

THE DYNAMICS OF EFFICIENT ASSET TRADING WITH HETEROGENEOUS BELIEFS*

Pablo F. BEKER #

Emilio ESPINO ##

DEPARTMENT OF ECONOMICS
UNIVERSITY OF WARWICK

DEPARTMENT OF ECONOMICS
UNIVERSIDAD TORCUATO DI TELLA

This Draft: February 19, 2009

Abstract

This paper analyzes the dynamic properties of portfolios that sustain dynamically complete markets equilibria when agents have heterogeneous priors. We argue that the conventional wisdom that belief heterogeneity generates continuous trade and significant fluctuations in individual portfolios may be correct but it also needs some qualifications. We consider an infinite horizon stochastic endowment economy where the actual process of the states of nature consists in i.i.d. draws. The economy is populated by many Bayesian agents with heterogeneous priors over the stochastic process of the states of nature. Our approach hinges on studying portfolios that support Pareto optimal allocations. Since these allocations are typically history dependent, we propose a methodology to provide a complete recursive characterization when agents know that the process of states of nature is i.i.d. but disagree about the probability of the states. We show that even though heterogeneous priors within that class can indeed generate genuine changes in the portfolios of any dynamically complete markets equilibrium, these changes vanish with probability one if the support of every agent's prior belief contains the true distribution. Finally, we provide examples in which asset trading does not vanish because either (i) no agent learns the true conditional probability of the states or (ii) some agent does not know the true process generating the data is i.i.d.

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

*We thank Rody Manuelli and Juan Dubra for detailed comments. All the remaining errors are ours.

Corresponding Author: 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

A long-standing tenet in economics is that belief heterogeneity plays a prime role in explaining the behavior of prices and quantities in financial markets. In spite of the emphasis that economists give to efficiency, surprisingly, very little is known about the implications of belief heterogeneity on dynamically complete markets. However, there are some notable exceptions. Sandroni [20] and Blume and Easley [4] provide an analysis of the asymptotic properties of consumption. Cogley and Sargent [7] focus on asset prices. Our paper, instead, focuses on the effect of belief heterogeneity on asset trading.

Before proceeding it is useful to recall what is known about asset trading in a dynamically complete markets equilibrium when agents have identical beliefs. Judd *et al.* [14] considered a stationary Markovian economy where agents have homogeneous and degenerate beliefs but different attitudes towards risk and 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. It follows that differences in risk aversion by itself cannot explain why investors change their portfolios over time.

We consider an exchange economy where both the endowments as well as the assets returns are i.i.d. draws from a common probability distribution. Investors who are infinitely lived do not know the one-period-ahead conditional probability of the states of nature and update their priors in a Bayesian fashion as data unfolds.¹ We begin with two examples of dynamically complete markets equilibrium that illustrate that the conventional wisdom that belief heterogeneity causes significant trade may be correct but it also needs some qualifications. In example 1, agents know the true process is i.i.d. and they only disagree about the probability of the states of nature. In the long run, conditional probabilities, wealth and portfolios converge. In example 2, agents do not know the true process is i.i.d., conditional probabilities converge and yet wealth bounces back and forth between them infinitely often so that each of them holds almost all the wealth infinitely many times. This second example shows that even though agents may learn, prior belief heterogeneity may indeed generate significant fluctuations in the wealth distribution and the corresponding portfolios that do not exhaust in the long run. We argue that the different dynamics in the two examples reflect differences in the limit behavior of the likelihood ratio of the agents' priors.

This paper links the evolution of the wealth distribution and the corresponding

¹To avoid any confusion, we use 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 these parameters, we call the latter the agent's prior belief.

portfolios in any dynamically complete markets equilibrium to the evolution of the likelihood ratio. This is useful because the likelihood ratio is an exogenous variable and several properties of its limit behavior are well understood from the statistics literature (see Phillips and Ploberger [19] and references therein.) Our examples 1 and 2 raise the question of what type of belief heterogeneity matters for asset trading. In order to answer this question, we first carefully assess a class of priors satisfying two assumptions that are ubiquitous in the literature. Namely, every agent knows the likelihood function generating the data and there is at least one agent who learns in the sense that her one-period-ahead conditional probability converges to the truth. In our setup, this is ensured by assuming that every agent knows the data is generated by i.i.d. draws from a common (unknown) distribution and the support of their prior beliefs contains the true probability distribution of the states of nature as in example 1. We first show that even though heterogeneous priors in that class can indeed generate changes in the portfolios of a dynamically complete markets equilibrium, these changes vanish with probability one. Very importantly, we fully characterize the dynamics of portfolios and its corresponding limit. Afterwards, we show by means of two additional examples that if one wants to argue that heterogeneity of priors can have enduring implications on the volume of trade in a stationary environment then one needs to relax one of the aforementioned assumptions; that is, either (i) no agent learns the true conditional probability of states or (ii) some agent does not know the likelihood function generating the data.

Since solving directly for the portfolios of a dynamically complete markets equilibrium is not always possible, we follow an indirect approach developed by Espino and Hintermaier [9]. This approach hinges on studying portfolios that support Pareto optimal allocations. The difficulty is that belief heterogeneity makes optimal allocations history dependent 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. This ratio represents the novel margin of heterogeneity among agents considered in this paper, which we call *the \mathcal{B} -margin of heterogeneity*. The evolution of the \mathcal{B} -margin determines the dynamics of the optimal distribution rule of consumption and, consequently, the evolution of the wealth distribution in any dynamically complete markets equilibrium. The law of motion of this margin is typically history dependent and, very importantly, the current state and the current \mathcal{B} -margin are not enough to summarize the history. Under the assumption that every agent knows the data is generated by i.i.d. draws from a common (unknown) distribution but have different beliefs over

the unknown parameters, this history dependence can be succinctly captured by the agents' beliefs (over the parameters). This assumption allow us to use a strategy similar to Lucas and Stokey's [16] to obtain a recursive characterization of the set of Pareto optimal allocations in our stochastic framework.² The key 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 state of nature, the agents' prior beliefs over probabilities and, very importantly, the current \mathcal{B} -margin, i.e. the likelihood ratios of the agents' priors that summarize how the weight attached to each agent depends on history. We argue that the sequential formulation of the planner's problem is equivalent to a recursive dynamic program where the planner, who takes a vector of welfare weights as given, allocates current feasible consumption and assigns next period attainable utility levels among agents. The planner's optimal choice of continuation utilities induces a law of motion for welfare weights that is isomorphic to the evolution of the likelihood ratio of the agents' priors. 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 supporting portfolios) converges if and only if both the \mathcal{B} -margin vanishes and the agents' beliefs over the parameters become homogeneous.

When the agents know that the true process consists in i.i.d. draws from a common distribution and the true distribution is in the support of their priors, the well-known consistency property of Bayesian learning implies that the agents' prior belief become homogeneous with probability one. To get a thorough understanding of the limiting behavior of portfolios, therefore, what remains to be explained is the asymptotic behavior of the \mathcal{B} -margin. 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 Sandroni [20]. When the agents' prior beliefs have a positive and continuous density with support containing the true parameter, the hypothesis in Sandroni [2] are not satisfied and so we apply a result in Phillips and Ploberger [19] to show that the likelihood ratio of the agents' priors still converges with probability one. The important message here is that the heterogeneity of priors by itself can generate changes in portfolios but these changes necessarily vanish because the \mathcal{B} -margin vanishes. Furthermore, we show that portfolios converge to those of a rational expectations equilibrium of an economy where the investors' relative wealth is determined by the

²Lucas and Stokey [16] characterize recursively optimal programs in a deterministic setting where recursive preferences induce the dependence upon histories.

densities of their prior beliefs evaluated at the true parameter and the date zero welfare weight.³

To conclude we analyze the exact role played by the aforementioned assumptions on priors and we argue that it is critical that they are coupled together. We do so by providing two additional examples, each of which relax one of these assumptions, in which the \mathcal{B} -margin does not vanish and consequently portfolios change infinitely often. In example 3, agents know the data is generated by i.i.d. draws from a common distribution but they do not have the true parameter in the support of their prior beliefs and so no agent learns. We assume that their 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. Finally, example 4 underscores the importance of assuming that every agent knows the process of states consists in i.i.d. draws for the portfolios to converge. To stretch the argument to the limit, we consider an example in which only one agent does not know the data is generated by i.i.d. draws. 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 is organized as follows. In Section 2 we review the related literature. In section 3 we describe the model. In section 4 we present a simple example that illustrate the main ideas in this paper. The recursive characterization of Pareto optimal allocations is in section 5. Section 6 characterizes the asymptotic behavior of the agents' financial wealth and their corresponding supporting portfolios. Finally, sections 7 and 8 discuss when the agents' portfolio converge and when it does not. Conclusions are in section 9. Proofs are gathered in the Appendix.

2 Related Literature

This paper relates to two branches of the literature on the effect of belief heterogeneity in 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 [13] and Harris and Raviv [12] who study the implications of belief heterogeneity on asset prices and trading volume, respectively, are

³In particular, even though agents learn the true probability of states of nature, these limiting portfolios need not coincide with those of an otherwise identical economy that starts with homogeneous priors and zero financial wealth.

the leading articles of the first branch. These 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. Since they are risk neutral and have heterogeneous belief, they have different marginal valuations and so trade occurs if and only if agents "switch sides" regarding their valuation of the asset. In addition, Harrison and Kreps' [13] show that a *speculative premium* might arise, in the sense that the asset price might be strictly greater than every trader's fundamental valuation. Since each investor is absolutely convinced her model is the correct one, their disagreement does not vanish as the data unfold.

