios converge. Thus, trading is purely determined by the heterogeneity of priors.

The relevant margin of heterogeneity, described by likelihood ratios and one-period-ahead conditional probabilities, changes as time and uncertainty unfold. Consequently, (4) suggests that the conventional wisdom that changes in portfolios are fundamentally driven by heterogeneity in priors is correct as long as this margin of heterogeneity persists. Bayesian updating, however, imposes a strong structure on the limit behavior of beliefs, in the sense that agents typically end up agreeing on one-period-ahead conditional probabilities. What is pending to explain is the limit behavior of likelihood ratios when one-period-ahead conditional probabilities converge. Before addressing this issue in a general setting, we consider some examples to illustrate some widespread conjectures.

Benchmark Case: Homogeneous Priors

This is a particular case of the framework analyzed by Judd *et al.* [15]. Agents have identical one-period-ahead conditional probabilities of state 1 after observing partial history s^t, $p_i\left(1 | s^t\right)$. Then, the likelihood ratio $\frac{P_{j,t}(s)}{P_{i,t}(s)} = 1$ for all t and s. Consequently,

$$a_{i,t}^{\xi'}(s) = 0 \text{ for all } t,\, s \text{ and } \xi',$$

and thus portfolios are fixed forever. In every equilibrium, agents consume their endowment every period and, then, consumption and Arrow Securities prices are simple random variables with support depending only on the aggregate endowment. More precisely,

$$\begin{aligned} c_{i,t}(s) &= \lambda_i \; y_t(s) \\ m_t^{\xi'}(s) &= \beta \; \frac{1}{2} \; \frac{y_t(s)}{y(\xi')}. \end{aligned} \qquad \square$$

From this result and as a direct consequence of the convergence of the one-period-ahead conditional probabilities, one might hastily make the following conjectures:

♦ CONJECTURE I: Portfolios converge to a fixed vector while consumption and Arrow security prices converge to some simple random variable depending only on the aggregate endowment.

♦ CONJECTURE II: Limiting portfolios, consumption and Arrow security prices are those of an otherwise identical economy where agents begin with homogeneous priors.

Example 1 shows that Conjecture II might fail even if Conjecture I holds.

Example 1: Heterogeneous Priors I

The agents' one-period-ahead conditional probabilities of state 1 are given by

$$p_1\left(1|\, s^t\right) = \frac{n_t^1(s)}{t} \text{ and } p_2\left(1|\, s^t\right) = \frac{n_t^2(s)+2}{t+4},$$

where $n_t^{\xi}(s)$ stands for the number of times state $\xi \in \{1,2\}$ has been realized up to date t on path s. Since we assume $A.0$ holds with $\theta^*(1) = \frac{1}{2}$, the Strong Law of Large Numbers implies that $p_i\left(1|\, s^t\right) \to \frac{1}{2}$ $(P^{\theta^*} - a.s.)$ as $t \to \infty$, for every agent $i \in \{1,2\}$.

By the Kolmogorov's Extension Theorem (Shiryaev [25, Theorem 3, p. 163]), there exists a unique P_i on $(S^\infty, \mathcal{F})$ associated to the agents' one-period-ahead conditional probabilities. Moreover, P_i satisfies $A.1$ and A_2 and prior beliefs over θ have density $f^i(\theta) \equiv \theta^{\alpha_1^i - 1}(1-\theta)^{\alpha_2^i - 1}$ on $(0,1)$, where $\alpha_1^i = \alpha_2^i = i$.[9] The likelihood ratio is

$$\frac{P_{1,t}(s)}{P_{2,t}(s)} = \frac{\int_0^1 P_t^\theta(s)\, d\theta}{\int_0^1 P_t^\theta(s)\;\; \theta(1-\theta)\, d\theta} = \frac{\frac{\Gamma\left[n_t^1(s)+1\right]\, \Gamma\left[n_t^2(s)+1\right]}{\Gamma[t+2]}}{\frac{\Gamma\left[n_t^1(s)+2\right]\, \Gamma\left[n_t^2(s)+2\right]}{\Gamma[t+4]}} = \frac{(t+3)\;(t+2)}{\left(n_t^1(s)+1\right)\, \left(n_t^2(s)+1\right)},$$

where Γ stands for the Gamma function.[10] The Strong Law of Large Numbers can

[9] Prior beliefs over θ follow a Beta distribution $B\left(\alpha_1^i, \alpha_2^i\right)$ on $(0,1)$, as in Morris [18].

[10] Recall that if n is an integer, then $\Gamma(n) = (n-1)!$

be applied once again to show that

$$\frac{P_{1,t}(s)}{P_{2,t}(s)} \to 4 = \frac{f^1\left(\frac{1}{2}\right)}{f^2\left(\frac{1}{2}\right)} \qquad P^{\theta^*} - a.s.$$

It follows from (4) that portfolios converge to a fixed vector, that is

$$a_{1,t}^{\xi'}(s) \to \frac{1}{1-\beta}\ y(\xi')\ \lambda_1\ \left(\left(\lambda_1 + \lambda_2\ \frac{1}{4}\right)^{-1} - 1\right), \quad \xi' \in \{1,2\} \qquad P^{\theta^*} - a.s.$$

Although security prices, asset holdings and consumption all converge, we want to underscore that only prices converge to those of an otherwise identical economy with homogeneous priors. Indeed,

$$\begin{aligned} c_{1,t}(s) &\to \tfrac{\lambda_1}{\lambda_1+\lambda_2\ \frac{1}{4}}\ y_t(s) > \lambda_1\ y_t(s), \\ m_t^{\xi'}(s) &\to \beta\ \tfrac{1}{2}\ \tfrac{y_t(s)}{y(\xi')}, \end{aligned}$$

and thus Conjecture I holds but Conjecture II does not. The reason is that even though in the limit economy agents have identical beliefs, the agents' financial wealth need not be zero as in the economy that starts with homogenous priors. In fact, the limit financial wealth distribution is endogenous and depends critically on priors as we show in Section 6. □

The following example shows that Conjecture I might be false as well.

Example 2: Heterogeneous Priors II

The agents' one-period-ahead-conditional probabilities of state 1 are given by

$$p_1\left(1|\,s^t\right) = \frac{1}{1+e^{\sqrt{1/t}}} \quad \text{and} \quad p_2\left(1|\,s^t\right) = \frac{e^{\sqrt{1/t}}}{1+e^{\sqrt{1/t}}}.$$

Observe that one-period-ahead conditional probabilities converge to $\frac{1}{2}$ for both agents and have the same entropy. That is,

$$E^{P^{\theta^*}}\left(\log p_{1,t+1}|\,\mathcal{F}_t\right) = E^{P^{\theta^*}}\left(\log p_{2,t+1}|\,\mathcal{F}_t\right).$$

The ratio of one-period-ahead conditional probabilities, $\frac{p_{1,t}(s)}{p_{2,t}(s)}$, is a random variable that takes values in $\left\{e^{\sqrt{1/t}}, \frac{1}{e^{\sqrt{1/t}}}\right\}$. The logarithm of the likelihood ratio can be written as the sum of conditional mean zero random variables as follows

$$\begin{aligned}\log\left(\frac{P_{1,t}(s)}{P_{2,t}(s)}\right) &= \log\prod_{k=1}^{t}\frac{p_{1,k}(s)}{p_{2,k}(s)} \\ &= \sum_{k=1}^{t}\left[1_{s_k=1}(s)\ \log\left(e^{\sqrt{1/t}}\right)+\left(1-1_{s_k=1}(s)\right)\ \log\left(\frac{1}{e^{\sqrt{1/t}}}\right)\right] \\ &= \sum_{k=1}^{t}x_k(s)\end{aligned}$$

where $x_k(s)\in\left\{-\sqrt{\frac{1}{k}},\sqrt{\frac{1}{k}}\right\}$, $E^{P^{\theta^*}}\left(x_k\right|\mathcal{F}_{k-1})(s)=0$ and $Var^{P^{\theta^*}}\left(x_k\right|\mathcal{F}_{k-1})(s)=E^{P^{\theta^*}}\left(x_k^2\middle|\mathcal{F}_{k-1}\right)(s)=\frac{1}{k}$. Consequently, the log-likelihood ratio is the sum of uniformly bounded random variables with zero conditional mean. Additionally, since the sum of conditional variances of x_k diverges with probability 1, it follows by Freedman [10, Proposition 4.5 (a)] that

$$\sup_t\sum_{k=1}^{t}x_k(s)=+\infty \text{ and } \inf_t\sum_{k=1}^{t}x_k(s)=-\infty \quad P^{\theta^*}-a.s.$$

and, therefore,

$$\liminf\frac{P_{1,t}(s)}{P_{2,t}(s)}=0 \text{ and } \limsup\frac{P_{1,t}(s)}{P_{2,t}(s)}=+\infty \quad P^{\theta^*}-a.s.$$

This behavior of the likelihood ratio implies that individual portfolios fluctuate infinitely often. In particular,

$$\liminf a_{i,t}^{\xi'}(s)=-\frac{1}{1-\beta}\ \lambda_i\ y(\xi') \quad \text{and} \quad \limsup a_{i,t}^{\xi'}(s)=\frac{1}{1-\beta}\ (1-\lambda_i)\ y(\xi').$$

Individual portfolios, therefore, are highly volatile because each agent's debt attains its so-called *natural debt limit* infinitely often. Consequently, Conjecture I does not hold in this example and, a priori, this is rather surprising since every agent learns the true one-period-ahead-conditional probabilities.□

Why does Conjecture I hold in example 1 while it fails in example 2? The main difference is that priors satisfy $A.1$ in example 1 but not in example 2. It turns out that when $A.1$ holds for every agent, the likelihood ratios always converge and, thus, Conjecture I holds in general.

However, to generalize these lessons to the setting described in section 2 one faces two difficulties that we avoid in the examples by carefully choosing preferences, individual endowments and priors. First, equilibrium portfolios in a more general setup are typically history dependent. Closed-form solutions for asset demands as in (4) are useful to tackle this difficulty but are a particular feature derived from log preferences and constant individual endowment shares. Second, likelihood ratios are typically complicated objects which makes the analysis of their behavior a nonstandard task. Closed-form representation for the likelihood ratio, as in the examples above, simplifies the analysis of its asymptotic properties but it is a consequence of the particular family of priors that we choose.

The rest of the paper tackles the difficulties to extend the lessons from the examples to the more general setup described in section 2. Here we offer an outline. We begin with a recursive characterization of efficient allocations and their corresponding supporting portfolios under the assumption that A.1 holds. In section 4, we show that the evolution of any Pareto optimal allocation is driven solely by the evolution of the likelihood ratios of the agents' priors and the agents' posterior beliefs over the unknown parameters, as in the examples. In section 5, we prove that the agents' financial wealth converges if and only if both the likelihood ratio as well as their beliefs (over the unknown parameters) converge. Afterwards, we tackle the difficulties associated with the lack of closed form for the likelihood ratios. In section 6, we consider a broad class of priors containing those satisfying A.1 and A.2. We apply recent results in probability theory to prove that the likelihood ratios converge with probability one, as in example 1. Finally, in section 7, we argue that is key that A.1 holds for every agent. More precisely, we construct priors such that A.1 does not hold for only one agent while it does for the other. We show that the likelihood ratio does not converge and, consequently, neither their financial wealth, nor their consumption nor their portfolios converge, as in example 2.

4 A Recursive Approach to Pareto Optimality

In this section, we provide a recursive characterization of the set of Pareto optimal allocations and prove a version of the Principle of Optimality for economies with heterogeneous prior beliefs. Throughout this section we assume that $A.0$ and $A.1$ hold.

4.1 Pareto Optimal Allocations

A feasible allocation $\{c_i^*\}_{i=1}^I$ is *Pareto optimal (PO)* if there is no alternative feasible allocation $\{\widehat{c}_i\}_{i=1}^I$ such that $U_i^{P_i}(\widehat{c}_i) > U_i^{P_i}(c_i^*)$ for all $i \in \mathcal{I}$.

Given the state of nature and prior beliefs at date zero, $s_0 = \xi$ and $\mu_0 \equiv$

$\left(\mu_{1,0},...,\mu_{I,0}\right)=\mu$, define the *utility possibility correspondence* by

$$\mathcal{U}(\xi,\mu)=\{u\in\mathbb{R}^I:\ \exists\ \{c_i\}_{i=1}^I\in Y^\infty,\ U_i^{P_i}(c_i)\geq u_i\quad\forall i,\ s_0=\xi,\ \mu_0=\mu\}.$$

Now we show that the utility possibility correspondence is well-behaved, i.e. it is compact and convex-valued. Convexity follows from the strict concavity of the utility functions while compactness is a direct consequence of the compactness of the consumption set and the continuity of the utility functions.

Lemma 3 $\mathcal{U}(\xi,\mu)$ *is compact and convex-valued for all* (ξ,μ)

Lemma 3 suggests that the set of PO allocations can be characterized as the solution to the following planner's problem. Given μ_0, s_0 and welfare weights $\alpha\in\mathbb{R}_+^I$, define

$$v^*(s_0,\mu_0,\alpha)\equiv\sup_{\{c_i\}_{i=1}^I\in Y^\infty}\quad\sum_{i=1}^I\alpha_i\ E^{P_i}\left(\sum_t\beta^t\ u_i(c_{i,t})\right),\tag{5}$$

It is straightforward to prove that this problem can be written as

$$v^*(\xi,\mu,\alpha)=\sup_{u\ \in\ \mathcal{U}(\xi,\mu)}\quad\sum_{i=1}^I\alpha_i\cdot u_i,\tag{6}$$

The maximum in (6) is attained since the problem consists in maximizing a continuous function on a set that is compact by Lemma 3.

First order conditions are necessary and sufficient to characterize the solution for the planner's problem and, consequently, the set of PO allocations. These conditions can be written as

$$\frac{\alpha_i\ P_{i,t}(s)}{\alpha_j\ P_{j,t}(s)}\ \frac{u_i'(c_{i,t}(s))}{u_j'(c_{j,t}(s))}=1\quad\text{for all } i,\ j\in\mathcal{I}\text{, for all } t \text{ and all } s.\tag{7}$$

$$\sum_{i=1}^I c_{i,t}(s)=y(s_t).\tag{8}$$

Here we explain in detail why conditions (7) and (8) imply that PO allocations are history dependent in general. Since $\frac{\alpha_j}{\alpha_i}=\frac{u_i'(c_{i,0})}{u_j'(c_{j,0})}$, the planner distributes consumption such that the ratio of marginal valuations of any two agents -which, we recall, include priors that could be subjectively held- is constant along time. Consequently, under the optimal distribution rule, the ratio of marginal utilities, $\frac{u_i'(c_{i,t}(s))}{u_j'(c_{j,t}(s))}$, must be proportional to the likelihood ratio of the agents' priors, $\frac{P_{j,t}(s)}{P_{i,t}(s)}$. Whenever this ratio is constant along time (for instance, when all agents have the same priors), the optimal distribution rule is both *time and history independent*. Therefore, individual consumption depends only upon the current shock s_t (because it determines aggregate output) and the fixed vector of welfare weights α. When agents have heterogeneous priors, instead, the likelihood ratio is typically *history dependent*.

Now we argue that this history dependence can be handled with a properly chosen set of state variables. Note that since condition (7) holds if and only if

$$\begin{aligned}\frac{u_i'(c_{i,t}(s))}{u_j'(c_{j,t}(s))} &= \frac{P_{i,s^t}(C\left(s_1,...,s_k\right))}{P_{j,s^t}(C\left(s_1,...,s_k\right))}\,\frac{u_i'(c_{i,t+k}(s))}{u_j'(c_{j,t+k}(s))} \\ &= \frac{\int_{\Delta^{K-1}}\theta\left(s_1\right)...\theta\left(s_k\right)\ \mu_{i,s^t}\left(d\theta\right)}{\int_{\Delta^{K-1}}\theta\left(s_1\right)...\theta\left(s_k\right)\ \mu_{j,s^t}\left(d\theta\right)}\,\frac{u_i'(c_{i,t+k}(s))}{u_j'(c_{j,t+k}(s))},\end{aligned}$$

then the planner does not need to know the partial history itself in order to continue the date 0 optimal plan from date t onwards. Indeed, since $\mu_{i,s^t}\left(d\theta\right) = \frac{\theta(s_t)\ \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}}\theta(s_t)\ \mu_{i,s^{t-1}}(d\theta)}$, it is sufficient that he knows the ratio of marginal utilities at date t induced by the original plan, $\frac{u_i'(c_{i,t}(s))}{u_j'(c_{j,t}(s))}$, the state of nature at date t, s_t, and the posterior beliefs, $\mu_{s^{t-1}}\left(d\theta\right)$. Moreover, since the ratio of marginal utilities at date t equals the likelihood ratio weighted by the date zero welfare weights, $\frac{\alpha_j\ P_{j,t}(s)}{\alpha_i P_{i,t}(s)}$, the difficulties stemming from the optimal plan history dependence can be handled by using $(\alpha_1 P_{1,t}(s),...,\alpha_I P_{I,t}(s),\mu_{s^{t-1}})$ as state variables summarizing the history and the state of nature at date t, s_t, describing aggregate resources.

From the discussion above, we conclude that a PO allocation cannot be fully characterized using only the agents' beliefs over the unknown parameters (that is, $\mu_{s^{t-1}}$) and s_t as state variables as in the single agent setting (see, for example, Easley and Kiefer [7]). In a multiple agent setting, instead, the planner needs to distribute consumption and because of this one needs to introduce $(\alpha_1 P_{1,t}(s),...,\alpha_I P_{I,t}(s))$ as an additional state variable.

In Section 4.2 below we present a formal exposition of this result. But first, we establish some properties of the value function v^*.

Lemma 4 *The value function $v^*(\xi,\alpha,\mu)$ is bounded and continuous for all (ξ,α,μ). Moreover, v^* is homogeneous of degree 1 (hereafter HOD 1) and increasing in α.*

To conclude this section, we characterize the utility possibility correspondence and show that the set of PO allocations can be parametrized by welfare weights α.

Lemma 5 *$u\in\mathcal{U}(\xi,\mu)$ if and only if $u\geq 0$ and $v^*(\xi,\alpha,\mu)\geq\alpha\, u$ for all $\alpha\in\Delta^{I-1}$.*

4.2 Recursive Characterization of PO Allocations

Given that PO allocations are typically history dependent, standard recursive methods cannot be applied. We tackle this issue by adapting the method developed by Lucas and Stokey [17]. They characterize recursively optimal allocation problems in a deterministic setting when the history dependence is induced by recursive preferences. We use the same strategy to characterize recursively the set of PO allocations in our stochastic framework where the history dependence is due to prior belief heterogeneity.

In order to extend the Principle of Optimality to our economy, we first provide a recursive characterization of the frontier of $\mathcal{U}(\xi,\mu)$. For each agent i, the law of motion of beliefs is given by

$$\mu_i'(\xi,\mu)\,(B) = \frac{\int_B \theta(\xi)\ \mu_i\,(d\theta)}{\int_{\Delta^{K-1}} \theta(\xi)\ \mu_i\,(d\theta)} \text{ for any } B \in \mathcal{B}(\Delta^{K-1}). \tag{9}$$

Given μ_i', we define agent i's one-period-ahead conditional probabilities recursively as

$$p^r(\xi')(\mu_i'(\xi,\mu)) = \int_{\Delta^{K-1}} \theta(\xi')\mu_i'(\xi,\mu)\,(d\theta)\,,$$

Define $\|f\| = \sup_{(\xi,\alpha,\mu)} \mid f(\xi,\alpha,\mu) : \alpha \in \Delta^{I-1} \mid$ and let

$$\begin{aligned} F &\equiv \{f : S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1}) \to \mathbb{R}_+ :\ f \text{ is continuous and } \|f\| < \infty\}. \\ F_H &\equiv \{f \in F : f \text{ is increasing and HOD 1 in } \alpha\} \end{aligned}$$

F_H is a closed subset of the Banach space F and thus a Banach space itself. Continuity is with respect to the weak topology and thus the metric on F is induced by $\|.\|$.

For any $v \in F_H$, define the operator

$$(Tv)\,(\xi,\alpha,\mu) = \max_{(c,u'(\xi'))} \sum_{i\in\mathcal{I}} \alpha_i \left\{ u_i(c_i) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi,\mu))u_i'(\xi') \right\}, \tag{10}$$

subject to

$$\sum_{i=1}^I c_i = y(\xi) \qquad \text{for all } \xi, \qquad c_i \geq 0, \qquad u'(\xi') \geq 0 \qquad \text{for all } \xi', \tag{11}$$

$$v(\xi',\alpha'(\xi'),\mu'(\xi,\mu)) \geq \sum_{i=1}^I \alpha_i'(\xi')u_i'(\xi') \qquad \text{for all } \alpha'(\xi') \text{ and } \xi'. \tag{12}$$

In the following proposition we establish that the operator T is a contraction on F_H and then we apply standard arguments to show that the operator has a unique fixed point in F_H.

Proposition 6 *There is a unique function $v \in F_H$ solving (10)-(12)and the corresponding policy functions are continuous.*

Let $v \in F_H$ be the unique solution to (10) - (12), i.e. $v = Tv$, where

$$(c, \alpha', u') : S \times \mathbb{R}^I_+ \times \mathcal{P}(\Delta^{K-1}) \to \mathbb{R}_+ \times \mathbb{R}^I_+ \times \mathbb{R}^I_+$$

denote the corresponding set of policy functions. Given (s_0, α_0, μ_0), we say that a set of policy functions $(c, \alpha', u^{'})$ generates an allocation $\widehat{c}$ if

$$\begin{aligned} \widehat{c}_{i,t}(s) &= c_i(s_t, \alpha_t(s)), \\ \alpha_{t+1}(s) &= \alpha'(s_t, \alpha_t(s), \mu_{s^{t-1}})(s_{t+1}), \\ \mu_{s^t} &= \mu'(s_t, \mu_{s^{t-1}}), \end{aligned}$$

for all i and all $t \geq 0$ and $s \in S^\infty$ where $\mu_{s^{-1}} = \mu_0$.

In the recursive dynamic program defined by (10) - (12), the planner takes as given (ξ, α, μ) and allocates current consumption and continuation utility levels among agents. It follows from convexity of $\mathcal{U}(\xi, \mu)$ that for a given vector of utility levels in the frontier, there is an associated vector of welfare weights (which is unique up to a normalization). Therefore, the optimal choice of continuation utility levels induces a law of motion for welfare weights. Now we show that there is a one-to-one mapping between the set of PO allocations and the allocations generated by the optimal policy functions solving (10) - (12).

Proposition 7 (Principle of Optimality) *$v^* \in F_H$ is the unique solution to (10) - (12). Moreover, an allocation $(c^*_i)^I_{i=1}$ is PO given (ξ, α, μ) if and only if it is generated by the set of policy functions solving (10) - (12).*

Informally, this result can be grasped as follows. The characterization of the solution to the sequential formulation of the planner's problem hints that once the planner knows both the likelihood ratio weighted by the date zero welfare weights and the beliefs at date t, he can continue the optimal plan from date t onwards. It is key to understand that the consumption plan from date $t+1$ onwards can be summarized by its associated utility level. Proposition 7 shows that the date zero optimal plan is consistent in the sense that the continuation plan is indeed the solution from date t onwards.

Now we define the set of policy functions solving problem (10) - (12) . The law of motion for agent i's beliefs at (ξ, μ) is given by (9) and $c_i(\xi, \alpha)$ is the unique solution to

$$c_i(\xi, \alpha) + \sum_{h \neq i} (u_h')^{-1}(\frac{\alpha_i}{\alpha_h} \; u_i'(c_i(\xi, \alpha))) = y(\xi). \tag{13}$$

for each $i \in \mathcal{I}$, where $(u_h')^{-1}$ denotes the inverse of u_h'. Finally, the law of motion for welfare weights is

$$\alpha_i'(\xi, \alpha, \mu)(\xi') = \frac{\alpha_i \; p^r(\xi')(\mu_i'(\xi, \mu))}{\sum_h \alpha_h \; p^r(\xi')(\mu_h'(\xi, \mu))} = \frac{\alpha_i \int \theta(\xi')\mu_i'(\xi, \mu) \, (d\theta)}{\sum_h \alpha_h \int \theta(\xi')\mu_h'(\xi, \mu) \, (d\theta)}. \tag{14}$$

It follows by standard arguments that (13) is the corresponding consumption policy function. The (normalized) law of motion for the welfare weights (14) follows from the first order conditions with respect to the continuation utility levels for each individual. Observe that the normalization is harmless since optimal policy functions are HOD zero with respect to α. (see Lucas and Stokey [17] for related results).

5 Determinants of the Financial Wealth Distribution

In this section we study the determinants of the financial wealth distribution that supports a dynamically complete markets equilibrium allocation. First, we characterize individual financial wealth recursively as a time invariant function of (ξ, α, μ). The current state, ξ, captures the impact of changes in aggregate output while (α, μ) summarizes and isolates the dependence upon history introduced by the evolving degree of heterogeneity. Later, we employ a properly adapted recursive version of the Negishi's approach to pin down the PO allocation that can be decentralized as a competitive equilibrium without transfers.

Given (ξ, α, μ), we construct individual consumption using (13) and define the stochastic discount factor by

$$M(\xi, \alpha, \mu)(\xi') = \beta \; p^r(\xi')(\mu_1'(\xi, \mu)) \; \frac{u_1'(c_1(\xi', \alpha'(\xi, \alpha, \mu)(\xi'))}{u_1'(c_1(\xi, \alpha))}, \tag{15}$$

where the choice of agent 1 to define M is without loss of generality since Pareto optimality implies that the intertemporal marginal rates of substitution are equalized across agents.

The functional equation that determines agent $i's$ financial wealth is

$$A_i(\xi, \alpha, \mu) = c_i(\xi, \alpha) - y_i(\xi) + \sum_{\xi'} M(\xi, \alpha, \mu)(\xi')\ A_i(\xi', \alpha', \mu'), \tag{16}$$

where $\mu'(\xi, \mu)$ and $\alpha'(\xi, \alpha, \mu)(\xi')$ are given by (9) and (14), respectively. Note that (16) computes recursively the present discounted value of agent i's excess demand.

In Proposition 8, we show that A_i is well-defined. Furthermore, we apply Negishi's approach to show that there exist a welfare weight such that A_i is zero for every i.

Proposition 8 *Suppose A.0 and A.1 hold. Then, there is a unique continuous function A_i solving (16). Moreover, for each (s_0, μ_0) there exists $\alpha_0 = \alpha(s_0, \mu_0) \in \mathbb{R}^I_+$ such that $A_i(s_0, \alpha_0, \mu_0) = 0$ for all i.*

5.1 The Fixed Equilibrium Portfolio Property

We say that the *fixed equilibrium portfolio (FEP hereafter)* property holds if there exists $\{a_i(1), ..., a_i(K)\} \in \Re^K$ such that $a_i(\xi) = A_i(\xi, \alpha, \mu)$ for all (ξ, α, μ) and all i. If the *FEP* property holds, any portfolio that decentralizes a PO allocation with a fixed set of non-redundant assets is kept constant over time and across states.

Judd *et al.* [15] show that the *FEP* property is always satisfied after a once-and-for all initial rebalancing when agents have homogeneous priors. Indeed, in their setting the solution to (16) is independent of (α, μ), i.e. $a_i(\xi) = A_i(\xi, \alpha, \mu)$ for all (ξ, α, μ), and, therefore, the agents' financial wealth is a vector in $\mathbb{R}^K$ in any dynamically complete markets equilibrium.

In our setting, instead, portfolios typically change as the welfare weights determining the evolution of the wealth distribution change as time and uncertainty unfold. Therefore, the *FEP* property does not hold in a dynamically complete markets equilibrium when priors are heterogenous implying that the result in Judd *et al.* [15] is not robust to the introduction of this margin of heterogeneity. However, since agents observe the same data and update their priors in a Bayesian fashion, a pending deeper question is whether this trading activity fades out as this margin of heterogeneity vanishes. Stated in a more technical language: Do welfare weights necessarily converge, exhausting changes in portfolios? Our recursive approach permits to study this issue directly.

The following proposition, a direct consequence of the continuity of A_i, relates the asymptotic behavior of α_t and μ_{s^t} with the set of paths where the *FEP* property holds asymptotically. Given (s_0, α_0, μ_0), define

$$\alpha_{t+1}(s) = \alpha'(s_t, \alpha_t(s), \mu_{s^{t-1}})(s_{t+1}),$$

where $\mu_{s^t} = \mu'(s_t, \mu_{s^{t-1}})$, $\alpha(s_0) = \alpha_0$ and $\mu_{s^{-1}} = \mu_0$.

Proposition 9 *Suppose A.0 and A.1 hold. If $\alpha_t(s)$ and μ_{s^t} converge on s, then $A_i(\xi, \alpha_t(s), \mu_{s^{t-1}})$ converges on s for every i and $\xi \in S$ and, consequently, the FEP property holds asymptotically.*

Proposition 9 underscores that if a PO allocation can be decentralized through a sequence of markets, the associated wealth distribution converges to a fixed vector for each ξ on every path s on which both $\alpha_t(s)$ and $\mu_{s^{t-1}}$ converge. Consequently, asset trading reduces to the minimum.

The limit behavior of $\mu_{s^{t-1}}$ under Bayesian updating is well understood (see Lemma 2). On the other hand, very little is known about the evolution of the welfare weights that decentralize a PO allocation. We address this issue in the following section.