The possibility that agents learn is addressed by Morris [17] who extends Harrison and Kreps' [13] model to consider agents that have heterogeneous and non-degenerate prior beliefs over the probability distribution of dividends. He characterizes the set of prior beliefs for which the speculative premium is positive. He assumes the true process is i.i.d., investors know this fact but they have heterogeneous prior beliefs about the distribution of these draws with support containing the true distribution. Since they are Bayesian, they eventually learn the true distribution. Consequently, risk neutrality implies the price converges and the speculative premium vanishes. We underscore that asset trading does not vanish because there is always a period in the future when the asset changes hands once again. Morris' [17] 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. In this paper, we argue that these forces do not operate in a setting where agents are risk-averse and allocations are Pareto optimal. More precisely, Pareto optimality implies that the agent's intertemporal marginal rates of substitution must be equalized and, unlike in Morris [17] where they switch persistently, every trader's valuation of any future income stream *always* coincide. Consequently, there is *never* a speculative premium in spite of belief heterogeneity. Our analysis makes evident that the speculative premium is not necessarily driven by belief heterogeneity but, more importantly, by the

differences in the agents' intertemporal marginal valuations due to the existence of short-sale constraints.⁴ In our setting, portfolios might still change persistently but these changes depend 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.

Belief heterogeneity may have fundamental implications on the behavior of asset markets even in the absence of the aforementioned capital market imperfections. In the context of the Lucas [15] tree model, Cogley and Sargent [6] and [7] focus on the effects of learning and prior belief heterogeneity, respectively, on asset prices under the assumption that agents know the true likelihood function. In [6], they consider an economy 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.⁵ Cogley and Sargent [7] analyze the robustness of that finding by considering an economy with complete markets with some agents who know the true probability distribution (i.e., they add belief heterogeneity). For a plausible calibration of their model, they show that unless the agents with correct beliefs own a large fraction of the initial wealth, it takes a long time for the effect of pessimism to be erased. Their work is close in spirit to ours in that they use a general equilibrium model without any additional market imperfection. Since they are principally interested in studying the market prices of risk, however, they are silent about the implications of belief heterogeneity for trading volume. Consequently, the asset trading implications stemming purely from differences in priors are still an open question.

The second branch of the literature related to our paper analyses the market selection hypothesis and is exemplified by the work of Sandroni [20] and Blume and Easley [4]. Sandroni [20] shows that, controlling for discount factors, if the true distribution is absolute continuous with respect to some trader's prior then she survives and any other trader survives if and only if the true distribution is absolute continuous with respect to her prior as well.⁶ He also considers some cases in which the true distribution is not absolute continuous with respect to any agent's prior. He shows that

⁴Indeed, in any economy where trading constraints are occasionally binding for different agents, the agent who prices the asset changes and thus the "speculative premium" can arise naturally.

⁵Their 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.

⁶An agent is said to survive if her consumption does not converge to zero.

the entropy of priors determines survival and, therefore, an agent who persistently makes wrong predictions vanishes in the presence of a learner. Absolute continuity 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.⁷ This is precisely the case that Blume and Easley [4] consider and they prove that among Bayesian learners who have the truth in the support of their priors, only those with the lowest dimensional support can have positive consumption in the long 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.⁸

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 process consists in i.i.d. draws. In addition, it includes cases in which some agents have the truth in the support of their priors while some other agent do not learn the true one-period-ahead conditional probability and yet the latter survives as in our example 4. To the best of our knowledge this is the first example of its kind in the literature.

Our results characterizing the portfolios that support a Pareto optimal allocation are a novel contribution to the literature since neither Sandroni [20] nor Blume and Easley [4] analyze portfolio dynamics. Indeed, the mapping between consumption and its supporting portfolio is only simple when agents have degenerate homogeneous priors as in Judd *et al.* [14]. This is most evident when one consider the case where agents have homogeneous but non-degenerate prior beliefs. In this case, the distribution of consumption is time independent while the supporting portfolios are not because the state prices change as agents learn. We also contribute to the analysis of the asymptotic behavior of portfolios since it is not evident that Sandroni’s [20] and Blume and Easley’s [4] results on the limit behavior of consumption imply that (i) portfolios must converge when likelihood ratios do and, very importantly,

⁷In 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.

⁸This distinction is economically relevant because both in Blume and Easley’s [4] 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.

(ii) when portfolios converge, what the limiting portfolios are. The recursive characterization of the financial wealth distribution that we obtain allows to answer these two questions. First, since it provides a continuous mapping between the portfolios supporting a PO allocation and the investors' likelihood ratios, it makes evident that they converge together. Second, it makes possible to single out the PO allocation that can be decentralized as a competitive equilibrium without transfers by means of the application of a recursive version of the Negishi's approach. This allocation is parametrized by its corresponding welfare weight that depends upon date 0 prior beliefs, individual endowments and aggregate resources. Finally, and very importantly, it allows to conclude that the limiting wealth distribution is pinned down by the densities of their prior beliefs evaluated at the true parameter and the corresponding date 0 welfare weights.

3 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.

3.1 The Environment

Time is discrete and indexed by $t = 0, 1, 2, \dots$. The set of possible states of nature at date $t \geq 1$ is $S_t \equiv \{1, \dots, K\}$. The state of nature at date zero is known and denoted by $s_0 \in \{1, \dots, K\}$. We define the set of partial histories up to date t as $S^t = \{s_0\} \times (\times_{k=1}^t S_k)$ with typical element $s^t = (s_0, \dots, s_t)$. $S^\infty \equiv \{s_0\} \times (\times_{k=1}^\infty S_k)$ is the set of infinite sequences of the states of nature and $s = (s_0, s_1, s_2, \dots)$, 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 $C(s^t) \equiv \{s \in S^\infty : s = (s^t, s_{t+1}, \dots)\}$ of all paths whose $t + 1$ 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 $\mathcal{F}_0 \subset \dots \subset \mathcal{F}_t \subset \dots \subset \mathcal{F}$ where $\mathcal{F}_0 \equiv \{\emptyset, S^\infty\}$ is the trivial σ -algebra and \mathcal{F} is the σ -algebra generated by the algebra $\bigcup_{t=0}^\infty \mathcal{F}_t$.

For any probability measure $\Pi : \mathcal{F} \rightarrow [0, 1]$ on (S^∞, \mathcal{F}) , $\Pi_{s^t} : \mathcal{F} \rightarrow [0, 1]$ denotes its posterior distribution after observing s^t .⁹ Let $\Pi_t(s)$ be the probability of the finite history s^t , i.e. the \mathcal{F}_t -measurable function defined by $\Pi_t(s) \equiv \Pi(C(s^t))$ for all $t \geq 1$ and $\Pi_0 \equiv 1$. Let π_t be the \mathcal{F}_t -measurable function defined by $\pi_t(s) \equiv \frac{\Pi_t(s)}{\Pi_{t-1}(s)}$. That is, given the partial history s^{t-1} up to date $t - 1$, π_t is the one-period-ahead

⁹Formally, $\Pi_{s^t}(A) \equiv \frac{\Pi(A_{s^t})}{\Pi(C(s^t))}$ for every $A \in \mathcal{F}$, where $A_{s^t} \equiv \{s \in S^\infty : s = (s^t, s'), s' \in A\}$.

conditional probability of the states at date t and $\pi(\xi | s^{t-1})$ denotes its realization at $s_t = \xi$ after the partial history s^{t-1} . Finally, for any random variable $x : S^\infty \rightarrow \mathfrak{R}$, $E^\Pi(x)$ denotes its mathematical expectation with respect to Π .

Let Δ^{K-1} be the $K - 1$ dimensional unit simplex in \mathfrak{R}^K , $\mathcal{B}(\Delta^{K-1})$ be its Borel sets and $\mathcal{P}(\Delta^{K-1})$ be the set of probability measures on $(\Delta^{K-1}, \mathcal{B}(\Delta^{K-1}))$. Consider a set of probability measures on (S^∞, \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 $\mathcal{B}(\Delta^{K-1})$ -measurable for each $B \in \mathcal{F}$. This set includes the subset of probability measures on (S^∞, \mathcal{F}) uniquely induced by i.i.d. draws from a common distribution $\theta : 2^K \rightarrow [0, 1]$, where $\theta(\xi) > 0$ for all $\xi \in \{1, \dots, K\}$, 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 $\theta^* \gg 0$.

We assume the true process of states of nature is i.i.d. to ease the exposition. However, all our results hold true for any time-homogeneous Markov process.

3.2 The Economy

There is a single perishable consumption good every period. The economy is populated by I (types of) infinitely-lived agents where $i \in \mathcal{I} = \{1, \dots, I\}$ 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.