6 Limiting Welfare Weights under A.1

From condition (14) and Proposition 7, the ratio of welfare weights is

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\int_{\Delta^{K-1}} \theta\left(s_t\right) \ \mu_{i,s^{t-1}}\left(d\theta\right)}{\int_{\Delta^{K-1}} \theta\left(s_t\right) \ \mu_{j,s^{t-1}}\left(d\theta\right)} \frac{\alpha_{i,t-1}(s)}{\alpha_{j,t-1}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{P_{i,t}\left(s\right)}{P_{j,t}\left(s\right)}, \tag{17}$$

and, therefore, the asymptotic behavior of $\alpha_t(s)$ depends on the limit behavior of the likelihood ratios $\frac{P_{i,t}(s)}{P_{j,t}(s)}$.

Here, we show that when agents agree on the likelihood function generating the data (A.1 holds), likelihood ratios converge and so do welfare weights. However, we need to distinguish the case where the support of the agents' prior beliefs is countable from that when it is uncountable. When the support is countable, the true probability distribution is *always* absolutely continuous with respect to the agents' priors and, therefore, the convergence of likelihood ratios follows from the well-known result in Blackwell and Dubins [1]. The assumption of countable support, however, seems too strong since it rules out, for instance, the case of prior beliefs that satisfy assumption A.2. When A.2 holds, the probability distribution that generates the data is *never* absolutely continuous with respect to the agents' priors and so Blackwell and Dubins' result does not apply.[11] Nonetheless, we show that likelihood ratios converge applying a recent result by Phillips and Ploberger [21].

[11] Blume and Easley [3] also emphasize this point.

6.1 Countable Support

We first consider the case where the support of every agent's prior belief is countable (i.e., for every i, the set $B \in \mathcal{B}(\Delta^{K-1})$ such that $\mu_{i,0}(B) = 1$ is countable) and, therefore, the true probability distribution is absolutely continuous with respect to the agents' priors. As Blackwell and Dubins [1] show, this condition is equivalent to the convergence to a positive (finite) number of the ratio of the agent's prior through date t to the true probability distribution of the first t states. Indeed, in Proposition 10 we show that for every agent i,

$$\frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} \rightarrow \mu_{i,0}(\theta^*) \qquad P^{\theta^*} - a.s. \tag{18}$$

Proposition 10 *Suppose A.0 and A.1 hold. If the support of every agent's prior belief is countable and $\mu_{i,0}(\theta^*) > 0$, then (18) holds.*

In turn, Proposition 10 implies that the agent's likelihood ratios also have a finite positive limit and consequently

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{P_{i,t}(s)}{P_{j,t}(s)} \rightarrow \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{\mu_i(\theta^*)}{\mu_j(\theta^*)} \qquad P^{\theta^*} - a.s.$$

Since A_h is homogeneous of degree zero and $\mu_{h,s^{t-1}}$ converges weakly to δ_{θ^*} for every agent h, it follows by Lemma 2 that for every state $\xi \in \{1, ..., K\}$,

$$A_h(\xi, \alpha_t(s), \mu_{s^{t-1}}) \rightarrow A_h(\xi, \alpha^*, \delta_{\theta^*}) \qquad P^{\theta^*} - a.s.$$

where, for every h, $\alpha_h^* = \alpha_{h,0}\, \mu_{h,0}(\theta^*)$ and $\alpha_{h,0}$ is the welfare weight defined in Proposition 8. Therefore, we obtain the following result which completely characterizes the limiting properties of the economy.

Theorem 11 *Suppose A.0 and A.1 hold. If the support of every agent's prior belief is countable and $\mu_{i,0}(\theta^*) > 0$, then every efficient allocation converges to the Pareto optimal allocation parametrized by $\left[\alpha_{1,0}\ \mu_{1,0}(\theta^*), ..., \alpha_{I,0}\ \mu_{I,0}(\theta^*)\right]$, $P^{\theta^*} - a.s.$ Furthermore, the FEP property holds asymptotically where $a_h(\xi) = A_h(\xi, \alpha^*, \delta_{\theta^*})$ for all ξ and $h \in \mathcal{I}$.*

6.2 Uncountable Support

Now we turn to the case where the agent's prior satisfies $A.1$ and $A.2$. Since Blackwell and Dubins' result does not apply, we invoke a result in Phillips and Ploberger [21, Theorem 4.1] (stated in the appendix for completeness) to establish that there exists a sequence of measures $Q_{h,t}$ on $(S^{\infty}, \mathcal{F}_t)$ that approximates $P_{h,t}$ in the sense that the likelihood ratio $\frac{P_{h,t}}{Q_{h,t}}$ converges to 1.

Although there are several alternative asymptotically equivalent forms for $Q_{h,t}$, we find the following representation particularly useful

$$\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)} = \frac{\sqrt{2\ \pi}\ f^h(\theta^*)}{B_t^{1/2}(s)}\ e^{l_t(\widehat{\theta}_t(s))}, \tag{19}$$

where $l_t(\theta) \equiv \ln \frac{P_t^{\theta}}{P_t^{\theta^*}}$, $\widehat{\theta}_t$ is the Maximum Likelihood Estimator (MLE) of θ, $B_t(\theta)$ is the conditional quadratic variation of the score and $B_t = B_t(\theta^*)$.

In fact, under assumptions $A.0$, $A.1$ and $A.2$, the aforementioned result by Phillips and Ploberger can be handled to show that

$$\frac{1}{\frac{\sqrt{2\ \pi}\ f^h(\theta^*)}{(t\ \phi^*)^{1/2}}}\ \frac{P_{h,t}(s)}{P_t^{\widehat{\theta}_t}(s)} \to 1 \qquad P^{\theta^*} - a.s., \tag{20}$$

where ϕ^* is a constant depending upon θ^* that we define properly in the Appendix.

Proposition 12 *Suppose A.0 and A.1 hold. If every agent's prior belief satisfies A.2, then (20) holds.*

This result can be manipulated to show that if agent i and j's priors satisfy A.1 and A.2, then

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}}\ \frac{P_{i,t}(s)}{P_{j,t}(s)} \to \frac{\alpha_{i,0}}{\alpha_{j,0}}\ \frac{f_i(\theta^*)}{f_j(\theta^*)} \qquad P^{\theta^*} - a.s.$$

By a reasoning analogous to the one we used in the countable case, it follows that for every state $\xi \in \{1, ..., K\}$,

$$A_h(\xi, \alpha_t(s), \mu_{s^{t-1}}) \to A_h(\xi, \alpha^*, \delta_{\theta^*}) \qquad P^{\theta^*} - a.s.$$

where, for every h, $\alpha_h^* = \alpha_{h,0}\ f_i(\theta^*)$ and $\alpha_{h,0}$ is the welfare weight defined in Proposition 8. We summarize all these results in the following theorem.

Theorem 13 *Suppose A.0 and A.1 hold. If every agent's prior belief satisfies A.2, then every efficient allocation converges to the Pareto optimal allocation parametrized by* $[\alpha_{1,0}\ f_1(\theta^*), ..., \alpha_{I,0}\ f_I(\theta^*)]$, $P^{\theta^*} - a.s.$ *Furthermore, the FEP property holds asymptotically where* $a_h(\xi) = A_h(\xi, \alpha^*, \delta_{\theta^*})$ *for all* ξ *and* $h \in \mathcal{I}$.

6.3 Discussion

Theorems 11 and 13 argue forcefully that when the true parameter is in the support of every agent's prior belief and they know the true likelihood function generating the data (i.e., $A.1$ holds), the equilibrium allocation of the economy with heterogeneous priors converges to that of an economy with correct priors where the wealth distribution is determined by $\{\alpha_h^*\}_{h\in\mathcal{I}}$. That is, the density of the agents' prior beliefs, evaluated at the true parameter, is sufficient to pin down the limiting wealth distribution. This result is particularly appealing since it only requires to know exogenous parameters describing the economy at date zero. Indeed, it allows to compute the limiting allocation without solving for the equilibrium.

The mechanics to obtain our results, then, is to exploit A_h's homogeneity of degree zero to normalize welfare weights and then to show the convergence of these normalized welfare weights. To get a thorough understanding, it is key first to recognize that the driving force of the equilibrium allocation dynamics is the evolution of the welfare weights. Observe that agent i' welfare weight, α_i, is the planner's current valuation of an additional unit of agent i's utility. By consistency, then, $\alpha_i\beta\int\theta\left(\xi'\right)\mu_i'(\xi,\alpha,\mu)\,(d\theta)$ is the planner's current valuation of an additional unit of agent i's next period utility at state ξ'. This is the economics behind the law of motion (17), before normalizing the welfare weights. Secondly, since the evolution of these weights is fully driven by the behavior of likelihood ratios, we are lead to study their dynamics.

However, the study of the limit behavior of these ratios is a non-trivial task. The first problem one faces is that both the numerator and the denominator are vanishing and, consequently, it is crucial to understand their relative rate of convergence. Evidently, this asks for an appropriate normalization. While looking for the proper normalization, we found some technical difficulties that forced us to treat separately the cases with countable and uncountable support. In the countable case, the analysis in Blackwell and Dubins [1] suggests that $P_t^{\theta^*}$ is the normalization that works. In the uncountable case, on the other hand, the work of Phillips and Ploberger [21] suggests that $Q_{h,t}$ is the proper normalization. Therefore, as long as $A.1$ holds and the true parameter is in the support of every agent's prior belief, we can conclude that relative welfare weights converge to positive numbers for both the countable and the uncountable case.

So far we have made two critical assumptions regarding the support of the agent's prior belief, namely, (i) it contains the true parameter and (ii) it has the same dimension for every agent. The logic behind these two assumptions is as follows. As Blume and Easley [3] and Sandroni [22] argue forcefully, when some agent learns the truth, (i) and (ii) are necessary to rule out that the likelihood ratio converges to zero for some pair of agents and, therefore, to rule out that the welfare weight goes to zero for some agent. Evidently, consumption vanishes and their wealth approaches the so-called natural debt limit (see condition (16)) for those agents whose welfare weights converge to zero. The limiting economy, therefore, mimics the economy where those agents' property rights on their individual endowments have been redistributed among the remaining agents. But then those agents are basically irrelevant to understand the properties of the long-run behavior of the individuals' portfolios supporting PO allocations.

7 Persistent Trade

In this section we give examples to illustrates the necessity of assuming that the support of every agent's contains the true parameter (section 7.1) and that every agent knows the true likelihood function (section 7.2) for the *FEP* property to hold asymptotically.

7.1 Example 3: Dogmatic Priors

Judd *et al.* [15] show that, after a once-and-for all initial rebalancing, the *FEP* property holds for economies with homogeneous priors. On the one hand, we have shown forcefully that the *FEP* property holds asymptotically provided that the agents have priors satisfying $A.1$ and the support of their prior beliefs contains the true parameter. Here we show that this last condition is necessary in the sense that when it is not satisfied, the *FEP* property may not hold even if agents' priors satisfy $A.1$, no matter how close they are to the truth and with respect to each other.

We assume there are only two agents whose priors beliefs are point masses on θ_1 and θ_2, respectively, where $\theta_1 \neq \theta_2$ and $\theta^* \ln \frac{\theta_1}{\theta_2} + (1-\theta^*) \ln \frac{1-\theta_1}{1-\theta_2} = 0$. Since agents have heterogeneous "dogmatic" priors with the same entropy, then it can be shown that both agents survive.[12] The ratio of one-period-ahead conditional probabilities, $\frac{p_1(\xi'|s^t)}{p_2(\xi'|s^t)}$, is a simple random variable that takes values in $\left\{\frac{\theta_1}{\theta_2}, \frac{1-\theta_1}{1-\theta_2}\right\}$. The logarithm of the likelihood ratio can be written as the sum of conditional zero mean random

[12] An agent survive on a path if his consumption does not converge to zero on that path. See Blume and Easley [4] for a general analysis of optimal consumption paths in i.i.d. economies where agents have degenerate prior beliefs.

variables as follows

$$\begin{aligned}
\log\left(\frac{P_{1,t}(s)}{P_{2,t}(s)}\right) &= \log\prod_{k=1}^{t}\left(\frac{\theta_1}{\theta_2}\right)^{1_{s_k=1}(s)}\left(\frac{1-\theta_1}{1-\theta_2}\right)^{1-1_{s_k=1}(s)} \\
&= \sum_{k=1}^{t}\left[1_{s_k=1}(s)\ \log\left(\frac{\theta_1}{\theta_2}\right)+\left(1-1_{s_k=1}(s)\right)\ \log\left(\frac{1-\theta_1}{1-\theta_2}\right)\right] \\
&= \sum_{k=1}^{t} x_k(s),
\end{aligned}$$

where $E^{P^{\theta^*}}\left(x_k|\,\mathcal{F}_{k-1}\right)(s) = 0$ and $var^{P^{\theta^*}}\left(x_k|\,\mathcal{F}_{k-1}\right)(s) = E^{P^{\theta^*}}\left(x_k^2\big|\,\mathcal{F}_{k-1}\right)(s) = E^{P^{\theta^*}}\left(x_k^2\right) > 0$. So, the log likelihood ratio is the sum of uniformly bounded random variables with zero conditional mean and conditional variance bounded away from zero. Once again, it follows by Freedman [10, Proposition 4.5 (a)] that

$$\sup_t \sum_{k=1}^{t} x_k(s) = +\infty \text{ and } \inf_t \sum_{k=1}^{t} x_k(s) = -\infty \quad P^{\theta^*}-a.s.,$$

and, therefore,

$$\liminf \frac{P_{1,t}(s)}{P_{2,t}(s)} = 0 \text{ and } \limsup \frac{P_{1,t}(s)}{P_{2,t}(s)} = +\infty \quad P^{\theta^*}-a.s.$$

Inspecting condition (17), it is evident that welfare weights do not converge in this example. Since prior beliefs are degenerate at θ_i, posteriors are also degenerated at θ_i and, therefore, agent i's financial wealth is $A_i(\xi, \alpha_t(s), (\delta_{\theta_1}, \delta_{\theta_2}))$. We can conclude that the *FEP* property does not necessarily hold asymptotically.

7.2 Example 4: Different Likelihood Functions

In example 2 we show that the *FEP* property does not hold asymptotically when no agent satisfies $A.1$. To underscore the importance of assuming that $A.1$ holds for every agent, here consider, instead, the case in which $A.1$ does not hold for one agent while it holds for the other. One agent, on the one hand, has a prior satisfying $A.1$ and $A.2$ and, therefore, he ends up learning the true parameter with the implication that his one-period-ahead conditional probabilities converge to the truth. The other agent, on the other hand, does not know the likelihood function generating the data (i.e., he has a wrong "model" in mind). For some partial histories his one-period-ahead conditional probabilities are correct while for some others they are incorrect. The appealing feature of this example is not only that he survives but also the *FEP* property does not hold since agent 2 generates genuine asset trading infinitely often.

For simplicity, we assume there are only two states of nature every period, that is $K=2$. For a fixed prior satisfying $A.1$ and $A.2$ for agent 1, let $\theta^* \in \Delta^1$ be an element of the support of agent 1's prior belief such that $\mu_{1,s^t} \xrightarrow{w} \delta_{\theta^*}$, $P^{\theta^*}-a.s.$ By Lemma 2 we know θ^* lies in a $\mu_{1,0}-$full measure subset of Δ^1. Choose also $\theta \in \Delta^1$ and for each partial history s^t define

$$\widetilde{p}^{\theta}\left(\xi \left| s^t\right.\right) \equiv \begin{cases} \theta^*(\xi) & \text{if } \Pi_{k=1}^t \frac{\widetilde{p}^{\theta}\left(s_k \left| s^{k-1}\right.\right)}{p_1\left(s_k\left|s^{k-1}\right.\right)} \leq 1 \\ \theta(\xi) & \text{if } \Pi_{k=1}^t \frac{\widetilde{p}^{\theta}\left(s_k \left| s^{k-1}\right.\right)}{p_1\left(s_k\left|s^{k-1}\right.\right)} > 1 \end{cases}$$

Clearly, $\widetilde{p}^{\theta}\left(\xi\left|s^t\right.\right)$ is given by the true one period-ahead conditional probability whenever the likelihood ratio $\Pi_{k=1}^t \frac{\widetilde{p}^{\theta}\left(s_k \left| s^{k-1}\right.\right)}{p_1\left(s_k\left|s^{k-1}\right.\right)}$ is smaller than or equal to one. When that ratio is strictly greater than one, on the other hand, $\widetilde{p}^{\theta}\left(\xi\left|s^t\right.\right)$ is given by $\theta \in \Delta^1$.

Now we construct a probability measure on $(S^\infty, \mathcal{F})$ with the property that, after each partial history s^t, its one period-ahead conditional probability coincides with $\widetilde{p}^{\theta}\left(\cdot\left|s^t\right.\right)$. We define probability measures $\left\{\widetilde{P}_t^{\theta}\right\}_{t=1}^{\infty}$ on $\{(S^\infty, \mathcal{F}_t)\}_{t=0}^{\infty}$ as follows:

$$\begin{aligned} \widetilde{P}_1^{\theta}(s) &\equiv \widetilde{p}^{\theta}\left(s_1\left|s_0\right.\right) \\ \widetilde{P}_{t+1}^{\theta}(s) &\equiv \widetilde{p}^{\theta}\left(s_{t+1}\left|s^t\right.\right)\ \widetilde{P}_t^{\theta}(s) \qquad \forall s=\left(s^t,...\right) \text{ and } \forall t \geq 1. \end{aligned}$$

By the Kolmogorov's Extension Theorem there exists a unique probability measure $\widetilde{P}^{\theta}$ on $(S^\infty, \mathcal{F})$ that coincides with $\left\{\widetilde{P}_t^{\theta}\right\}_{t=1}^{\infty}$ when restricted to $\{(S^\infty, \mathcal{F}_t)\}_{t=0}^{\infty}$.

Remark 1: If $\theta=\theta^*$, then $\widetilde{P}^{\theta^*}=P^{\theta^*}$.

Clearly, the family $\left\{\widetilde{P}^{\theta} : \theta \in \Delta^1\right\}$ consist of probability measures on $(S^\infty, \mathcal{F})$ such that for each $B \in \mathcal{F}$, $\theta \rightarrow \widetilde{P}^{\theta}(B)$ is $\mathcal{B}\left(\Delta^1\right)-$measurable.

7.2.1 Agent 2's priors

Now we are ready to define agent 2's priors. Clearly, there exists $0<\varepsilon<1$ such that

$$\varepsilon < \min\{\theta^*, 1-\theta^*\} \leq \max\{\theta^*, 1-\theta^*\} < 1-\varepsilon.$$

Define

$$m^* \equiv \underset{\varepsilon \leq m \leq 1-\varepsilon}{\arg\min}\left(\theta^*\ \log m + (1-\theta^*)\ \log(1-m)\right),$$

and let m_t^* denote the i.i.d. random variable that takes values m^* and $1-m^*$ in states 1 and 2, respectively. Let θ_t^* denote the i.i.d. random variable that takes values θ^* and $1-\theta^*$ in states 1 and 2, respectively.

Agent 1 knows the likelihood function generating the data and so his prior is

$$P_1(B) \equiv \int_{\Delta^1} P^{\theta}(B)\,\mu_{1,0}(d\theta) \qquad \text{for any } B \in \mathcal{F}.$$

The prior of agent 2 is

$$P_2(B) \equiv \int_{\Delta^{K-1}} \widetilde{P}^{\theta}(B)\,\mu_{2,0}(d\theta) = \widetilde{P}^{m^*}(B),$$

and agent 2's one period ahead conditional probabilities are given by

$$p_{2,t}(s) \equiv \frac{P_{2,t+1}(s)}{P_{2,t}(s)} = \frac{\widetilde{P}^{m^*}_{t+1}(s)}{\widetilde{P}^{m^*}_{t}(s)} = \widetilde{p}^{m^*}\left(s_{t+1}\left|s^t\right.\right).$$

Remark 2: Notice that agent 2's one period-ahead probabilities are infinitely often bounded away from the true one period-ahead conditional probability and so he never learns the true parameter. At first reading this seems to contradict Lemma 2 above. However, that Lemma only asserts that for almost all possible parameters, according to agent 2' prior belief, he almost surely learn the parameter value. But according to agent 2's prior belief, θ^* is in a zero measure set and so there is no reason to expect consistency when θ^* is the true parameter generating the data.

The following proposition shows that the likelihood ratio of 2's prior to 1's prior fluctuates between 1 and $+\infty$. The intuition behind this result is as follows. On the one hand, the likelihood ratio cannot be both bounded away from and greater than one eventually. If this were the case, agent 2's one-period-ahead conditional probability would be bounded away from the truth eventually. Since agent 1's one-period-ahead conditional probability converges to the truth, the likelihood ratio would converge to zero almost surely. But this contradicts the assumption that the likelihood ratio is greater than one eventually. On the other hand, the set of paths where the likelihood ratio is greater than one infinitely often has full measure. To see this, consider its complement, the set of paths where the likelihood ratio is smaller or equal to one in finite time. On those paths, agent 2's one-period-ahead conditional probability is correct in finite time and, since agent 1's prior satisfies $A.1$ and $A.2$, the likelihood ratio diverges almost surely, contradicting the initial assumption.[13] Therefore, the set of paths where the likelihood ratio is smaller than or equal to one in finite time must have zero measure. Finally, since the ratio of one-period-ahead conditional probabilities is bounded away from one infinitely often, the likelihood

[13] If agent 1 had a prior belief with countable support (so that $A.2$ does not hold) then the truth would be absolute continuous with respect to 1's prior and so the likelihood ratio would not diverge even if 2 were correct every period.

ratio exceeds any pre-specified upper bound infinitely often on the set of paths where the likelihood ratio is greater than one infinitely often.[14] Thus, it must diverge along some subsequence of periods.

Proposition 14 *Suppose A.0 holds. If agent* 1*'s prior satisfies* $A.1$ *and* $A.2$*, then* $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$ *and* $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty \quad P^{\theta^*} - a.s.$

7.2.2 Dynamics of Portfolios: On the Failure of the *FEP* property

We consider again the economy described in Section 3 where portfolios are given by (4). Proposition 14 makes it clear that agents 1 and 2 survive. Agent 1's one-period-ahead conditional probabilities converge to the truth while agent 2 makes mistakes infinitely often. However, agent 2's one-period-ahead conditional probabilities are also correct infinitely often. Whether this is sufficient to offset the disadvantage stemming from his mistakes depends on the speed of agent 1's learning process. Assumption $A.2$ ensures that this convergence rate is small enough to make both agents survive. Moreover, since the likelihood ratio fluctuations do not damp out, wealth fluctuations do not damp out either. It follows immediately that the *FEP* property fails and, consequently, asset trading purely generated by heterogeneous priors does not vanish. We summarize these results in the following proposition; the proof is omitted since it is a direct consequence of Proposition 14 and the arguments above.

Proposition 15 *Suppose A.0 holds. If agent* 1*'s prior satisfies* $A.1$ *and* $A.2$*, then,* $P^{\theta^*} - a.s.$*:*

(a) agents 1 *and* 2 *survive on s.*

(b) the wealth of agent 1 *is infinitely often close to its lower bound on s.*

(c) the FEP property does not hold on s.

7.2.3 Further Remarks

In Sandroni's [22] terminology, agent 1 *eventually makes accurate next period predictions* while agent 2 does not and yet both agent survive. At a first glance,

[14] To see why, consider the event where the ratio of agent 2's to agent 1's one-period-ahead conditional probabilities is bounded away from one. The conditional probability of that event is bounded away from zero infinitely often on the set of paths where the likelihood ratio is greater than one infinitely often. This is because only agent 1's one-period-ahead conditional probability converges to the truth on those paths. Therefore, the conditional probability of the event "the likelihood ratio exceeds a pre-specified upper bound in a fixed number of periods" is also bounded away from zero infinitely often on the set of paths where the likelihood ratio is greater than one infinitely often. To clinch the result we need to argue that such event actually occurs infinitely often. An application of Levy's conditional form of the Second Borel-Cantelli Lemma shows the latter is true on the set of paths where the likelihood ratio exceeds one infinitely often.

then, Proposition 15 (a) might seem to contradict the results in Sandroni. However, no contradiction exists because this example does not satisfy the assumptions of his propositions. Indeed, his first result applies to the case in which the truth is absolutely continuous with respect to some agent's prior, an assumption that is not satisfied in this example (again, $A.0$, $A.1$ and $A.2$ rule out absolute continuity for agent 1). His second result concerns economies where agents whose one-period-ahead conditional probabilities converge to the truth coexist with others whose one-period-ahead conditional probabilities are bounded away from the truth eventually. This result does not apply either because agent 2 does not belong to any of these categories.

This example does not fit in the general setting described by Blume and Easley [3] either since they only consider economies where every agent's prior satisfies $A.1$.[15] That is, the margin of heterogeneity in priors they consider is the one arising from differences in the dimension of the agents' support. However, since nothing in the Savage approach to decision making imposes assumption $A.1$, it is also important to address the effect of the margin of heterogeneity stemming from differences in the agents' likelihood functions (i.e., agents having different models). Since we assumed that agent 1's prior satisfies $A.1$ while agent 2's prior does not (because he does not know the true likelihood function), our example explores that margin. Our findings, stated in Proposition 15, strongly suggest that the additional assumption $A.1$ shuts down a margin of heterogeneity that might be critical not only for survival but also for asset pricing and trading volume.

8 Concluding Remarks

If agents know the likelihood function generating the data and every agent has the true probability distribution over states of nature in the support of her prior beliefs, then investors change their portfolios with the arrival of new information but these changes necessarily vanish in any dynamically complete markets equilibrium. Therefore, persistent changes in portfolios can be attributed to differences of opinion about the content of new information only if one assumes that either (i) no agent has the true parameters in the support of her prior beliefs or (ii) agents disagree on the likelihood function generating the data or (iii) the probability of the states of nature changes along time.

[15] They do have an example in which agent 1 satisfies $A.1$ while agent 2 does not and yet the latter survives. However, their example differs from ours in that agent 2 not only learns the true one-period-ahead conditional probability but also, and most importantly, likelihood ratios converge with probability one.

9 Appendix A

In this Appendix we show that (4), used throughout Examples 1 - 4, denotes the equilibrium Arrow security holdings.

First, observe that the planner's problem is

$$v^*(s_0, \mu_0, \alpha) = \max_{\{c_i\}_{i=1}^2 \in Y^\infty} \sum_{i=1}^2 \alpha_i \ E^{P_i}\left(\sum_{t=0}^\infty \beta^t \log c_{i,t}\right)$$

The first order conditions imply that

$$\alpha_i \ \beta^t \ P_{i,t}(s) \frac{1}{c_{i,t}(s)(\alpha)} = \eta_t(s)(\alpha) \qquad \text{for all } i,\ t \text{ and } s, \tag{21}$$

where $\eta_t(s)(\alpha)$ denotes the Lagrange multiplier corresponding to the feasibility constraint at date t on s. The corresponding optimal allocation is fully characterized by

$$c_{i,t}(s)(\alpha) = \frac{\alpha_i \ P_{i,t}(s)}{\alpha_i \ P_{i,t}(s) + \alpha_j \ P_{j,t}(s)} \ y_t(s) \qquad \text{for all } i,\ t \text{ and } s. \tag{22}$$

Let $\eta_{i,t}(s)(\alpha) = \frac{\eta_t(s)(\alpha)}{P_{i,t}(s)}$ and $q_{i,t}(s)(\alpha) = \frac{\eta_{i,t}(s)(\alpha)}{\eta_{i,0}(s)(\alpha)}$. Now, define

$$\begin{aligned} A_{i,0}(\alpha) &= E^{P_i}\left(\sum_{t=0}^\infty \ q_{i,t}(\alpha) \ \left(c_{i,t}(\alpha) - y_{i,t}\right)\right) \\ &= E^{P_i}\left(\sum_{t=0}^\infty \ q_{i,t}(\alpha) \ \left(\frac{\alpha_i \ P_{i,t}}{\alpha_i \ P_{i,t} + \alpha_j \ P_{j,t}} - \lambda_i\right)\right). \end{aligned}$$

Using (21) and (22) it is easy to check that

$$A_{i,0}(\alpha) = \frac{y(s_0)}{1-\beta} \ \left(\alpha_i - \lambda_i\right).$$

It is a routine exercise to show that the PO allocation corresponding to $(\alpha_1, \alpha_2) = (\lambda_1, \lambda_2)$ can be decentralized a competitive equilibrium with sequential trading where a full set of Arrow securities can be traded. To pin down the corresponding asset holdings, we first compute the value of excess demand at date t on path s

$$\begin{aligned} A_{i,t}(s) &= E^{P_i}\left(\sum_{j=0}^\infty \frac{q_{i,t+j}(\lambda_1,\lambda_2)}{q_{i,t}(\lambda_1,\lambda_2)} \ \left(c_{i,t+j}(\lambda_1, \lambda_2) - \lambda_i y_{t+j}\right) \middle| \mathcal{F}_t\right)(s) \\ &= \frac{\lambda_i \ y(s_t)}{1-\beta}\left[\left(\lambda_i + \lambda_j \ \frac{P_{j,t-1}(s)}{P_{i,t-1}(s)} \ \frac{p_j(s_t|s^{t-1})}{p_i(s_t|s^{t-1})}\right)^{-1} - 1\right]. \end{aligned}$$