3.2.1 Preferences

We assume that agents' preferences satisfy Savage's [21] 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.¹⁰

We denote by P_i the probability measure on (S^∞, \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, for every $c_i \in \mathbf{C}$ her preferences are represented by

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

¹⁰See Blume and Easley [3] for a complete discussion on the implications of Savage's axioms.

where $\beta \in (0, 1)$ and $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuously differentiable, strictly increasing, strictly concave and $\lim_{x \rightarrow 0} \frac{\partial u_i(x)}{\partial x} = +\infty$ 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, the probability of every event $A \in \mathcal{F}$ is

$$\Pi(A) = \int_{\Delta^{K-1}} \Pi^\theta(A) \mu(d\theta), \quad (1)$$

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 process generating the data is i.i.d. but does not know the true probability of the states of nature. We state this assumption as A.1.¹¹

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

We want to emphasize that A.1 says that even though 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 $\mu^\theta : \mathcal{F} \rightarrow [0, 1]$ where

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

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

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

where $\mu_{i,0} \in \mathcal{P}(\Delta^{K-1})$ is given at date 0 and $\pi^\theta(s_t | s^{t-1}) \equiv \Pi^\theta(C(s^t)) / \Pi^\theta(C(s^{t-1}))$.

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 [23, Theorems 3.2 and 3.3]), and since our example 4 in Section 8.2 does not satisfy A.1 but it does satisfy (1), we state the consistency result in the following Lemma to make precise its scope.

¹¹The celebrated De Finetti theorem states that this is equivalent to the prior being exchangeable.

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

REMARK 1: It is ubiquitous in the learning literature related to asset pricing to assume both that (i) every agent knows the likelihood function generating the data and (ii) some agent learns the true conditional probability of the states. The latter is guaranteed in our setup by strengthening A.1 to require that the true parameter, θ^* , is in the support of some agent's prior. The case where this holds for every agent is considered in sections 5, 6 and 7. Section 8 deals with the cases in which either (i) or (ii) does not hold.

3.2.2 Endowments

Agent i 's endowment at date t is a time-homogeneous function of the current state of nature, that is $y_i(s_t) > 0$ for all $s_t \in \{1, \dots, K\}$ and the aggregate endowment is $y(s_t) \equiv \sum_{i=1}^I y_i(s_t) \leq \bar{y} < \infty$. An allocation $\{c_i\}_{i=1}^I \in \mathbf{C}^I$ is *feasible* if $c_i \in \mathbf{C}$ for all i and $\sum_{i=1}^I c_{i,t}(s) \leq y(s_t)$ for all $s \in S^\infty$. Let Y^∞ denote the set of feasible allocations.

4 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. In Section 3 we assumed that the range of utility functions was \mathfrak{R}_+ . This lower bound on utility will be used in the proofs of Theorems 3 and 4. We have verified the conclusions of those Theorems directly for all of the examples in this section and in section 8.

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 in zero net supply. 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 after partial history s^t are denoted by $m_i^{\xi'}(s)$ and $a_{i,t}^{\xi'}(s)$, respectively. We assume that agents have no endowment of Arrow securities, i.e. they have zero financial wealth at date 0.

In Appendix A we show that equilibrium consumption and portfolios are

$$\begin{aligned} c_{i,t}(s) &= \left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \right)^{-1} \lambda_i y_t(s), \\ a_{i,t}^{\xi'}(s) &= \frac{1}{1-\beta} y(\xi') \lambda_i \left(\left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \frac{p_j(\xi'|s^t)}{p_i(\xi'|s^t)} \right)^{-1} - 1 \right), \quad \xi' \in \{1, 2\}, \end{aligned} \quad (3)$$

where $P_{i,t}(s) = P_i(C(s^t))$ and $p_i(\xi'|s^t) = P_i(C(s^t, \xi'))/P_i(C(s^t))$. Observe that individual portfolios at date t are completely determined by the likelihood ratio at

$t + 1, \frac{P_{j,t+1}}{P_{i,t+1}}$. Portfolios converge if and only if the likelihood ratio converges. Thus, changes in portfolios are purely determined by the heterogeneity of priors.

The relevant margin of heterogeneity described by likelihood ratios changes as time and uncertainty unfold. Consequently, (3) 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 the one-period-ahead conditional probability. What is pending to explain is the limit behavior of likelihood ratios when one-period-ahead conditional probabilities converge.

Benchmark Case: Homogeneous Priors

Agents have identical one-period-ahead conditional probabilities of state 1 after observing partial history $s^t, p_i(1|s^t)$. 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,

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

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 and zero financial wealth.

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

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

where $n^\xi (s^t)$ stands for the number of times state $\xi \in \{1, 2\}$ has been realized at the partial history s^t . Since we assume A.0 holds with $\theta^* (1) = \frac{1}{2}$, the Strong Law of Large Numbers implies that $p_i (1 | s^t) \rightarrow \frac{1}{2} (P^{\theta^*} - a.s.)$ as $t \rightarrow \infty$, for every agent $i \in \{1, 2\}$. Therefore, both agents learn the true one-period-ahead conditional probability.

By the Kolmogorov's Extension Theorem (Shiryaev [22, Theorem 3, p. 163]), there exists a unique P_i on (S^∞, \mathcal{F}) associated to the agent's one-period-ahead conditional probability. Moreover, P_i satisfies A.1 and A.2 and agents' prior beliefs over θ have densities $f_1 (\theta) = 1$ and $f_2 (\theta) = 6 \theta (1 - \theta)$ on $(0, 1)$, respectively.¹² The likelihood ratio is

$$\frac{P_{1,t} (s)}{P_{2,t} (s)} = \frac{\int_0^1 P_t^\theta (s) d\theta}{\int_0^1 P_t^\theta (s) 6 \theta (1 - \theta) d\theta} = \frac{1}{6} \frac{\frac{\Gamma[n^1(s^t)+1] \Gamma[n^2(s^t)+1]}{\Gamma[t+2]}}{\frac{\Gamma[n^1(s^t)+2] \Gamma[n^2(s^t)+2]}{\Gamma[t+4]}} = \frac{1}{6} \frac{(t+3) (t+2)}{(n^1(s^t)+1) (n^2(s^t)+1)},$$

where Γ stands for the Gamma function.¹³ The Strong Law of Large Numbers can be applied once again to show that

$$\frac{P_{1,t} (s)}{P_{2,t} (s)} \rightarrow \frac{2}{3} = \frac{f^1 (\frac{1}{2})}{f^2 (\frac{1}{2})} \quad P^{\theta^*} - a.s.$$

It follows from (3) that portfolios converge to a fixed vector, that is

$$a_{1,t}^{\xi'} (s) \rightarrow \frac{1}{1 - \beta} y(\xi') \lambda_1 \left(\left(\lambda_1 + \lambda_2 \frac{3}{2} \right)^{-1} - 1 \right), \quad \xi' \in \{1, 2\} \quad P^{\theta^*} - a.s.$$

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,

$$c_{1,t} (s) \rightarrow \frac{\lambda_1}{\lambda_1 + \lambda_2 \frac{3}{2}} y(s_t) < \lambda_1 y(s_t),$$

$$m_t^{\xi'} (s) \rightarrow \beta \frac{1}{2} \frac{y(s_t)}{y(\xi')},$$

and thus Conjecture I holds but Conjecture II does not. The reason is that in the economy that starts with homogenous prior beliefs the agents' financial wealth is zero while in the limit economy prior beliefs are homogeneous but the agents' financial wealth is not zero. In this example limit asset prices are identical to those of an otherwise identical economy that starts with homogenous prior beliefs because logarithmic preferences make intertemporal marginal rates of substitution, and thus asset prices, independent of the wealth distribution. In general, however, asset prices

¹²That is, agent i 's prior beliefs over θ follow a Beta distribution $B(i, i)$ on $(0, 1)$, as in Morris [17].

¹³Recall that if n is an integer, then $\Gamma(n) = (n - 1)!$

do depend on the wealth distribution. In Section 6 we fully characterize the limit wealth distribution and argue that it depends critically on date 0 priors. \square

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

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

i.e., agents believe that the states of nature are independent draws from time-varying distributions. Observe that one-period-ahead conditional probabilities converge to $\frac{1}{2}$ for both agents, i.e. agents learn, and have the same entropy. That is,

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

The ratio of one-period-ahead conditional probabilities at date t after partial history s^{t-1} is a random variable, $\frac{p_{1,t}}{p_{2,t}}$, 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

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

where $x_k(s) \in \{-\sqrt{1/k}, \sqrt{1/k}\}$, $E^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = 0$ and $Var^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = E^{P^{\theta^*}}(x_k^2 | \mathcal{F}_{k-1})(s) = 1/k$. 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 [11, Proposition 4.5 (a)] that

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

and, therefore,

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

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

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

Since each agent's debt attains its so-called *natural debt limit* infinitely often, individual portfolios are highly volatile. 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 probability. The fact that the one-period-ahead-conditional probabilities converge certainly means that trade in each period becomes eventually very small. However, since the likelihood ratio of agents' beliefs fails to converge, this small trade compounds over large periods of time and so (in a sufficiently long span of time) there are wide fluctuations in the distribution of wealth. \square

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 3 one faces two difficulties that we avoid in the examples by carefully choosing preferences, individual endowments and priors. First, equilibrium portfolios are typically history dependent in a more general setup. Closed-form solutions for asset demands as in (3) are useful to tackle this difficulty but they are a particular feature derived from logarithmic 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 3. 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 5, 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' beliefs over the unknown parameters, as in the examples. In section 6, 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 7, we consider a broad class of priors satisfying A.1. We apply recent results in probability theory

to prove that the likelihood ratios converge with probability one, as in example 1. Finally, in section 8 we explain the exact role played by the assumptions that every agent knows the likelihood function generating the data and that some agent learns and we argue that it is critical that they are coupled together. We do so by providing two additional examples, each of which relax one of these assumptions, in which the likelihood ratio does not converge and consequently portfolios change infinitely often as in example 2.

5 A Recursive Approach to Pareto Optimality

In this section, we provide a recursive characterization of the set of Pareto optimal allocations providing 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. It is well known that under A.1, Bayes' rule implies that prior beliefs evolve according to

$$\mu_{i,s^t}(d\theta) = \frac{\theta(s_t) \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta)}, \quad (4)$$

where $\mu_{i,0} \in \mathcal{P}(\Delta^{K-1})$ is given at date 0.

Lemma 2 *Suppose agent i 's prior satisfies A.1. Then, for every $B \in \mathcal{F}$*

$$P_{i,s^t}(B) = \int_{\Delta^{K-1}} P_{s^t}^\theta(B) \mu_{i,s^t}(d\theta). \quad (5)$$

5.1 Pareto Optimal Allocations

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

It is well known 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

$$v^*(s_0, \mu_0, \alpha) \equiv \sup_{\{c_i\}_{i=1}^I \in Y^\infty} \sum_{i=1}^I \alpha_i E^{P_i} \left(\sum_t \beta^t u_i(c_{i,t}) \right). \quad (6)$$

Unlike the case where agents have homogeneous beliefs, the recursive characterization of PO allocations in our economy is rather tricky because belief heterogeneity makes optimal allocations history dependent. This can be seen from the following

(necessary and sufficient) first order conditions to the planner's problem:

$$\frac{\alpha_i P_{i,t}(s)}{\alpha_j P_{j,t}(s)} \frac{\frac{\partial u_i(c_{i,t}(s))}{\partial c_{i,t}}}{\frac{\partial u_j(c_{j,t}(s))}{\partial c_{j,t}}} = 1 \quad \text{for all } i, j \in \mathcal{I}, \text{ for all } t \text{ and all } s, \quad (7)$$

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

Since $\frac{\alpha_j}{\alpha_i} = \frac{\partial u_i(c_{i,0})}{\partial c_{i,0}} / \frac{\partial u_j(c_{j,0})}{\partial c_{j,0}}$, the planner distributes consumption among agents to make the ratio of marginal valuations of any two agents -which, we recall, include priors that could be subjectively held- to be constant along time. Consequently, under the optimal distribution rule of consumption, the ratio of marginal utilities, $\frac{\partial u_i(c_{i,t}(s))}{\partial c_{i,t}} / \frac{\partial u_j(c_{j,t}(s))}{\partial c_{j,t}}$, must be proportional to the likelihood ratio of the agents' priors, $P_{j,t}(s) / P_{i,t}(s)$. This ratio represents the novel margin of heterogeneity among agents considered in this paper, which we call *the \mathcal{B} -margin of heterogeneity*. The \mathcal{B} -margin is purely driven by heterogeneity in priors and its evolution determines the dynamics of the optimal distribution rule of consumption. Indeed, when all agents have the same priors the \mathcal{B} -margin remains constant along time and the optimal distribution rule of consumption is both *time and history independent*. Consequently, individual consumption depends only upon the current shock s_t (because it determines aggregate output) and the date 0 vector of welfare weights α . When agents have heterogeneous priors, instead, the \mathcal{B} -margin is *history dependent* and so is the optimal distribution rule of consumption.

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

$$\frac{\frac{\partial u_i(c_{i,t}(s))}{\partial c_{i,t}}}{\frac{\partial u_j(c_{j,t}(s))}{\partial c_{j,t}}} = \frac{\int_{\Delta_{K-1}} \theta(s_{t+1}) \dots \theta(s_{t+k}) \mu_{i,s^t}(d\theta) \frac{\partial u_i(c_{i,t+k}(s))}{\partial c_{i,t+k}}}{\int_{\Delta_{K-1}} \theta(s_{t+1}) \dots \theta(s_{t+k}) \mu_{j,s^t}(d\theta) \frac{\partial u_j(c_{j,t+k}(s))}{\partial c_{j,t+k}}},$$

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, it is sufficient that he knows the ratio of marginal utilities that the original plan induces at date t , $\frac{\partial u_i(c_{i,t}(s))}{\partial c_{i,t}} / \frac{\partial u_j(c_{j,t}(s))}{\partial c_{j,t}}$, the state of nature at date t , s_t , and the posterior beliefs, $\mu_{s^t-1}(d\theta)$, since $\mu_{i,s^t}(d\theta) = \frac{\theta(s_t) \mu_{i,s^t-1}(d\theta)}{\int_{\Delta_{K-1}} \theta(s_t) \mu_{i,s^t-1}(d\theta)}$. 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 $(\alpha_1 P_{1,t}(s), \dots, \alpha_I P_{I,t}(s), \mu_{s^t-1})$ 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, Easley and Kiefer [8]). In a multiple agent setting, instead, the planner needs to distribute consumption and because of this one needs to introduce $(\alpha_1 P_{1,t}(s), \dots, \alpha_I P_{I,t}(s))$ as an additional state variable, which can be interpreted as the date t welfare weights, $\alpha_{i,t}(s) = \alpha_i P_{i,t}(s)$. These weights evolve according to the law of motion

$$\alpha_{i,t}(s) = \alpha_{i,t-1}(s) \int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta) \quad \text{where } \alpha_{i,0}(s) = \alpha_i. \quad (9)$$

In Section 5.2 below we present a formal exposition of this result.

5.2 Recursive Characterization of PO Allocations

Given that in an environment with heterogeneous beliefs and learning 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 [16].

In Appendix B we show that v^* is the unique solution of the functional equation¹⁴

$$v(\xi, \alpha, \mu) = \max_{(c, w'(\xi'))} \sum_{i \in \mathcal{I}} \alpha_i \left\{ u_i(c_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu)(d\theta) w'_i(\xi') \right\}, \quad (10)$$

subject to

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

$$\alpha'(\xi') \equiv \arg \min_{\tilde{\alpha} \in \Delta^{I-1}} \left[v(\xi', \tilde{\alpha}, \mu'(\xi, \mu)) - \sum_{i=1}^I \tilde{\alpha}_i w'_i(\xi') \right] \geq 0 \quad \text{for all } \xi', \quad (12)$$

where

$$\mu'_i(\xi, \mu)(B) = \frac{\int_B \theta(\xi) \mu_i(d\theta)}{\int_{\Delta^{K-1}} \theta(\xi) \mu_i(d\theta)} \quad \text{for any } B \in \mathcal{B}(\Delta^{K-1}). \quad (13)$$

In the recursive dynamic program defined by (10) - (13), 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 \mathcal{B} -margin of heterogeneity. The planner takes as given (ξ, α, μ) and allocates current consumption and continuation utility levels among agents. That is, instead of allocating consumption from tomorrow on, the planner assigns to each agent the utility level associated with the corresponding continuation sequence of consumption. Indeed, the optimization problem defined in condition (12) characterizes the set of continuation utility levels

¹⁴In sections 5.2 and 6, we abuse notation and let c to be a non-negative vector and c_i its i^{th} component.

attainable at $(\xi', \mu'(\xi, \mu))$ (see Lemma 14 in Appendix B).¹⁵ The weights $\alpha'(\xi')$ that attain the minimum in (12) will then be the new weights used in selecting tomorrow's allocation.

The (normalized) law of motion for the welfare weights, $\alpha'_i(\xi, \alpha, \mu)(\xi')$, follows from the first order conditions with respect to the continuation utility levels for each individual and is given by

$$\alpha'_i(\xi, \alpha, \mu)(\xi') \equiv \frac{\alpha_i \int \theta(\xi') \mu'_i(\xi, \mu) (d\theta)}{\sum_h \alpha_h \int \theta(\xi') \mu'_h(\xi, \mu) (d\theta)}. \quad (14)$$

Observe that the normalization is harmless since optimal policy functions are homogeneous of degree zero with respect to α .

It follows by standard arguments that the corresponding consumption policy function, $c_i(\xi, \alpha)$, is the unique solution to

$$c_i(\xi, \alpha) + \sum_{h \neq i} \left(\frac{\partial u_h}{\partial c_h} \right)^{-1} \left(\frac{\alpha_i}{\alpha_h} \frac{\partial u_i(c_i(\xi, \alpha))}{\partial c_i} \right) = y(\xi). \quad (15)$$

for each $i \in \mathcal{I}$, where $\left(\frac{\partial u_h}{\partial c_h} \right)^{-1}$ denotes the inverse function of $\frac{\partial u_h}{\partial c_h}$.

Given (s_0, α_0, μ_0) , we say the policy functions (c, α') coupled with μ' generates an allocation \hat{c} if

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

for all i and all $t \geq 0$ and $s \in S^\infty$ where $\alpha_0(s) = \alpha_0$ and $\mu_{s^{-1}} = \mu_0$. The following Theorem shows 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) - (13).

Theorem 3 (The Principle of Optimality) *An allocation $(c_i^*)_{i=1}^I$ is PO given (ξ, α, μ) if and only if it is generated by the policy functions solving (10) - (13).*

¹⁵To understand condition (12) notice that the utility possibility set, i.e. the set of expected lifetime utility levels that are attainable by mean of feasible allocations, is convex, compact and contains its corresponding frontier. The frontier of a convex set can always be parametrized by supporting hyperplanes. Moreover, under our assumptions, the corresponding parameters can be restricted to lie in the unit simplex and, therefore, they can be interpreted as welfare weights. Thus, a utility level vector w is in the utility possibility set if and only if for every welfare weight α the hyperplane parametrized by α and passing through w , αw , lies below the hyperplane generated by the utility levels attained by the PO allocation corresponding to that welfare weight α , attaining the value $v(\xi, \alpha, \mu)$. This is why we must have $\alpha w \leq v(\xi, \alpha, \mu)$ for all α or, equivalently, $\min_{\tilde{\alpha}} [v(\xi, \tilde{\alpha}, \mu) - \tilde{\alpha} w] \geq 0$. See Appendix B for technical details.