Thus, equilibrium portfolios are

$$a_{i,t}^{\xi'}(s) = \frac{\lambda_i \ y(\xi')}{1-\beta} \quad \left(\left(\lambda_i + \lambda_j \ \frac{P_{j,t}(s)}{P_{i,t}(s)} \ \frac{p_j(\xi'|s^t)}{p_i(\xi'|s^t)}\right)^{-1} - 1\right), \ \ \xi' \in \{1, 2\}.$$

10 Appendix B

Proof of Lemma 1:. Observe first that

$$p_i(s_k \left| s^{k-1}\right.) = \int_{\Delta^{K-1}} \Pi^\theta(s_k \left| s^{k-1}\right.)\mu_{i,s^{k-1}}(d\theta)$$

for any $1 \leq k \leq t$. Then, we have that

$$\begin{aligned}
\int_{\Delta^{K-1}} \Pi^\theta_{s^t}(B)\,\mu_{i,s^t}(d\theta) &= \int_{\Delta^{K-1}} \Pi^\theta_{s^t}(B)\,\frac{\Pi^\theta(s_t\left|s^{t-1}\right.)\mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}}\Pi^\theta(s_t\left|s^{t-1}\right.)\mu_{i,s^{t-1}}(d\theta)} \\
&= \frac{1}{p_i(s_t\left|s^{t-1}\right.)}\cdots\frac{1}{p_i(s_1\left|s_0\right.)}\int_{\Delta^{K-1}} \Pi^\theta_{s^t}(B)\;\Pi^\theta(s_t\left|s^{t-1}\right.)\ldots\Pi^\theta(s_1\left|s_0\right.)\,\mu_{i,0}(d\theta) \\
&= \frac{1}{P_i(C(s^t))}\int_{\Delta^{K-1}}\frac{\Pi^\theta(B_{s^t})}{\Pi^\theta(C(s^t))}\;\Pi^\theta\left(C(s^t)\right)\mu_{i,0}(d\theta) \\
&= \frac{1}{P_i(C(s^t))}\int_{\Delta^{K-1}}\Pi^\theta(B_{s^t})\;\;\mu_{i,0}(d\theta) \\
&= \frac{1}{P_i(C(s^t))}\int_{\Delta^{K-1}}\Pi^\theta\left(C(s^t)\cap B_{s^t}\right)\;\;\mu_{i,0}(d\theta) \\
&= \frac{P_i\left(C(s^t)\cap B_{s^t}\right)}{P_i(C(s^t))} = P_{i,s^t}(B)
\end{aligned}$$

since $B_{s^t} \subset C(s^t)$. ∎

Proof of Lemma 3. Boundedness follows because Y^∞ is bounded. Convexity follows from the strict concavity of u_i.

To prove that $\mathcal{U}(\xi,\mu)$ is closed, take any sequence $\{u^n\}$ such that $u^n \in \mathcal{U}(\xi,\mu)$ for all n and $u^n \to \overline{u} \in \mathbb{R}^I_+$. Take the corresponding sequence $\{c^n\} \subset Y^\infty$. Since Y^∞ is compact under the sup-norm, there exists a convergent subsequence $\{c^{n_k}\}$ such that $c^{n_k} \to \overline{c} \in Y^\infty$. Thus, it follows by definition that $U_i^{P_i}(c_i^{n_k}) \geq u_i^{n_k}$ for all k and for all i. Since $U_i^{P_i}$ is continuous and $\mathbf{C}$ is compact, then $U_i^{P_i}$ is continuous under the sup-norm. Thus, it follows that $U_i^{P_i}(\overline{c}_i) \geq \overline{u}_i$, for all i. Consequently, $\overline{u} \in \mathcal{U}(\xi,\mu)$ by definition and $\mathcal{U}(\xi,\mu)$ is closed. ∎

Proof of Lemma 4. Boundedness follows because Y^∞ is bounded and $\beta \in (0,1)$.

Let $Y^k \equiv \{c \in Y : c_i(s^t) \equiv c_{i,t}(s) = 0 \text{ for all } t \geq k\}$ be the $k-$truncated set of feasible allocations. Note that $Y^k \subset Y^{k+1} \subset Y^\infty$ and define

$$v_k^*(\xi,\mu,\alpha) \equiv \max_{c\,\in\,Y^k} \sum_{i\in\mathcal{I}} \alpha_i\; U_i^{P_i}(c_i)$$

Suppose that $\left\{(\mu_i^n)_{i=1}^I\right\}$ is a sequence of probability measures such that μ_i^n converges weakly to $\overline{\mu}_i \in \mathcal{P}(\Delta^{K-1})$ for all i. Given k, note that

$$\sum_{t=0}^{k} \beta^t \int_{\Delta^{K-1}} \left(\sum_{s^t} P^\theta \left(C(s^t) \right) u_i(c_i(s^t)) \right) \mu_i^n \left(d\theta \right),$$

converges to

$$\sum_{t=0}^{k} \beta^t \int_{\Delta^{K-1}} \left(\sum_{s^t} P^\theta \left(C(s^t) \right) u_i(c_i(s^t)) \right) \overline{\mu}_i \left(d\theta \right),$$

since $P^\theta \left(C(s^t) \right)$ is continuous and bounded for all t and s^t. Thus, it follows from the Maximum Theorem that $v_k^*(\xi, \mu, \alpha)$ is continuous in (μ, α) for all ξ.

Note that $v_k^*(\xi, \mu, \alpha) \leq v_{k+1}^*(\xi, \mu, \alpha) \leq v^*(\xi, \mu, \alpha)$ for all (ξ, μ, α). Hence, $v_k^*(\xi, \mu, \alpha) \to v^*(\xi, \mu, \alpha)$ for each (ξ, μ, α) since there exists some $c^* \in Y^\infty$ attaining $v^*(\xi, \mu, \alpha)$. Now we show that this convergence is uniform.

Given any (ξ, μ, α), let $c^* \in Y^\infty$ attain $v^*(\xi, \mu, \alpha)$ and define c^{*^k} as its $k-$truncated version. Then,

$$0 \leq v^*(\xi, \mu, \alpha) - v_k^*(\xi, \mu, \alpha) \leq \sum_{i=1}^{I} \alpha_i (U_i^{P_i}(c_i^*) - U_i^{P_i}(c_i^{*^k})) \leq \frac{\beta^k}{1-\beta} \max_i u_i(\overline{y}).$$

Since $\beta \in (0, 1)$, this convergence is uniform (i.e., the RHS is independent of (ξ, μ, α)) and thus $v^*(\xi, \mu, \alpha)$ is continuous. ■

Proof of Lemma 5. For any $u \in \mathcal{U}(\xi, \mu)$, it follows by definition (6) that $v^*(\xi, \mu, \alpha) \geq \alpha\, u$ for all $\alpha \in \Delta^{I-1}$. To show the converse, suppose that $u \geq 0$ and $v^*(\xi, \mu, \alpha) \geq \alpha\, u$ for all $\alpha \in \Delta^{I-1}$ but $u \notin \mathcal{U}(\xi, \mu)$. This implies that $\nexists\ \widetilde{u} \in \mathcal{U}(\xi, \mu)$ such that $\widetilde{u} \geq u$. Since $\mathcal{U}(\xi, \mu)$ is convex, it follows by the separating hyperplane theorem that $\exists\ \omega \in \mathbb{R}_+^I / \{0\}$ such that $\omega u \geq \omega \widetilde{u}$ for all $\widetilde{u} \in \mathcal{U}(\xi, \mu)$. Since $\mathcal{U}(\xi, \mu)$ is closed, $\omega u > \omega \widetilde{u}$ for all $\widetilde{u} \in \mathcal{U}(\xi, \mu)$, where ω can be normalized such that $\omega \in \Delta^{I-1}$. But then $v^*(\xi, \omega, \mu) \geq \omega\, u > \omega \widetilde{u}$ for all $\widetilde{u} \in \mathcal{U}(\xi, \mu)$. This contradicts (6). ■

Proof of Proposition 6. We first show that $T : F_H \to F_H$.

Suppose that $f \in F_H$. Since $u_i(c_i) \leq \max u_i(\overline{y})$ and $0 \leq u_i'(\xi') \leq \left\| f(\xi') \right\|$ for all i and all ξ', it follows that $\|Tf\| < \infty$. Since $\mu^{'}(\xi, \mu)$ is weakly continuous in μ for all ξ (Easley and Kiefer [7, Theorem 1]), it follows by the Maximum Theorem that $(Tf)\,(\xi, \alpha, \mu)$ is continuous in (α, μ) for all ξ (Easley and Kiefer [7, Theorem 3]). Note that this implies that there exists a solution that attains $(Tf)\,(\xi, \alpha, \mu)$.

Observe that $\mu^{'}(\xi, \mu)$ does not depend on α and consequently the constraint correspondence is independent of welfare weights. Thus, it follows from standard arguments that $(Tf)\,(\xi, \alpha, \mu)$ is HOD 1 and increasing in α. Consequently, $T : F_H \to F_H$.

Now we show that the operator T satisfies Blackwell's sufficient conditions.

(i) Monotonicity. Suppose that $f \leq g$. Then, if for all ξ'

$$\min_{\alpha'(\xi') \in \Delta^{I-1}} \left[f(\xi', \alpha'(\xi'), \mu'(\xi, \mu)) - \sum_{i=1}^{I} \alpha'_i(\xi') u'_i(\xi') \right] \geq 0,$$

it follows that

$$\begin{aligned} & \min_{\alpha'(\xi') \in \Delta^{I-1}} \left[g(\xi', \alpha'(\xi'), \mu'(\xi, \mu)) - \sum_{i=1}^{I} \alpha'_i(\xi') u'_i(\xi') \right] \\ \geq \quad & \min_{\alpha'(\xi') \in \Delta^{I-1}} \left[f(\xi', \alpha'(\xi'), \mu'(\xi, \mu)) - \sum_{i=1}^{I} \alpha'_i(\xi') u'_i(\xi') \right] \geq 0. \end{aligned}$$

and then the constraint set is enlarged. Consequently, $(Tv)(\xi, \alpha, \mu) \leq (Tg)(\xi, \alpha, \mu)$ for all (ξ, α, μ).

(ii) Discounting. Consider any arbitrary $a > 0$ and let $\left(\widehat{c}, \widehat{u}'(\xi')\right)$ attain $T(f + a)$. Fix $(\xi', \mu'(\xi, \mu))$, denote $f(\alpha) = f(\xi', \alpha, \mu')$ and define

$$\begin{aligned} U^a &\equiv \{u \in \mathbb{R}^I_+ : \ f(\alpha) + a \geq \alpha \cdot u, \qquad \forall \alpha \in \Delta^{I-1}\}, \\ B &\equiv \{u \in \mathbb{R}^I_+ : \ u \leq u' + a, \text{ for some } u' \in U^0\}. \end{aligned}$$

To show that $B \subset U^a$, notice that $u \in B$ implies that $\alpha \cdot u \leq \alpha \cdot (u' + a) \leq \alpha \cdot f(\alpha) + a$ for all $\alpha \in \Delta^{I-1}$, since $u' \in U^0$ implies $\alpha \cdot u' \leq \alpha \cdot f(\alpha)$ for all $\alpha \in \Delta^{I-1}$.

To check that $U^a \subset B$, consider any $u \in U^a$. There are three cases to consider corresponding to different regions in Figure 1 below. (i) If $u \leq a$ (see Region I, Figure I), let $u' = 0 \in U^0$ and thus $u \in B$ (see Region I). (ii) If $u \geq a$ (see Region II), let $u' = u - a \geq 0$ and thus $u' \in U^0$ since for any $\alpha \in \Delta^{I-1}$, $\alpha \cdot u' = \alpha \cdot (u - a) = \alpha \cdot u - a \leq f(\alpha)$.

(iii) To consider the third case (see Regions III and IV), suppose to simplify that $I = 2$ and let $u_1 \geq a$ and $u_2 < a$. Fix u_2, let $\alpha \in [0, 1]$ and define

$$\begin{aligned} U^a_1(u_2) &\equiv \{u_1 \geq 0 : \ f(\alpha, 1-\alpha) + a \geq \alpha u_1 + (1-\alpha) u_2, \qquad \forall \alpha \in [0,1]\} \\ &= \{u_1 \geq 0 : \ f(\alpha, 1-\alpha) + (a - u_2) \geq \alpha(u_1 - u_2), \qquad \forall \alpha \in [0,1]\}. \end{aligned}$$

Define $\overline{u}^a_1(u_2) \equiv \sup U^a_1(u_2)$ and note that

$$\begin{aligned} \overline{u}^a_1(u_2) &= \min_{0 \leq \alpha \leq 1} \left(\frac{f(\alpha, 1-\alpha)}{\alpha} + \frac{(a - u_2)}{\alpha} \right) + u_2 \\ &= \min_{0 \leq \alpha \leq 1} \left(f(1, \frac{1}{\alpha} - 1) + \frac{(a - u_2)}{\alpha} \right) + u_2 \\ &= f(1, 0) + (a - u_2) + u_2 = f(1, 0) + a, \end{aligned}$$

where the second line follows from HOD 1 and the last line from the monotonicity assumption about f and $(a-u_2) > 0$. Very importantly, note that $\overline{u}_1^a(u_2)$ is independent of u_2 for all $u_2 \leq a$, i.e. $\overline{u}_1^a(u_2) = \overline{u}_1^a = f(1,0) + a$ for all $u_2 \leq a$.

Define $u' = (f(1,0), 0) \geq 0$ and let $\alpha \in \Delta^{I-1}$. If $\alpha_1 = 0$, then $\alpha \cdot u' = 0 \leq f(\alpha)$. If $\alpha_1 > 0$, then

$$f(\alpha) = \alpha_1 \ f\left(1, \frac{\alpha_2}{\alpha_1}\right) \geq \alpha_1 \ f(1,0) = \alpha \cdot u',$$

and thus $u' \in U^0$.[16] Finally, notice that $u \leq (\overline{u}_1^a, a) = u' + a$ and $u' \in U^0$. Consequently, we can conclude that $B = U^a$. See Figure 1 below.