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 levels which in turn can be summarized by a vector of welfare weights. Theorem 3 shows that the date zero optimal plan is consistent in the sense that the continuation plan is indeed the solution from date t onwards.

5.2.1 Discussion: An Alternative Approach

There is an alternative approach to state the dynamic program defined by (10) - (13): instead of parametrizing allocations with welfare weights, the planner chooses current feasible consumption and continuation utilities for both agents in order to maximize the utility of agent 1 subject to two restrictions: (i) the utility of agent 2 is above some prespecified level (the so-called promise keeping constraint) and (ii) continuation utility levels lie in next period utility frontier. Very importantly, this last condition implies that the corresponding value function defines the constraint set.¹⁶ Since both in our approach as well as in the alternative one the corresponding value function defines the constraint set, neither of the two dynamic programs is standard in the sense that it is not obvious that any of the corresponding operators satisfies one of Blackwell’s sufficient conditions, namely, discounting. Indeed, for any function v that defines the constraint set there might be some $a > 0$ such that $v + a$ enlarges the feasible set of choices of continuation utilities with respect to v . The key to show discounting in our approach is to restrict the set of functions to be homogeneous of degree 1 with respect to the state variables, i.e. the welfare weights, (a property that is satisfied by v^* , see Lemma 13 in Appendix B).¹⁷ Since $v + a$ is an affine linear transformation of v , the choice of current consumption is the same for v and $v + a$. In addition, homogeneity of degree 1 of the value function with respect to the welfare weights implies that w' is the optimal choice for the constraint set defined by v if and only if $w' + a$ is the optimal choice for the constraint set defined by $v + a$. This explains why homogeneity of degree 1 of the value function with respect to the welfare weights is key to show that discounting holds in our setting.¹⁸

¹⁶Since we parametrized the utility levels with their associated welfare weights, our approach amounts to replacing the promise keeping constraint by using the associated Lagrange multipliers as state variables.

¹⁷Lucas and Stokey [16] do not make this restriction and so it is unclear whether discounting holds in their approach.

¹⁸There is no obvious condition equivalent to homogeneity in the alternative approach described above and then one needs to find the solution differently. One plausible strategy would be to follow

6 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 the states (ξ, α, μ) . 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 $c_i(\xi, \alpha)$ and define the state price by¹⁹

$$M(\xi, \alpha, \mu)(\xi') = \beta \int \theta(\xi') \mu'_i(\xi, \mu) (d\theta) \frac{\partial u_1(c_1(\xi', \alpha'(\xi, \alpha, \mu)(\xi')))/\partial c_1}{\partial u_1(c_1(\xi, \alpha))/\partial c_1}, \quad (16)$$

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

$$A_i(\xi, \alpha, \mu) = c_i(\xi, \alpha) - y_i(\xi) + \sum_{\xi'} M(\xi, \alpha, \mu)(\xi') A_i(\xi', \alpha', \mu'), \quad (17)$$

where $\mu'(\xi, \mu)$ and $\alpha'(\xi, \alpha, \mu)(\xi')$ are given by (13) and (14), respectively. Note that (17) computes recursively the present discounted value of agent i 's excess demand at the PO allocation.

In Theorem 4, 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 .

Theorem 4 *Suppose A.0 and A.1 hold. Then, there is a unique continuous function A_i solving (17). 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 .*

REMARK 2: Observe that if agents have both homogeneous priors and dogmatic beliefs (i.e., $\mu_{i,0} = \mu^\theta$ for all i and for some $\theta \in \Delta^{K-1}$), it follows immediately that $\alpha_{i,t+1}(s) = \alpha_0$ and $\mu_{i,s^t} = \mu^\theta$ for all s and all $t \geq 0$ and for each agent i . Therefore, $c_{i,t}(s) = c_i(s_t, \alpha_0)$ for all s and all $t \geq 0$. Very importantly, the solution to (17) reduces to a vector in \mathbb{R}^K and, thus, $A_i(\alpha_0, \mu^\theta) \in \mathbb{R}^K$ such that $A_i(s_0, \alpha_0, \mu^\theta) = 0$ for all i .

the seminal idea pioneered by Abreu, Pearce and Stacchetti [1] and construct an alternative operator that iterates directly on the utility possibility correspondence. Then, the value function (and the corresponding policy functions) could be recovered from the frontier of the fixed point of that operator on utility correspondences.

¹⁹The 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.

6.1 The Fixed Equilibrium Portfolio Property

We say that the *fixed equilibrium portfolio (FEP hereafter)* property holds if there exists $\{a_i(1), \dots, a_i(K)\} \in \mathfrak{R}^K$ 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. We say that the *B-margin of heterogeneity vanishes on a path s* if $\alpha_t(s)$ converge on s . Additionally, the *FEP property holds asymptotically on s* if individual financial wealth, $A_{i,t}(s) = A_i(\xi, \alpha_t(s), \mu_{s^{t-1}})$, converges on s for every i and $\xi \in S$.

Judd *et al.* [14] show that the *FEP* property is always satisfied after a once-and-for all initial rebalancing when agents have homogeneous priors (the *B*-margin of heterogeneity is constant) and degenerate beliefs. Specifically, Remark 6 implies that $a_i(\xi) = A_i(\xi, \alpha_0, \mu^\theta)$ for all ξ since in their setting $\mu_{i,0} = \mu^\theta$ for some $\theta \in \Delta^{K-1}$ and for all i , therefore, the agents' financial wealth is a vector in \mathbb{R}^K in any dynamically complete markets equilibrium.

When agents have homogeneous but non-degenerate prior beliefs, the welfare weights are constant along time, $\alpha_t(s) = \alpha_0$, and the distribution of consumption is given by $c_i(\xi) = c_i(\xi, \alpha_0)$ for each i and so it remains unchanged as time and uncertainty unfold. However, the wealth distribution, $A_{i,t}(s) = A_i(\xi, \alpha_0, \mu_{s^{t-1}})$, is still history dependent because the agents' learning process make state prices history dependent. Consequently, the *FEP* property does not necessarily hold. To get a thorough understanding of why the wealth distribution changes even though the distribution of consumption does not, consider a two-agent and two-state economy where $y_i(\xi) = 1$ if $\xi = i$ and 0 otherwise. Both agents' prior belief is that state 1 is very likely (i.e. μ_0 concentrate most of its mass around $\theta(1) = 1$) while the true probability of each state is, say, $\theta^*(\xi) = 1/2$ for $\xi \in \{1, 2\}$. Thus, agent 1 is richer than agent 2 and this implies that agent 1's date-zero welfare weight, $\alpha_{1,0}(s_0, \mu_0)$, is relatively larger than $\alpha_{1,0}(s_0, \mu^{\theta^*})$ and she consumes accordingly forever. For t sufficiently large, $\mu_{s^{t-1}}$ gets close to μ^{θ^*} and the present discounted value of agent 1's endowment will not be enough to afford the fixed consumption bundle $c_1(\xi, \alpha_0)$. Consequently, she must have accumulated sufficient financial wealth to make the consumption parametrized by $\alpha_0(s_0, \mu_0)$ affordable (that is, $A_{i,t}(s) = A_i(\xi, \alpha_0, \mu_{s^{t-1}}) > 0$ for all $t > T$). The remarkable feature is that the evolution of (homogeneous) beliefs has an impact only on portfolios but not on consumption. This makes evident that Judd *et al*'s [14] results need both homogeneous as well as degenerate beliefs.

In our setting, instead, portfolios typically change as time and uncertainty unfolds because the changes in the *B*-margin of heterogeneity affects the dynamics of the wealth distribution through the evolution of the welfare weights. 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.* [14] is not robust to the introduction of this margin of heterogeneity. However, since agents observe the same data, have the true model in the support of their priors and update their priors in a Bayesian fashion, a much deeper question is whether this trading activity fades out as the \mathcal{B} -margin of heterogeneity vanishes. Our recursive approach permits to study this issue directly.

The following proposition, a direct consequence of the continuity of A_i , underscores that if a PO allocation can be decentralized through a sequence of competitive markets, the associated wealth distribution converges to a fixed vector for each ξ whenever the \mathcal{B} -margin of heterogeneity is exhausted. Consequently, asset trading reduces to the minimum.

Proposition 5 *Suppose A.0 and A.1 hold. If the \mathcal{B} -margin of heterogeneity vanishes on a path s , then the FEP property holds asymptotically on s ($P^{\theta^*} - a.s.$)*

7 Limiting Welfare Weights

In this section, we analyze how the \mathcal{B} -margin of heterogeneity evolves over time when A.1 holds and every agent has the true parameter in the support of their priors.

From condition (14) and Theorem 3, the ratio of welfare weights is

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}} \theta(s_t) \mu_{j,s^{t-1}}(d\theta)} \frac{\alpha_{i,t-1}(s)}{\alpha_{j,t-1}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{P_{i,t}(s)}{P_{j,t}(s)}, \quad (18)$$

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 that the data is generated by i.i.d. draws from a common distribution (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 Sandroni [20]. 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 both A.1 and A.2 hold, the probability distribution that generates the data is *never* absolutely continuous with respect to the agents' priors and so Sandroni's result does not apply.²⁰ Nonetheless, we show that likelihood ratios converge applying a recent result by Phillips and Ploberger [19].

²⁰Blume and Easley [4] also emphasize this point.

7.1 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 Sandroni [20] shows, 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 6 we show that for every agent i ,

$$\frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} \rightarrow \mu_{i,0}(\theta^*) \quad P^{\theta^*} - a.s. \quad (19)$$

Proposition 6 *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 (19) holds.*

In turn, Proposition 6 implies that the agent's likelihood ratios also have a finite positive limit. Indeed,

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

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

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

where, for every h , $\alpha_h^* = \alpha_{h,0} \mu_{h,0}(\theta^*)$ and $\alpha_{h,0}$ is the welfare weight defined in Theorem 4. Therefore, we obtain the following result which completely characterizes the limiting properties of the economy.

Theorem 7 *Suppose A.0 and A.1 hold. If the support of every agent's prior belief is countable and $\mu_{h,0}(\theta^*) > 0$, then every efficient allocation converges to the Pareto optimal allocation parametrized by $[\alpha_{1,0} \mu_{1,0}(\theta^*), \dots, \alpha_{I,0} \mu_{I,0}(\theta^*)]$, $P^{\theta^*} - a.s.$ Furthermore, the FEP property holds asymptotically with $a_h(\xi) = A_h(\xi, \alpha^*, \mu^{\theta^*})$ for all ξ and $h \in \mathcal{I}$.*

7.2 Uncountable Support

Now we turn to the case where the agent's prior satisfies A.1 and A.2. Since Sandroni's result does not apply, we invoke a result in Phillips and Ploberger [19,

Theorem 4.1] (stated in the appendix for completeness) to establish that under assumptions A.0, A.1 and A.2

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

where $\hat{\theta}_t$ is the Maximum Likelihood Estimator (MLE) of θ and ϕ^* is a constant depending upon θ^* that we define properly in the Appendix.

Proposition 8 *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

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

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

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

where, for every h , $\alpha_h^* = \alpha_{h,0} f_i(\theta^*)$ and $\alpha_{h,0}$ is the welfare weight defined in Theorem 4. We summarize all these results in the following theorem.

Theorem 9 *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 $[\alpha_{1,0} f_1(\theta^*), \dots, \alpha_{I,0} f_I(\theta^*)]$, $P^{\theta^*} - a.s.$ Furthermore, the FEP property holds asymptotically with $a_h(\xi) = A_h(\xi, \alpha^*, \mu^{\theta^*})$ for all ξ and $h \in \mathcal{I}$.*

7.3 Discussion

Theorems 7 and 9 argue forcefully that when the true parameter is in the support of every agent's prior belief and they know the data is generated by i.i.d. draws from a common distribution (i.e., A.1 holds), the equilibrium allocation of the economy with heterogeneous priors converges to that of an otherwise identical 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, and the date 0 welfare weights are 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 h 's welfare weight, α_h , is the planner's current valuation of an additional unit of agent h 's utility. By consistency, then, $\alpha_h \beta \int \theta(\xi') \mu'_h(\xi, \alpha, \mu)(d\theta)$ is the planner's current valuation of an additional unit of agent h 's next period utility at state ξ' . This is the economics behind the law of motion (18), 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 Sandroni [20] suggests that $P_t^{\theta^*}$ is the normalization that works. In the uncountable case, on the other hand, the work of Phillips and Ploberger [19] suggests that $P_t^{\widehat{\theta}_t}$ 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 [4] and Sandroni [20] 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 (17)) 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.

8 Persistent Trade

In this section we argue, by means of two examples, that to show the *FEP* property holds asymptotically it is necessary to assume that some agent learns the true conditional probability of the states (section 8.1) and that every agent knows the likelihood function generating the data (section 8.2).

8.1 Example 3: Dogmatic Priors

Judd *et al.* [14] show that the *FEP* property holds for economies with homogeneous and degenerate 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. These two assumptions ensures the agents learn the true conditional probability of the states. Here we show that this last property 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.

For simplicity, we assume there are two states of nature, that is $K = 2$, and 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.^{21, 22} The ratio of one-period-ahead conditional probabilities, $\frac{p_{1,t}}{p_{2,t}}$, 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 variables as follows:

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

where $E^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = 0$ and $var^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = E^{P^{\theta^*}}(x_k^2 | \mathcal{F}_{k-1})(s) = E^{P^{\theta^*}}(x_k^2) > 0$. So, the log likelihood ratio is the sum of uniformly bounded random

²¹See Blume and Easley [5] for a general analysis of optimal consumption paths in i.i.d. economies where agents have degenerate prior beliefs.

²²The assumption that priors are dogmatic is not necessary for the result to be true. The necessary condition is that no agent has the true in the support of her prior and that the limit of each agent's prior beliefs has identical entropy.

variables with zero conditional mean and conditional variance bounded away from zero. Once again, it follows by Freedman [11, 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 (18), 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 $A_i(\xi, \alpha_t(s), (\mu^{\theta_1}, \mu^{\theta_2}))$. We can conclude that the *FEP* property does not necessarily hold asymptotically.

8.2 Example 4: Different Likelihood Functions

In example 2 we show that the *FEP* property may not hold asymptotically when no agent satisfies *A.1*. To underscore the importance of assuming that *A.1* holds for *every* agent, here we 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 probability converge to the truth. The other agent, on the other hand, does not know that the data is generated by i.i.d. draws from a common distribution (i.e., he has a wrong "model" in mind). For some partial histories his one-period-ahead conditional probability is correct while for some others it is 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 $\mu_{1,s^t} \xrightarrow{w} \mu^{\theta^*}$, $P^{\theta^*} - a.s.$ By Lemma 1 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

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

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

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

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

By the Kolmogorov's Extension Theorem there exists a unique probability measure \tilde{P}^θ on (S^∞, \mathcal{F}) that coincides with $\{\tilde{P}_t^\theta\}_{t=1}^\infty$ when restricted to $\{(S^\infty, \mathcal{F}_t)\}_{t=0}^\infty$.²³

REMARK 3: If $\theta = \theta^*$, then $\tilde{P}^{\theta^*} = P^{\theta^*}$.

8.2.1 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

$$m^* \equiv \arg \min_{\varepsilon \leq m \leq 1 - \varepsilon} (\theta^* \log m + (1 - \theta^*) \log(1 - m)),$$

and let m_t^* denote the i.i.d. random variable that takes values m^* and $1 - m^*$ in states 1 and 2, respectively.²⁴ 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 data is generated by i.i.d. draws from an unknown common distribution and so his prior is

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

The prior of agent 2 is

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

and agent 2's one period ahead conditional probability is given by

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

REMARK 4: Notice that agent 2's one period-ahead probability is infinitely often bounded away from the true one period-ahead conditional probability and so he never

²³ Clearly, the family $\{\tilde{P}^\theta : \theta \in \Delta^1\}$ consist of probability measures on (S^∞, \mathcal{F}) such that for each $B \in \mathcal{F}$, $\theta \rightarrow \tilde{P}^\theta(B)$ is $\mathcal{B}(\Delta^1)$ -measurable.

²⁴ $(m^*, 1 - m^*)$ is the probability distribution that minimizes entropy among those for which the probability of state 1 is larger than ε . This property guarantees that m^* is different from θ^* eventually.

learns the true parameter. At first reading this seems to contradict Lemma 1 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 at least 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, i.e. one cannot have $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$. 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. So, it must be the case that $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$. 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.²⁵ Therefore, the set of paths where the likelihood ratio is smaller than or equal to one in finite time must have zero measure. To clinch the result notice that since the ratio of one-period-ahead conditional probabilities is bounded away from one infinitely often, the likelihood ratio exceeds any pre-specified upper bound infinitely often on the set of paths where the likelihood ratio is greater than one infinitely often.²⁶ Thus, it must diverge along some subsequence of periods, i.e. $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$.

Proposition 10 *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 $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$ P^{θ^*} - a.s.*

²⁵If 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.

²⁶To 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.

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

We consider again the economy described in Section 4 where portfolios are given by (3). Proposition 10 makes it clear that agents 1 and 2 survive. Agent 1's one-period-ahead conditional probability converge to the truth while agent 2 makes mistakes infinitely often. However, agent 2's one-period-ahead conditional probability 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 10 and the arguments above.

Proposition 11 *Suppose A.0 holds. If agent 1's prior satisfies A.1 and A.2, then, P^{θ^*} – 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 asymptotically on s .*

8.2.3 Further Remarks

In Sandroni's [20] terminology, agent 1 *eventually makes accurate next period predictions* while agent 2 does not and yet both agent survive. At a first glance, then, Proposition 11 (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 [4] either since they only consider economies where every agent's prior satisfies A.1.²⁷

²⁷They 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.

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 mis-specified models). We explore that margin in example 4 and our findings, stated in Proposition 11, 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.

9 Concluding Remarks

If agents know that the data is generated by i.i.d. draws from a common distribution 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 learns the true conditional probability of the states or (ii) some agent does not know the likelihood function generating the data or (iii) the conditional probability of the states of nature changes along time.