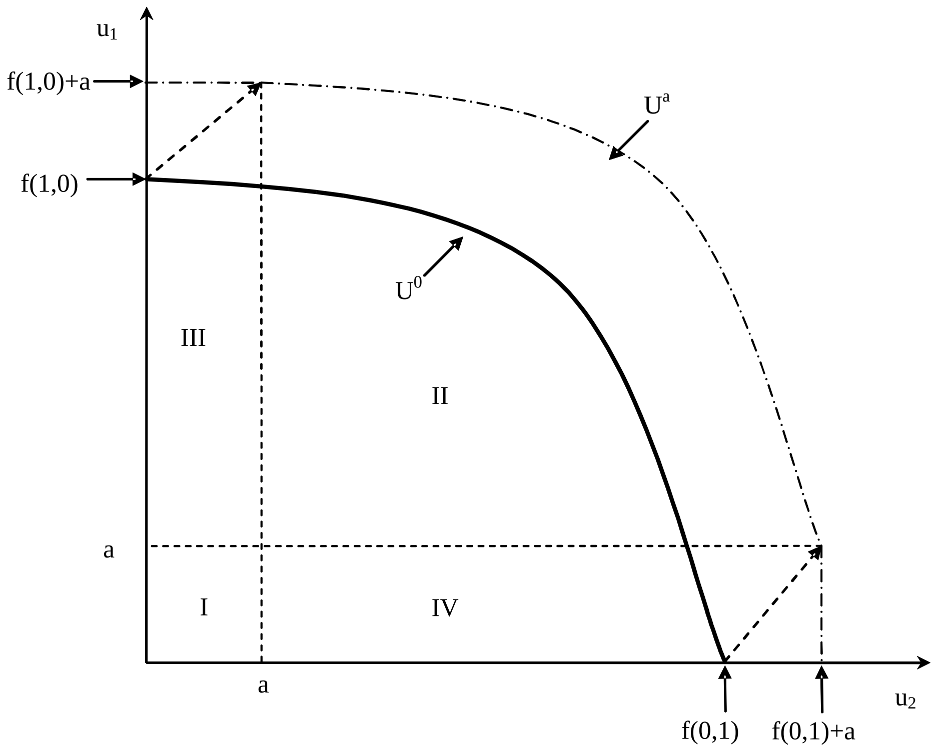

Figure 1: Figure 1

Notice that if $(\widehat{c}, \widehat{u}')$ attain $T(f+a)$, then there exists $\widetilde{u}'(\xi') \in U^0$ such that

[16] We underscore here that without assuming that f is HOD 1 and monotone (i.e., $f \in F_H$), this result does not necessarily hold. More precisely, these assumptions guarantee that $\arg\min\left(\frac{f(\alpha,1-\alpha)}{\alpha} + \frac{a-u_2}{\alpha}\right) = 1$. If any of these two assumptions is not satisfied (i.e., $f \notin F_H$), on the other hand, it is easy to construct examples such that

$$\overline{u}_1^a = \min\left(\frac{f(\alpha, 1-\alpha)}{\alpha} + \frac{a-u_2}{\alpha}\right) > \min \frac{f(\alpha, 1-\alpha)}{\alpha} + \min \frac{a-u_2}{\alpha} = \overline{u}_1^0 + a - u_2.$$

$\widetilde{u}'(\xi') \geq \widehat{u}'(\xi') - a$ for all ξ'. By monotonicity, $\left(\widehat{c}, \widetilde{u}'(\xi') + a\right)$ also attain $T(f+a)$. Observe that for any (ξ, α, μ), it follows by definition that

$$Tf(\xi, \alpha, \mu) \geq \sum_{i=1}^{I} \alpha_i \{u_i(\widehat{c}_i) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu), \alpha)\widetilde{u}_i'(\xi')\},$$

and thus

$$\begin{aligned} & T(f+a)(\xi, \alpha, \mu) - Tf(\xi, \alpha, \mu) \\ \leq \; & \sum_{i=1}^{I} \alpha_i \{u_i(\widehat{c}_i) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu), \alpha)(\widetilde{u}_i'(\xi') + a)\} \\ & - \sum_{i=1}^{I} \alpha_i \{u_i(\widehat{c}_i) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu), \alpha)\widetilde{u}_i'(\xi')\} \\ = \; & \beta a, \end{aligned}$$

and therefore, since (ξ, α, μ) was arbitrarily chosen, we can conclude that the operator T satisfies discounting. Consequently, it follows by the contraction mapping theorem that there exists a unique $v \in F_H$ such that $v = Tv$. ■

Proof of Proposition 7. Given $s_0 = \xi$ and $c_i \in C$, define for each ξ'

$${}_{\xi'}c_i = \{{}_{\xi'}c_i(s^t) = c_i(s^t) \text{ for all } t \geq 1 : (s_0, s_1) = \left(\xi, \xi'\right)\},$$

as the $\xi'-$continuation of c_i. Also, let

$$P_{i,\xi'}(s^t) = \frac{P_i(C(s^t))}{p^r(\xi')(\mu_i'(\xi, \mu))},$$

for all s^t such that $t \geq 1$. Note that

$$\begin{aligned} v^*(\xi, \mu, \alpha) &= \max_{u \in \mathcal{U}(\xi, \mu)} \sum_{i=1}^{I} \alpha_i u_i = \max_{c \in Y^\infty} \sum_{i \in \mathcal{I}} \alpha_i U_i^{P_i}(c_i) \\ &= \max_{c \in Y^\infty} \sum_{i=1}^{I} \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu)) U_i^{P_{i,\xi'}}({}_{\xi'}c_i) \right\} \\ &= \max_{\substack{c \in Y \\ \widetilde{u}'(\xi') \in \mathcal{U}(\xi', \mu'(\xi, \mu))}} \sum_{i=1}^{I} \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu)) \widetilde{u}_i'(\xi') \right\} \\ &= \max_{\substack{c \in Y \\ \widetilde{u}'(\xi') \geq 0}} \sum_{i=1}^{I} \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi, \mu)) \widetilde{u}_i'(\xi') \right\} \end{aligned}$$

subject to $\left[v^*(\xi', \alpha'(\xi'), \mu'(\xi,\mu)) - \sum_{i=1}^I \alpha_i'(\xi')u_i'(\xi')\right] \geq 0$ for all $\alpha'(\xi') \in \Delta^{I-1}$. Here, the second line follows from the definition of $U_i^{P_i}(c_i)$, the third follows from the definition of $\mathcal{U}(\xi', \mu'(\xi,\mu))$ and the last from Lemma 5. Consequently, v^* uniquely solves (RPP) by definition.

Now we claim that the set of policy functions (c, α', u') solving (RPP) generates a Pareto optimal allocation. Consider the allocation $\widehat{c}$ given by

$$\begin{aligned} \widehat{c}_{i,t}(s) &= c_i(s_t, \alpha_t(s)) \\ \alpha_{t+1}(s) &= \alpha'(s_t, \alpha_t(s), \mu_{s^{t-1}})(s_{t+1}) \\ \mu_{s^t} &= \mu'(s_t, \mu_{s^{t-1}}), \end{aligned}$$

with $\alpha(s_0) = \alpha_0$ and $\mu_{s^{-1}} = \mu_0$. Suppose that this allocation is not Pareto optimal. Then, there exists an alternative allocation $(c_i^*)_{i=1}^I$such that

$$\begin{aligned} & \sum_{i=1}^I \alpha_i \left\{ u_i(c_i^*(\xi)) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi,\mu)) U_i^{P_{i,\xi'}}({}_{\xi'}c_i^*) \right\} \\ > & \sum_{i=1}^I \alpha_i \left\{ u_i(\widehat{c}_i(\xi)) + \beta \sum_{\xi'} p^r(\xi')(\mu_i'(\xi,\mu)) U_i^{P_{i,\xi'}}({}_{\xi'}\widehat{c}_i) \right\} \\ = & v^*(\xi', \alpha', \mu'(\xi,\mu)) \end{aligned}$$

Observe that $\sum_{i=1}^I c_i^*(\xi) = y(\xi)$ and $\left(U_i^{P_{i,\xi'}}({}_{\xi'}c_i^*)\right)_{i=1}^I \in \mathcal{U}(\xi', \mu'(\xi,\mu))$ for all ξ'. It follows by Lemma 4 that

$$v^*(\xi', \alpha', \mu'(\xi,\mu)) \geq \sum_{i=1}^I \alpha_i' U_i^{P_{i,\xi'}}({}_{\xi'}c_i^*)$$

for all $\alpha' \in \Delta^{I-1}$ and all ξ'. But this contradicts that the policy functions (c, α', u') solves (RPP) for v^*.

On the other hand, since the argument holds for any arbitrary feasible $\widehat{c}$, the converse follows and, thus, we can conclude that any PO allocation $(c_i^*)_{i=1}^I$ coupled with its corresponding $\left(U_i^{P_{i,\xi'}}({}_{\xi'}c_i^*)\right)_{i=1}^I$ solve (10) - (12). ■

Proof of Proposition 8. Let F be defined as before. Consider the alternative operator $\widetilde{T}$ defined by

$$\begin{aligned} (\widetilde{T}M)(\xi,\alpha,\mu) = & \left(c_i(\xi,\alpha,\mu) - y_i(\xi)\right) u_1'(c_1(\xi,\alpha,\mu)) \\ & + \sum_{\xi'} \beta p_r(\xi')(\mu_1'(\xi,\mu)) M(\xi', \alpha'(\xi,\alpha,\mu)(\xi'), \mu'(\xi,\mu)). \end{aligned}$$

Step 1. First we check that $\widetilde{T}: F \to F$. Suppose that $M \in F$. Consider first

$$\sum_{\xi'} \beta \; p_r(\xi')(\mu_1'(\xi,\mu)) \; M(\xi', \alpha'(\xi,\alpha,\mu)(\xi'), \mu'(\xi,\mu)), \tag{23}$$

and observe that α' and μ' are both continuous. Also, it follows by definition that $p_r(\xi')(\mu_1'(\xi,\mu))$ is continuous. Thus, the expression 23 is continuous in (ξ,α,μ). Since M is bounded, its boundedness is a direct consequence of $\sum_{\xi'} P(\xi' \parallel \mu_i'(\xi,\mu)) = 1$.

Notice now that $u_1'(c_1(\xi,\alpha,\mu)) = \frac{\alpha_i}{\alpha_1} u_i'(c_i(\xi,\alpha,\mu))$ for all i. Since u_i is concave for all i, it follows that

$$0 \leq c u_i'(c) \leq u_i(c) \leq u_i(\overline{y}),$$

for all $c > 0$. Also, observe that this implies that

$$0 \leq y_i u_1'(c_1) \leq y u_1'(c_1) = \left(\sum_{i=1}^{I} c_i \right) u_1'(c_1) \leq u_1(\overline{y}) I.$$

Consequently, $(c_i(\xi,\alpha,\mu) - y_i(\xi))\, u_1'(c_1(\xi,\alpha,\mu))$ is uniformly bounded. Clearly, it is also continuous since the policy functions are continuous. Thus, we can conclude that $\widetilde{T}M \in F$.

Step 2. Now we check that $\widetilde{T}$ satisfies Blackwell's sufficient conditions and, thus, it is a contraction mapping.

We start with *discounting*. Consider any $a > 0$ and note that

$$\begin{aligned} \widetilde{T}(M+a)(\xi,\alpha,\mu) &= (c_i(\xi,\alpha,\mu) - y_i(\xi))\, u_1'(c_1(\xi,\alpha,\mu)) \\ &\quad + \sum_{\xi'} \beta p_r(\xi')(\mu_1'(\xi,\mu)) M(\xi', \alpha'(\xi,\alpha,\mu)(\xi'), \mu'(\xi,\mu)) + \beta a. \\ &= (\widetilde{T}(M)(\xi,\alpha,\mu) + \beta a. \end{aligned}$$

Monotonicity is obvious. If $M(\xi,\alpha,\mu) \geq D(\xi,\alpha,\mu)$ for all (ξ,α,μ), it is immediate that $(\widetilde{T}M)(\xi,\alpha,\mu) \geq (\widetilde{T}D)(\xi,\alpha,\mu)$ for all (ξ,α,μ).

Therefor, we can apply the contraction mapping theorem to conclude that $\widetilde{T}$ is a contraction with a unique solution $M_i \in F$ for each i.

To complete the proof, define $A_i(\xi,\alpha,\mu) = M_i(\xi,\alpha,\mu)/u_1'(c_1(\xi,\alpha,\mu))$. It can be checked immediately that A_i is a continuous function which is the unique fixed point of the operator T defined by (16) Notice that

$$\begin{aligned} \sum_i A_i(\xi,\alpha,\mu) &= \sum_i (c_i(\xi,\alpha,\mu) - y_i(\xi)) + \sum_i \sum_{\xi'} M(\xi,\alpha,\mu)(\xi') A_i(\xi',\alpha',\mu') \quad (24) \\ &= \sum_{\xi'} M(\xi,\alpha,\mu)(\xi') \sum_i A_i(\xi',\alpha',\mu'). \end{aligned}$$

Note that the operator defined by (24) has a unique solution as well. Since $R(\xi, \alpha, \mu) = 0$ for all (ξ, α, μ) solves (24), it follows by uniqueness that

$$\sum_i A_i(\xi, \alpha, \mu) = 0, \quad \text{for all } (\xi, \alpha, \mu).$$

Step 3. Finally, we show that there exists some $\alpha_0 = \alpha(s_0, \mu_0)$ such that $A_i(s_0, \alpha_0, \mu_0) = 0$ for all i, given (s_0, μ_0).

Note first that if $\alpha_i = 0$, then $c_i(\xi, \alpha, \mu) = 0$ and consequently $A_i(\xi, \alpha, \mu) < 0$ for all (ξ, μ). Define the vector-valued function g as follows:

$$g_i(\alpha) = \frac{\max[\alpha_i - A_i(s_0, \alpha, \mu_0), 0]}{\sum_i \max[\alpha_i - A_i(s_0, \alpha, \mu_0), 0]}, \tag{25}$$

for each i. Note that $H(\alpha) = \sum_i \max[\alpha_i - A_i(s_0, \alpha, \mu_0), 0]$ is positive for all $\alpha \in \Delta^{I-1}$. Also, $g_i(\alpha) \in [0, 1]$ and $\sum_i g_i(\alpha) = 1$ for all α. Thus, g is a continuous function mapping Δ^{I-1} into itself. The Brower's fixed point theorem implies that there exists some $\alpha_0 = \alpha(s_0, \mu_0)$ such that $\alpha_0 = g(\alpha_0)$.