10 Appendix A

In this Appendix we show that (3), 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

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

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

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

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

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

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

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

It is a routine exercise to show that the PO allocation corresponding to $(\alpha_1, \alpha_2) = (\lambda_1, \lambda_2)$ can be decentralized as a competitive equilibrium with sequential markets 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

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

Thus, equilibrium portfolios at date t are

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

11 Appendix B

In this Appendix we establish the results in Sections 4 and 5.

Proof of Lemma 2. Observe first that

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

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

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

■

Before proving Theorem 3, we need some definitions and preliminary results. We begin defining the *utility possibility correspondence*. In Lemma 12 we show that the utility possibility correspondence is well-behaved and in Lemma 13 we characterize the value function of the planner's problem. Lemma 14 provides a useful representation of the utility possibility frontier. The results in these three Lemmas allow us to define the operator given by (10) - (12). In Proposition 15 we argue that the operator has a unique fixed point and in Proposition 16 we show that the value function of the planner's problem v^* is the unique fixed point of the operator.

Given the state of nature and prior beliefs at date zero, $s_0 = \xi$ and $\mu_0 \equiv (\mu_{1,0}, \dots, \mu_{I,0}) = \mu$, define the *utility possibility correspondence* by

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

Now we show that the utility possibility correspondence is well-behaved, i.e. it is compact and convex-valued.

Lemma 12 $\mathcal{U}(\xi, \mu)$ is compact and convex-valued for all (ξ, μ) .

Proof of Lemma 12. 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 $\{w^n\}$ such that $w^n \in \mathcal{U}(\xi, \mu)$ for all n and $w^n \rightarrow \bar{w} \in \mathbb{R}_+^I$. 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 $c^{n_k} \rightarrow \bar{c} \in Y^\infty$. Thus, it follows by definition that $U_i^{P_i}(c_i^{n_k}) \geq w_i^{n_k}$ for all k and for all i . Since u_i is continuous and \mathbf{C} is compact, then $U_i^{P_i}$ is continuous under the sup-norm. Thus, it follows that $U_i^{P_i}(\bar{c}_i) \geq \bar{w}_i$, for all i . Consequently, $\bar{w} \in \mathcal{U}(\xi, \mu)$ by definition and $\mathcal{U}(\xi, \mu)$ is closed. ■

The following result characterizes the value function v^* and hints how to restrict the set of functions to consider in (10)-(13).

Lemma 13 *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 .*

Proof of Lemma 13. 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 $\{(\mu_i^n)_{i=1}^I\}$ is a sequence of probability measures such that μ_i^n converges weakly to $\bar{\mu}_i \in \mathcal{P}(\Delta^{K-1})$ for all i . Given k , note that

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

converges to

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

since $P^\theta(C(s^t))$ 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, $v_k^*(\xi, \mu, \alpha) \rightarrow v^*(\xi, \mu, \alpha)$ 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,

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

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. ■

We now show that the set of PO allocations can be parametrized by welfare weights α . Indeed, it is straightforward to prove that

$$v^*(\xi, \mu, \alpha) = \sup_{w \in \mathcal{U}(\xi, \mu)} \sum_{i=1}^I \alpha_i \cdot w_i, \quad (23)$$

Notice that the maximum in (23) is attained since the problem consists in maximizing a continuous function on a set that is compact by Lemma 12.

The next result shows that the continuation utility levels lie in the utility possibility set if and only if (12) is satisfied.

Lemma 14 $w \in U(\xi, \mu)$ if and only if $w \geq 0$ and $\min_{\tilde{\alpha} \in \Delta^{I-1}} \left[v(\xi, \tilde{\alpha}, \mu) - \sum_{i=1}^I \tilde{\alpha}_i w_i \right] \geq 0$.

Proof of Lemma 14. We first show that $w \in U(\xi, \mu)$ if and only if $w \geq 0$ and $v^*(\xi, \alpha, \mu) \geq \alpha w$ for all $\alpha \in \Delta^{I-1}$. To see this, observe first that for any $w \in \mathcal{U}(\xi, \mu)$, it follows by definition (23) that $v^*(\xi, \mu, \alpha) \geq \alpha w$ for all $\alpha \in \Delta^{I-1}$.

To show the converse, suppose that $w \geq 0$ and $v^*(\xi, \mu, \alpha) \geq \alpha w$ for all $\alpha \in \Delta^{I-1}$ but $w \notin \mathcal{U}(\xi, \mu)$. This implies that $\nexists \tilde{w} \in \mathcal{U}(\xi, \mu)$ such that $\tilde{w} \geq w$. Since $\mathcal{U}(\xi, \mu)$ is convex, it follows by the separating hyperplane theorem that $\exists \eta \in \mathbb{R}_+^I / \{0\}$ such that $\eta w \geq \eta \tilde{w}$ for all $\tilde{w} \in \mathcal{U}(\xi, \mu)$. Since $\mathcal{U}(\xi, \mu)$ is closed, $\eta w > \eta \tilde{w}$ for all $\tilde{w} \in \mathcal{U}(\xi, \mu)$, where η can be normalized such that $\eta \in \Delta^{I-1}$. But then $v^*(\xi, \omega, \mu) \geq \omega w > \omega \tilde{w}$ for all $\tilde{w} \in \mathcal{U}(\xi, \mu)$. This contradicts (23).

Finally, observe that $v^*(\xi, \alpha, \mu) \geq \alpha u$ for all $\alpha \in \Delta^{I-1}$, is satisfied if and only if

$$\min_{\tilde{\alpha} \in \Delta^{I-1}} \left[v(\xi, \tilde{\alpha}, \mu) - \sum_{i=1}^I \tilde{\alpha}_i w_i \right] \geq 0.$$

■

Now we define the domain where the operator is defined. Define

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

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 $\|\cdot\|$.

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

$$(Tv)(\xi, \alpha, \mu) = \max_{(c, w'(\xi'))} \sum_{i \in \mathcal{I}} \alpha_i \left\{ u_i(c_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) w'_i(\xi') \right\} \quad (24)$$

subject to

$$\sum_{i=1}^I c_i = y(\xi) \quad \text{for all } \xi, \quad c_i \geq 0, \quad w'(\xi') \geq 0 \quad \text{for all } \xi', \quad (25)$$

$$\min_{\tilde{\alpha} \in \Delta^{I-1}} \left[v(\xi', \tilde{\alpha}, \mu'(\xi, \mu)) - \sum_{i=1}^I \tilde{\alpha}_i w'_i(\xi') \right] \geq 0 \quad \text{for all } \xi', \quad (26)$$

where

$$\mu'_i(\xi, \mu)(B) = \frac{\int_B \theta(\xi) \mu_i(d\theta)}{\int_{\Delta^{K-1}} \theta(\xi) \mu_i(d\theta)} \quad \text{for any } B \in \mathcal{B}(\Delta^{K-1}). \quad (27)$$

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 . Furthermore, Proposition 16 below shows that the value function (6) is, indeed, the unique solution to (24)-(27).

Proposition 15 *There is a unique function $v \in F_H$ solving (24)-(27) and the corresponding policy functions are continuous.*

Proof of Proposition 15. We first show that $T : F_H \rightarrow F_H$.

Suppose that $f \in F_H$. Since $u_i(c_i) \leq \max u_i(\bar{y})$ and $0 \leq w'_i(\xi') \leq \|f(\xi')\|$ 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 [8, Theorem 1]), it follows by the Maximum Theorem that $(Tf)(\xi, \alpha, \mu)$ is continuous in (α, μ) for all ξ (Easley and Kiefer [8, 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, $T : F_H \rightarrow F_H$.

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

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

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

it follows that

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

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 $(\hat{c}, \hat{w}'(\xi'))$ attain $T(f+a)$. Fix $(\xi', \mu'(\xi, \mu))$, denote $f(\tilde{\alpha}) = f(\xi', \tilde{\alpha}, \mu')$ and define

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

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

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

(iii) To consider the third case (see Regions III and IV), suppose to simplify that $I = 2$ and let $w_1 \geq a$ and $w_2 < a$. Fix w_2 , let $\tilde{\alpha} \in [0, 1]$ and define

$$\begin{aligned} U_1^a(w_2) & \equiv \{w_1 \geq 0 : f(\tilde{\alpha}, 1 - \tilde{\alpha}) + a \geq \tilde{\alpha} w_1 + (1 - \tilde{\alpha}) w_2, \quad \forall \tilde{\alpha} \in [0, 1]\} \\ & = \{w_1 \geq 0 : f(\tilde{\alpha}, 1 - \tilde{\alpha}) + (a - w_2) \geq \tilde{\alpha} (w_1 - w_2), \quad \forall \tilde{\alpha} \in [0, 1]\}. \end{aligned}$$

Define $\bar{w}_1^a(w_2) \equiv \sup U_1^a(w_2)$ and note that

$$\begin{aligned} \bar{w}_1^a(w_2) & = \min_{0 \leq \tilde{\alpha} \leq 1} \left(\frac{f(\tilde{\alpha}, 1 - \tilde{\alpha})}{\tilde{\alpha}} + \frac{(a - w_2)}{\tilde{\alpha}} \right) + w_2 \\ & = \min_{0 \leq \tilde{\alpha} \leq 1} \left(f\left(1, \frac{1}{\tilde{\alpha}} - 1\right) + \frac{(a - w_2)}{\tilde{\alpha}} \right) + w_2 \\ & = f(1, 0) + (a - w_2) + w_2 = f(1, 0) + a, \end{aligned}$$

where the second line follows from HOD 1 and the last line from the monotonicity assumption about f and $(a - w_2) > 0$. Very importantly, note that $\bar{w}_1^a(w_2)$ is independent of w_2 for all $w_2 \leq a$, i.e. $\bar{w}_1^a(w_2) = \bar{w}_1^a = f(1, 0) + a$ for all $w_2 \leq a$.