Suppose now that $\alpha_{i,0} = 0$ for some i. By definition (25), this implies that $-A_i(s_0, \alpha_0, \mu_0) \leq 0$. But we have already argued that $-A_i(s_0, \alpha_0, \mu_0) > 0$ if $\alpha_{i,0} = g_i(\alpha_0) = 0$. This would lead to a contradiction and, hence, $\alpha_{i,0} > 0$ for all i. This implies that $\alpha_{i,0} - A_i(s_0, \alpha_0, \mu_0) > 0$ for all i. Therefore,

$$H(\alpha_0)\alpha_{i,0} = H(\alpha_0) g_i(\alpha_0) = \max[\alpha_{i,0} - A_i(s_0, \alpha_0, \mu_0), 0] = \alpha_{i,0} - A_i(s_0, \alpha_0, \mu_0).$$

This implies that $H(\alpha_0) = H(\alpha_0) \sum_i \alpha_{i,0} = \sum_i \alpha_{i,0} - \sum_i A_i(s_0, \alpha_0, \mu_0) = 1$. Therefore, $\alpha_{i,0} = \alpha_{i,0} - A^i(s_0, \alpha_0, \mu_0)$ for all i and thus $A^i(s_0, \alpha_0, \mu_0) = 0$ for all i. ■

Proof of Proposition 10. Since the support of agent i's prior belief is countable, then the true probability distribution over paths is absolutely continuous with respect to agent i's prior distribution. By Proposition B.2 in Sandroni [22], $P^{\theta^*} - a.s.$,

$$0 < \lim_{t \to \infty} \frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} < \infty, \tag{26}$$

and since P^{θ^*} is not absolutely continuous with respect to P^{θ} for all $\theta \neq \theta^*$, then P^{θ^*} is not absolutely continuous with respect to $\sum_{\theta \neq \theta^*} P_t^{\theta}(s) \frac{\mu_{i,0}(\theta)}{1-\mu_{i,0}(\theta)}$. It follows by Propositions B.1 and B.2 in Sandroni [22] that, $P^{\theta^*} - a.s.$,

$$\lim_{t \to \infty} \frac{\sum_{\theta \neq \theta^*} P_t^{\theta}(s) \frac{\mu_{i,0}(\theta)}{1-\mu_{i,0}(\theta^*)}}{P_t^{\theta^*}(s)} = 0. \tag{27}$$

Therefore, $P^{\theta^*} - a.s.$,

$$\begin{aligned}\lim_{t\to\infty}\frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} &= \mu_{i,0}(\theta^*) + \lim_{t\to\infty}\frac{\sum_{\theta\neq\theta^*} P_t^{\theta}(s)\,\mu_{i,0}(\theta)}{P_t^{\theta^*}(s)} \\ &= \mu_{i,0}(\theta^*) + \left(1-\mu_{i,0}(\theta^*)\right)\lim_{t\to\infty}\frac{\sum_{\theta\neq\theta^*} P_t^{\theta}(s)\,\frac{\mu_{i,0}(\theta)}{1-\mu_{i,0}(\theta^*)}}{P_t^{\theta^*}(s)} = \mu_{i,0}(\theta^*)\end{aligned}$$

where the last equality follows by (27). ■

The following Theorem which is due to Phillips and Ploberger [21, page 392] will be used in the proof of Proposition 12.

Theorem 16 (Phillips and Ploberger) *Assume the following conditions hold:*
(C1) $l_t(\theta)$ is twice continuously differentiable with derivatives $l_t^{(1)}(\theta)$ and $l_t^{(2)}(\theta)$.
(C2) Under P_t^{θ}, $l_t^{(1)}(\theta)$ is a zero mean L_2 martingale and $\lim_{t\to\infty} B_t(\theta)\to\infty$, $P^{\theta}-a.s.$
(C3) $\lim_{t\to\infty}\frac{l_t^{(2)}(\theta)}{B_t(\theta)}+1=0$ $P^{\theta}-a.s.$
(C4) There exist continuous functions $w_t(\theta,\theta')$ such that $w_t(\theta,\theta)=0$ and such that for some $\delta>0$ and for all $\theta,\theta'\in N_\delta(\theta^*)=\{\theta:|\theta-\theta^*|<\delta\}$ we have

$$\frac{l_t^{(2)}(\theta)-l_t^{(2)}(\theta')}{B_t}\leq w_t(\theta,\theta') \quad P^{\theta^*}-a.s.\ \textit{for each}\ t\geq 0,$$

$\lim_{t\to\infty} w_t(\theta,\theta') = w_\infty(\theta,\theta')$ $P^{\theta^*}-a.s.$ uniformly for $\theta,\theta'\in N_\delta(\theta^*)$ and $w_\infty(\theta,\theta)=0$.
(C5) $\lim_{t\to\infty}\widehat{\theta}_t=\theta^*$, $P^{\theta^*}-a.s.$
(C6) For any $\delta>0$ and $\omega_\delta=\{\theta:|\theta-\theta^*|\geq\delta\}$ we have

$$\lim_{t\to\infty} B_t^{1/2}\int_{\omega_\delta} f(\theta)\frac{P_t^{\theta}(s)}{P_t^{\theta^*}(s)}d\theta = 0 \quad P^{\theta^*}-a.s.$$

(C7) The density of the prior belief, $f(\theta)$, is continuous at θ^* with $f(\theta^*)>0$.

If $Q_{h,t}$ is the measure defined by the Radon Nykodim derivative in (19), then

$$\lim_{t\to\infty}\frac{\frac{P_{h,t}(s)}{P_t^{\theta^*}(s)}}{\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)}}=1 \quad P^{\theta^*}-a.s.$$

Proof of Proposition 12. We need to verify that (C.1) - (C.7) hold. Let n_t be the number of times that state 1 has occurred up to date t.

(C.1) holds trivially since $\ln\frac{P_t^{\theta}}{P_t^{\theta^*}}=\ln\frac{\theta^{n_t}\,(1-\theta)^{t-n_t}}{(\theta^*)^{n_t}\,(1-\theta^*)^{t-n_t}}$ is twice continuously differentiable.

(C.2) holds because $l_t^{(1)}(\theta) = \frac{n_t}{\theta} - \frac{t-n_t}{1-\theta}$ and so

$$\begin{aligned} E^{P^\theta_{s^{t-1}}}\left[l_t^{(1)}(\theta)\right] &= \theta\left(\frac{n_{1,t-1}+1}{\theta} - \frac{n_{2,t-1}}{1-\theta}\right) + (1-\theta)\left(\frac{n_{1,t-1}}{\theta} - \frac{n_{2,t-1}+1}{1-\theta}\right) \\ &= \left[n_{1,t-1} + 1 - \frac{\theta}{1-\theta} n_{2,t-1}\right] + \left[\frac{1-\theta}{\theta} n_{1,t-1} - n_{2,t-1} - 1\right] \\ &= \left[n_{1,t-1} + \frac{1-\theta}{\theta} n_{1,t-1}\right] - \left[\frac{\theta}{1-\theta} n_{2,t-1} + n_{2,t-1}\right] \\ &= \frac{n_{1,t-1}}{\theta} - \frac{n_{2,t-1}}{1-\theta} \\ &= l_{t-1}^{(1)}(\theta) \end{aligned}$$

Let $\varepsilon_k(\theta) = l_k^{(1)}(\theta) - l_{k-1}^{(1)}(\theta)$. Then $\varepsilon_k(\theta)$ takes values $\frac{1}{\theta}$ and $-\frac{1}{1-\theta}$ with probabilities θ and $1-\theta$. Therefore,

$$\begin{aligned} B_t(\theta) &= \sum_{k=1}^{t} E^{P^\theta_{s^{t-1}}}\left[\varepsilon_k(\theta)^2\right] \\ &= \sum_{k=1}^{t}\left(\theta\left(\frac{1}{\theta}\right)^2 + (1-\theta)\left(-\frac{1}{1-\theta}\right)^2\right) \\ &= \sum_{k=1}^{t}\left(\frac{1}{\theta} + \frac{1}{1-\theta}\right) \\ &= t\left(\frac{1}{\theta} + \frac{1}{1-\theta}\right) \end{aligned}$$

and we conclude that $B_t(\theta) \to \infty$ P^θ–a.s., as $t \to \infty$.

(C.3) Notice that

$$\frac{l_t^{(2)}(\theta)}{B_t(\theta)} = \frac{-\left(\frac{n_{1,t}}{\theta^2} + \frac{n_{2,t}}{(1-\theta)^2}\right)}{t\left(\frac{1}{\theta} + \frac{1}{1-\theta}\right)} \to -1 \quad P^\theta - a.s.,\ as\ t \to \infty.$$

so the desired result holds.

(C.4) Define $w_t(\theta,\theta') = w_\infty(\theta,\theta') \equiv \frac{\max\left\{\frac{1}{(\theta')^2} - \frac{1}{\theta^2}, \frac{1}{(1-\theta')^2} - \frac{1}{(1-\theta)^2}\right\}}{\frac{1}{\theta^*} + \frac{1}{1-\theta^*}}$. Clearly, $w_t(\theta,\theta')$ is continuous, $w_t(\theta,\theta) = w_\infty(\theta,\theta) = 0$ and, trivially, $w_t(\theta,\theta') \to w_\infty(\theta,\theta')$ $a.s.$ $\left(P^{\theta^*}\right)$ uniformly for every $\theta,\theta' \in N_\delta(\theta^*)$. In addition,

$$\begin{aligned} \frac{l_t^{(2)}(\theta) - l_t^{(2)}(\theta')}{B_t} &= \frac{n_{1,t}\left(\frac{1}{(\theta')^2} - \frac{1}{\theta^2}\right) + n_{2,t}\left(\frac{1}{(1-\theta')^2} - \frac{1}{(1-\theta)^2}\right)}{t\left(\frac{1}{\theta^*} + \frac{1}{1-\theta^*}\right)} \\ &= \frac{\frac{n_{1,t}}{t}\left(\frac{1}{(\theta')^2} - \frac{1}{\theta^2}\right) + \frac{n_{2,t}}{t}\left(\frac{1}{(1-\theta')^2} - \frac{1}{(1-\theta)^2}\right)}{\left(\frac{1}{\theta^*} + \frac{1}{1-\theta^*}\right)} \\ &\leq w_t(\theta,\theta') \quad P^{\theta^*} - a.s. \end{aligned}$$

(C.5) Notice that $\widehat{\theta}_t = \frac{n_{1,t}}{t} \to \theta^* \ P^{\theta^*} - a.s.$ by the SLLN.

(C.6) By the SLLN, we can take $T(s)$ such that for all $t \geq T(s)$ $a.s.$ P^{θ^*}, $\frac{n_{1,t}}{t} \in (\theta^* - \delta/2, \theta^* + \delta/2)$. In addition, there exists $\widetilde{\delta}$ such that for every $\theta \in \omega_\delta$, $\sup\limits_{x \in (\theta^* - \delta/2, \theta^* + \delta/2)} \frac{\theta^x (1-\theta)^{1-x}}{(\theta^*)^x (1-\theta^*)^{1-x}} \leq 1 - \widetilde{\delta}$. Then,

$$\begin{aligned} B_t^{1/2} \int_{\omega_\delta} f(\theta) \frac{P_t^\theta}{P_t^{\theta^*}} d\theta &= B_t^{1/2} \int_{\omega_\delta} f(\theta) \frac{\theta^{n_{1,t}} (1-\theta)^{n_{2,t}}}{(\theta^*)^{n_{1,t}} (1-\theta^*)^{n_{2,t}}} d\theta \\ &= B_t^{1/2} \int_{\omega_\delta} f(\theta) \left(\frac{\theta^{\frac{n_{1,t}}{t}} (1-\theta)^{\frac{n_{2,t}}{t}}}{(\theta^*)^{\frac{n_{1,t}}{t}} (1-\theta^*)^{\frac{n_{2,t}}{t}}} \right)^t d\theta \\ &\leq B_t^{1/2} \left(1 - \widetilde{\delta}\right)^t \\ &= \sqrt{t \ \phi^*} \ \left(1 - \widetilde{\delta}\right)^t \end{aligned}$$

where the inequality in the third line holds $P^{\theta^*} - a.s.$ The result follows because $\sqrt{t} \ \left(1 - \widetilde{\delta}\right)^t \to 0$ as $t \to \infty$.

(C.7) It follows by assumption (A.2). ■

Proof of Proposition 14. We begin with four claims that will be useful to prove the main result. Claim 17 shows that the set of paths where $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$ has full measure. Claim 20 argues that $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$ on the set of paths where the likelihood ratio is greater than one infinitely often. Claims 21 and 22 show that the latter set also has full measure.

Claim 17 $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \ P - a.s.$

Proof of Claim 17. Suppose not. Then, there exists a set of paths Ω_1 with $P^{\theta^*}(\Omega_1) > 0$ such that

$$\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1 \quad \forall s \in \Omega_1$$

Hence, there exists $T_2(s)$ such that for all $t \geq T_2(s)$

$$\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1 \quad \forall s \in \Omega_1$$

Since $p_{1,t}(s) \to \theta^*(s_t) \ P^{\theta^*} - a.s.$, there exists $T_1(s)$ such that for every $t \geq T_1(s)$, $\varepsilon < p_{1,t}(s) < 1 - \varepsilon$. Let $T(s) \equiv \max\{T_1(s), T_2(s)\}$. On the one hand, by the definition of P_2 one has that for every $T \geq T(s)$

$$\prod_{t=T(s)}^{T} \frac{m_t^*(s)}{p_{1,t}(s)} = \prod_{t=T(s)}^{T} \frac{p_{2,t}(s)}{p_{1,t}(s)} > \frac{P_{2,T(s)}(s)}{P_{1,T(s)}(s)} > 0 \quad \forall s \in \Omega_1$$

and so

$$\prod_{t=T(s)}^{T} \frac{m_t^*(s)}{p_{1,t}(s)} = \prod_{t=T(s)}^{T} \frac{p_{2,t}(s)}{p_{1,t}(s)} > \frac{P_{2,T(s)}(s)}{P_{1,T(s)}(s)} > 0 \quad \forall s \in \Omega_1$$

On the other hand, by the Strong Law of Large Numbers for uncorrelated random variables with uniformly bounded second moments,

$$\frac{1}{T - T(s)} \sum_{t=T(s)}^{T} \left(\log\left(\frac{m_t^*(s)}{p_{1,t}(s)} \right) - E^{P^{\theta^*}} \left[\log \frac{m_t^*}{p_{1,t}} | \mathcal{F}_{t-1} \right] \right) \to 0 \quad P^{\theta^*} - a.s,$$

and since $p_{1,t}(s) \to \theta^*(s_t)$ $P^{\theta^*} - a.s.$, we also have that

$$\frac{1}{T - T(s)} \sum_{t=T(s)}^{T} E^{P^{\theta^*}} \left[\log \frac{m_t^*}{p_{1,t}} | \mathcal{F}_{t-1} \right] \to E^{P^{\theta^*}} \left[\log \frac{m_t^*}{\theta_t^*} \right] < 0 \quad P^{\theta^*} - a.s.$$

Then it follows that

$$\frac{1}{T - T(s)} \sum_{t=T(s)}^{T} \log\left(\frac{m_t^*(s)}{p_{1,t}(s)} \right) \to E^{P^{\theta^*}} \left[\log \frac{m_t^*}{\theta_t^*} \right] < 0 \quad P^{\theta^*} - a.s,$$

and so

$$\sum_{t=T(s)}^{T} \log\left(\frac{m_t^*(s)}{p_{1,t}(s)} \right) \to -\infty \text{ as } T \to \infty \qquad P^{\theta^*} - a.s \quad \Leftrightarrow \prod_{t=T(s)}^{T} \frac{m_t^*(s)}{p_{1,t}(s)} \to 0 \text{ as } T \to \infty \qquad P^{\theta^*} - a.s,$$

But this implies that $P^{\theta^*}(\Omega_1) = 0$, a contradiction. ■

We continue with two results that we will need to prove Claim 20. The first is Levy's conditional form of the Second Borel-Cantelli Lemma which follows from a more general result due to Freedman ([9, Proposition 39]) and is stated without proof as Lemma 18. The second result, stated in Lemma 19, shows that on any path on which some event occurs infinitely often, the event consisting of the first event followed by any finite string of realizations of state 1 also occurs infinitely often.

For $E \in \mathcal{F}$ an event, let 1_E denote the indicator function. Recall that

$$\{\Omega_t \text{ i.o.}\} = \left\{ s : \sum_{t=1}^{\infty} 1_{\Omega_t}(s) = +\infty \right\}.$$

Also, define

$$\Omega_{1,t}^N = \{ s : s_{t-N} = ... = s_t = 1 \}.$$

Lemma 18 (Levy's Conditional form of the 2nd Borel-Cantelli Lemma) *Let* $\{\Omega_t\}_{t=0}^{\infty}$ *be a sequence of events adapted to the filtration* $\{\mathcal{F}_t\}_{t=0}^{\infty}$*. Then*

$$\sum_{t=1}^{\infty} 1_{\Omega_t}(s) = +\infty \quad P - \text{a.s. } s \in \left\{ \widetilde{s} : \sum_{t=1}^{\infty} E\left[1_{\Omega_t} | \mathcal{F}_{t-1} \right](\widetilde{s}) = +\infty \right\}.$$

Lemma 19 *Let* $\{\Omega_t\}_{t=0}^{\infty}$ *be a sequence of events adapted to the filtration* $\{\mathcal{F}_t\}_{t=0}^{\infty}$. *Then*

$$\forall N \geq 1 \quad \sum_{t=1}^{\infty} 1_{\Omega_{t-N}\cap\Omega_{1,t}^{N}}(s) = +\infty \quad P-\text{a.s. } s \in \{\Omega_t \text{ i.o.}\}.$$

Proof of Lemma 19. Notice that

$$s \in \Omega_{t-N}\cap\Omega_{1,t-1}^{N-1} \quad \Rightarrow \quad E\left[1_{\Omega_{t-N}\cap\Omega_{1,t}^{N}}\middle|\mathcal{F}_{t-1}\right](s) = P\left[s_t = 1\middle|\mathcal{F}_{t-1}\right](s) = \pi_1 > 0,$$

where we use the convention that $\Omega_{1,t}^{0} = \Omega$ to handle the case where $N = 1$, and $E\left[1_{\Omega_{t-N}\cap\Omega_{1,t}^{N}}\middle|\mathcal{F}_{t-1}\right](s)$ is non-negative otherwise.

For $s \in \{\Omega_t \text{ i.o.}\}$ arbitrarily chosen, there exists a sequence $\{t_k\}_{k=1}^{\infty}$ such that $s \in \Omega_{t_k}$ for every $k = 1, 2, \cdots$. Since $\Omega_{1,t}^{0} = \Omega$, $s \in \Omega_{(t_k+1)-1}\cap\Omega_{1,(t_k+1)-1}^{1-1}$ and therefore

$$\begin{aligned}\sum_{t=1}^{\infty} E\left[1_{\Omega_{t-1}\cap\Omega_{1,t}^{1}}\middle|\mathcal{F}_{t-1}\right](s) &\geq \sum_{k=1}^{\infty} E\left[1_{\Omega_{(t_k+1)-1}\cap\Omega_{1,t_k+1}^{1}}\middle|\mathcal{F}_{t_k}\right](s) \\ &\geq \sum_{k=1}^{\infty} P\left[s_t = 1\middle|\mathcal{F}_{t_k}\right](s) = +\infty\end{aligned}$$

and it follows by Lemma 18 that $\sum_{t=1}^{\infty} 1_{\Omega_{t-1}\cap\Omega_{1,t}^{1}}(s) = +\infty \quad P-\text{a.s. } s \in \{\Omega_t \text{ i.o.}\}$.

Suppose that the result holds for $N-1$. So, for P-a.s $s \in \{\Omega_t \text{ i.o.}\}$ arbitrarily chosen there exists $\{t_k\}_{k=1}^{\infty}$ such that $s \in \Omega_{t_k-(N-1)}\cap\Omega_{1,t_k}^{N-1} = \Omega_{(t_k+1)-N}\cap\Omega_{1,(t_k+1)-1}^{N-1}$ so that

$$\begin{aligned}\sum_{t=1}^{\infty} E\left[1_{\Omega_{t-N}\cap\Omega_{1,t}^{N}}\middle|\mathcal{F}_{t-1}\right](s) &\geq \sum_{k=1}^{\infty} E\left[1_{\Omega_{(t_k+1)-N}\cap\Omega_{1,t_k+1}^{N}}\middle|\mathcal{F}_{t_k}\right](s) \\ &\geq \sum_{k=1}^{\infty} P\left[s_t = 1\middle|\mathcal{F}_{t_k}\right](s) = +\infty\end{aligned}$$

and it follows by Lemma 18 that $\sum_{t=1}^{\infty} 1_{\Omega_{t-N}\cap\Omega_{1,t}^{N}}(s) = +\infty \quad P-\text{a.s. } s \in \{\Omega_t \text{ i.o.}\}$. That completes the induction argument and the proof. ■

Claim 20 $\limsup \frac{P_{j,t}(s)}{P_{1,t}(s)} = +\infty$ $P^{\theta^*}-a.s.$ $s \in \left\{\widetilde{s} : \frac{P_{j,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ } i.o.\right\}$

Proof of Claim 20. Let $s \in \left\{\widetilde{s} : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.}\right\}$ and $a > 1$. Since $p_{1,t}(s) \to \theta^*(s_t)$, there exists $T(s)$ such that for every $t \geq T(s)$, $\theta^*(s_t) - \frac{\varepsilon}{2} \leq p_{1,t}(s) \leq \theta^*(s_t) + \frac{\varepsilon}{2}$. Then there exists some state ξ, say state 1, such that $m_t^*(1) > \theta_t^*(1) + \frac{\varepsilon}{2}$. Let T^a be the smallest integer such that $\left(\frac{m^*}{\theta^*+\frac{\varepsilon}{2}}\right)^T > a$. Consider the event

$$\Omega_{1,t}^{T^a} \equiv \left\{\widetilde{s} : \frac{P_{2,t-1-T^a}(\widetilde{s})}{P_{1,t-1-T^a}(\widetilde{s})} > 1 \text{ and } \widetilde{s}_{t-T^a} = ... = \widetilde{s}_t = 1\right\}$$

By Lemma 19 it follows that

$$\Omega_{1,t}^{T^a} \text{ i.o. } P^{\theta^*} - a.s.\ s \in \left\{ \widetilde{s} : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$$

Therefore, $P^{\theta^*} - a.s.\ s \in \left\{ \widetilde{s} : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$, there exists a sub-subsequence $\{t_k\}_{k=0}^{\infty}$ such that $s \in \Omega_{1,t_k}^{T^a}$ and so

$$\begin{aligned}
\frac{P_{2,t_k}(s)}{P_{1,t_k}(s)} &= \frac{p_{2,t_k}(s)}{p_{1,t_k}(s)} \cdots \frac{p_{2,t_k - T^a}(s)}{p_{1,t_k - T^a}(s)} \frac{P_{2,t_k-1-T_a}(s)}{P_{1,t_k-1-T_a}(s)} \\
&= \frac{m^*}{p_{1,t_k}(s)} \cdots \frac{m^*}{p_{1,t_k - T^a}(s)} \frac{P_{2,t_k-1-T_a}(s)}{P_{1,t_k-1-T_a}(s)} \\
&> \frac{m^*}{p_{1,t_k}(s)} \cdots \frac{m^*}{p_{1,t_k - T^a}(s)} \\
&> \left(\frac{m^*}{\theta^* + \frac{\varepsilon}{2}} \right)^{T^a + 1} \\
&> a
\end{aligned}$$

where the first inequality uses the property that $\frac{P_{2,t_k-1-T_a}(s)}{P_{1,t_k-1-T_a}(s)} > 1$. It follows that

$$\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} > a \quad P^{\theta^*} - a.s. \quad s \in \left\{ \widetilde{s} : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$$

Since a was arbitrarily chosen, it follows that

$$\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty,\ P - a.s. \quad s \in \left\{ \widetilde{s} : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$$

as desired. ■

Claim 21 *$P^{\theta^*} - a.s.\ s \in \left\{ \widetilde{s} : \limsup \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} \le 1 \right\}$, there exists $T(s)$ such that $\frac{P_{2,t}(s)}{P_{1,t}(s)} \le 1\ \forall t \ge T(s)$*

Proof of Claim 21. Let $\Omega_1 \equiv \left\{ \widetilde{s} : \limsup \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} \le 1 \text{ and } \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$. Let $s \in \Omega_1$. Since $\Omega_1 \subset \left\{ s : \frac{P_{2,t}(\widetilde{s})}{P_{1,t}(\widetilde{s})} > 1 \text{ i.o.} \right\}$ then by Claim 20,

$$\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty \quad P^{\theta^*} - a.s. \quad s \in \Omega_1$$

and it follows that $P^{\theta^*}(\Omega_1) = 0$, as desired. ■

Claim 22 *$\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ i.o. $P^{\theta^*} - a.s.$*

Proof of Claim 22. Let $\Omega_1 \equiv \left\{ s : \exists T(s) \text{ such that } \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \ \forall t \geq T(s) \right\}$ and suppose that $P^{\theta^*}(\Omega_1) > 0$. Then, for every $s \in \Omega_1$

$$\prod_{k=T(s)}^{t} \frac{p_{2,k}(s)}{p_{1,k}(s)} \leq \frac{P_{1,T(s)-1}(s)}{P_{2,T(s)-1}(s)} < \infty \text{ for all } t \geq T(s)$$

By the definition of $p_{2,t}(s)$,

$$\prod_{k=T(s)}^{t} \frac{p_{2,k}(s)}{p_{1,k}(s)} = \prod_{k=T(s)}^{t} \frac{\theta_k^*(s)}{p_{1,k}(s)} \qquad \forall s \in \Omega_1$$

Since $A.2$ implies that P^{θ^*} is not absolutely continuous with respect to P_1, it follows by Propositions B.1 and B.2 in Sandroni [22] that

$$\prod_{k=T(s)}^{t} \frac{\theta_k^*(s)}{p_{1,k}(s)} \to +\infty \text{ as } t \to \infty$$

and so a contradiction is reached. It follows that $\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ i.o. $P^{\theta^*} - a.s.$ ■

Now we conclude the proof arguing that $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty \ P - a.s.$ Indeed, by Claim 22,and Claim 21, $P^{\theta^*} - a.s.$, $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ and by Claim 20 one concludes that $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty \ P^{\theta^*} - a.s.$

References

[1] BLACKWELL, D., AND L. DUBINS [1962]: "Merging of Opinions with Increasing Information." *Ann. Math. Statist.*, pp. 882-86.

[2] BLUME, L. AND D. EASLEY [2005]: "Rationality and Selection in Asset Markets," in *The Economy as an Evolving Complex System*, ed. by L. Blume, and S. Durfaul. Oxford University Press, Oxford.

[3] BLUME, L. AND D. EASLEY [2006]:"If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica*, Vol. 74, No. 4 (July), 929–966.

[4] BLUME, L. AND D. EASLEY [2006]: "The Market Organism: Long Run Survival in Markets with Heterogeneous Traders." Mimeo, Cornell University.

[5] COGLEY, T. AND T. SARGENT [2007]: "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?" mimeo, New York University.

[6] DUFFIE, D., [1988]: *Security Markets: Stochastic Models.* (Academic Press, London).

[7] EASLEY, D. AND N. M. KIEFER [1988]: "Controlling a Stochastic Process with Unknown Parameters." *Econometrica*, Vol. 56, No. 5, 1045-1064.

[8] ESPINO, E. AND T. HINTERMAIER. Forthcoming, 2007**.** "Asset Trading Volume in a Production Economy." *Economic Theory.*

[9] FREEDMAN, D. [1973]: "Another Note on the Borel-Cantelli Lemma and the Strong Law, with the Poisson approximation as a by-product ." *Annals of Probability*, vol. 1, No. 6, 910-925.

[10] FREEDMAN, D. A. [1975]: "On Tail Probabilities for Martingales." *The Annals of Probability*, vol. 3, No. 1, 100-118.

[11] HARRIS, M. AND A. RAVIV [1991]: "Differences of Opinion Make a Horserace." *Review of Financial Studies*, vol. 6 (3), 473-506.

[12] HARRISON, J. M. AND D. M. KREPS [1978]: "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *The Quarterly Journal of Economics*, vol. 92 (2), 323-336.

[13] HONG, H. AND J. STEIN [2007]: "Disagreement and the Stock Market." *Journal of Economic Perspectives*, Vol. 21, No. 2, Spring 2007, 109-128.

[14] KANDEL, E. AND N. PEARSON [1995]: "Differential Interpretation of Information and Trade in Speculative markets." *Journal of Political Economy*, vol. 103(4), 831-872.

[15] JUDD, K., KUBLER, F. AND SCHMEDDERS, K. [2003]: "Asset Trading Volume in Infinite-Horizon Economies with Dynamically Complete Markets and Heterogeneous Agents." *Journal of Finance*, 63, 2203-2217.

[16] LUCAS, R. E. JR. [1978]: "Asset Prices in an Exchange Economy." *Econometrica*, 46, 1429-1445.

[17] LUCAS, R. AND N. STOKEY [1984]: "Optimal Growth with Many Consumers." *Journal of Economic Theory*, 32, 139-171.

[18] MORRIS, S. (1995): "The Common Prior Assumption in Economic Theory." *Economics and Philosophy,* 11, 227-253.

[19] MORRIS, S. (1996): "Speculative Investor Behavior and Learning." *Quarterly Journal of Economics,* 111, 1111-1133.

[20] NEGISHI, T. [1960]: "Welfare Economics and the Existence of an Equilibrium for a Competitive Economy." *Metroeconomica,* 7, 92-97.

[21] PHILLIPS, P. C. B., AND W. PLOBERGER (1996): "An Asymptotic Theory of Bayesian Inference for Time Series." *Econometrica*, 64(2), 381–412.

[22] SANDRONI, A. (2000): "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica*, 68(6), 1303–42.

[23] SAVAGE, L. (1954): *On the Foundations of Statistics*. Wiley, New York, NY.

[24] SCHEINKMAN, J AND W. XIONG (2004): "Heterogeneous Beliefs, Speculation and Trading in Financial Markets." *Paris-Princeton Lectures on Mathematical Finance 2003, Lecture Notes in Mathematics 1847*, Springer-Verlag, Berlin.

[25] SHIRYAEV, A. (1991): *Probability*. Second Edition, New York: Springer Verlag.

[26] SCHWARTZ, L. (1965): "On Bayes Procedures." *Probability Theory and Related Fields*, 4(1), 10-26.