Define $w' = (f(1, 0), 0) \geq 0$ and let $\tilde{\alpha} \in \Delta^{I-1}$. If $\tilde{\alpha}_1 = 0$, then $\tilde{\alpha} \cdot w' = 0 \leq f(\tilde{\alpha})$. If $\tilde{\alpha}_1 > 0$, then

$$f(\tilde{\alpha}) = \tilde{\alpha}_1 f\left(1, \frac{\tilde{\alpha}_2}{\tilde{\alpha}_1}\right) \geq \tilde{\alpha}_1 f(1, 0) = \tilde{\alpha} \cdot w',$$

and thus $w' \in U^0$.²⁸ Finally, notice that $w \leq (\bar{w}_1^a, a) = w' + a$ and $w' \in U^0$. Consequently, we can conclude that $B = U^a$. See Figure 1 below.




Figure 1:

Notice that if (\hat{c}, \hat{w}') attain $T(f + a)$, then there exists $\tilde{w}'(\xi') \in U^0$ such that $\tilde{w}'(\xi') \geq \hat{w}'(\xi') - a$ for all ξ' . By monotonicity, $(\hat{c}, \tilde{w}'(\xi') + a)$ also attain $T(f + a)$. Observe that for any (ξ, α, μ) , it follows by definition that

$$Tf(\xi, \alpha, \mu) \geq \sum_{i=1}^I \alpha_i \{u_i(\hat{c}_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) \tilde{w}'_i(\xi')\},$$

²⁸We 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 $\arg \min \left(\frac{f(\tilde{\alpha}, 1 - \tilde{\alpha})}{\tilde{\alpha}} + \frac{a - w_2}{\tilde{\alpha}} \right) = 1$. 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

$$\bar{w}_1^a = \min \left(\frac{f(\tilde{\alpha}, 1 - \tilde{\alpha})}{\tilde{\alpha}} + \frac{a - w_2}{\tilde{\alpha}} \right) > \min \frac{f(\tilde{\alpha}, 1 - \tilde{\alpha})}{\tilde{\alpha}} + \min \frac{a - w_2}{\tilde{\alpha}} = \bar{w}_1^0 + a - w_2.$$

and thus

$$\begin{aligned}
& T(f + a)(\xi, \alpha, \mu) - Tf(\xi, \alpha, \mu) \\
& \leq \sum_{i=1}^I \alpha_i \{u_i(\widehat{c}_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) (\widetilde{w}'_i(\xi') + a)\} \\
& \quad - \sum_{i=1}^I \alpha_i \{u_i(\widehat{c}_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) \widetilde{w}'_i(\xi')\} \\
& = \beta a,
\end{aligned}$$

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 $v = Tv$. ■

Proposition 16 $v^* \in F_H$ is the unique solution to (24) - (27).

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

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

as the ξ' -continuation of c_i . Also, let

$$P_{i, \xi'}(s^t) = \frac{P_i(C(s^t))}{\int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta)},$$

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

$$\begin{aligned}
v^*(\xi, \mu, \alpha) & = \max_{u \in \mathcal{U}(\xi, \mu)} \sum_{i=1}^I \alpha_i u_i = \max_{c \in Y^\infty} \sum_{i \in \mathcal{I}} \alpha_i U_i^{P_i}(c_i) \\
& = \max_{c \in Y^\infty} \sum_{i=1}^I \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) U_i^{P_{i, \xi'}}(\xi' c_i) \right\} \\
& = \max_{\substack{c \in Y \\ w'(\xi') \in \mathcal{U}(\xi', \mu'(\xi, \mu))}} \sum_{i=1}^I \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) w'_i(\xi') \right\} \\
& = \max_{\substack{c \in Y \\ w'(\xi') \geq 0}} \sum_{i=1}^I \alpha_i \left\{ u_i(c_i(\xi)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) w'_i(\xi') \right\}
\end{aligned}$$

subject to $\left[v^*(\xi', \widetilde{\alpha}, \mu'(\xi, \mu)) - \sum_{i=1}^I \widetilde{\alpha}_i w'_i(\xi') \right] \geq 0$ for all $\widetilde{\alpha} \in \Delta^{I-1}$. 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. ■

Theorem 3 uses our previous results to conclude that the corresponding version of the principle of optimality holds in our setting.

Proof of Theorem 3. Now we show that the set of policy functions (c, α', w') solving (RPP) generates a Pareto optimal allocation. Suppose, on the contrary, that the allocation $(c_i^*)_{i=1}^I$ generated by the policy functions is not Pareto optimal. Then, there exists an alternative allocation $(\widehat{c}_i^*)_{i=1}^I$ such that

$$\begin{aligned} & \sum_{i=1}^I \alpha_i \left\{ u_i(c_i^*(\xi)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi', c_i^*) \right\} \\ & > \sum_{i=1}^I \alpha_i \left\{ u_i(\widehat{c}_i(\xi)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi', \widehat{c}_i) \right\} \\ & = v^*(\xi', \alpha', \mu'(\xi, \mu)) \end{aligned}$$

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 14 that

$$v^*(\xi', \tilde{\alpha}, \mu'(\xi, \mu)) \geq \sum_{i=1}^I \tilde{\alpha}_i U_i^{P_i, \xi'}(\xi', c_i^*)$$

for all $\tilde{\alpha} \in \Delta^{I-1}$ and all ξ' . But this contradicts that the policy functions (c, α', w') 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 (24) - (27). ■

Proof of Theorem 4. Let F be defined as before. Consider the alternative operator \widetilde{T} defined by

$$\begin{aligned} (\widetilde{T}M)(\xi, \alpha, \mu) &= (c_i(\xi, \alpha) - y_i(\xi)) u'_1(c_1(\xi, \alpha)) \\ &+ \sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu'_1(\xi, \mu) (d\theta) M(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)). \end{aligned}$$

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

$$\sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu'_1(\xi, \mu) (d\theta) M(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)), \quad (28)$$

and observe that α' and μ' are both continuous. Also, it follows by definition that $\int_{\Delta^{K-1}} \theta(\xi') \mu'_1(\xi, \mu) (d\theta)$ is continuous. Thus, the expression 28 is continuous in (ξ, α, μ) . Since M is bounded, its boundedness follows from

$$\sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_1(\xi, \mu) (d\theta) = 1.$$

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

$$0 \leq c \frac{\partial u_i(c)}{\partial c_i} \leq u_i(c) \leq u_i(\bar{y}),$$

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

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

Consequently, $(c_i(\xi, \alpha) - y_i(\xi)) \frac{\partial u_1(c_1(\xi, \alpha))}{\partial c_1}$ is uniformly bounded. Clearly, it is also continuous since the policy functions are continuous. Thus, we can conclude that $\tilde{T}M \in F$.

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

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

$$\begin{aligned} \tilde{T}(M + a)(\xi, \alpha, \mu) &= (c_i(\xi, \alpha) - y_i(\xi)) \frac{\partial u_1(c_1(\xi, \alpha))}{\partial c_1} \\ &\quad + \sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) M(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)) + \beta a. \\ &= (\tilde{T}(M)(\xi, \alpha, \mu) + \beta a. \end{aligned}$$

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

Therefor, we can apply the contraction mapping theorem to conclude that \tilde{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 (17) Notice that

$$\begin{aligned} \sum_i A_i(\xi, \alpha, \mu) &= \sum_i (c_i(\xi, \alpha) - y_i(\xi)) + \sum_i \sum_{\xi'} M(\xi, \alpha, \mu)(\xi') A_i(\xi', \alpha', \mu') \\ &= \sum_{\xi'} M(\xi, \alpha, \mu)(\xi') \sum_i A_i(\xi', \alpha', \mu'). \end{aligned} \tag{29}$$

Note that the operator defined by (29) has a unique solution as well. Since $R(\xi, \alpha, \mu) = 0$ for all (ξ, α, μ) solves (29), 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 $c_i(\xi, \alpha) = 0$ and consequently $A_i(\xi, \alpha, \mu) < 0$ for all (ξ, μ) . Define the vector-valued function g on Δ^{I-1} as follows:

$$g_i(\alpha) = \frac{\max[\alpha_i - A_i(s_0, \alpha, \mu_0), 0]}{\sum_i \max[\alpha_i - A_i(s_0, \alpha, \mu_0), 0]}, \quad (30)$$

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 Brouwer'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 (30), 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 . ■

12 Appendix C

In this Appendix we establish the results in Sections 6 and 7.

Proof of Proposition 6. 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 [20], $P^{\theta^*} - a.s.$,

$$0 < \lim_{t \rightarrow \infty} \frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} < \infty, \quad (31)$$

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 [20] that, $P^{\theta^*} - a.s.$,

$$\lim_{t \rightarrow \infty} \frac{\sum_{\theta \neq \theta^*} P_t^\theta(s) \frac{\mu_{i,0}(\theta)}{1 - \mu_{i,0}(\theta)}}{P_t^{\theta^*}(s)} = 0. \quad (32)$$

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

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

where the last equality follows by (32). ■

The following Theorem, due to Phillips and Ploberger [19, page 392], will be used in the proof of Proposition 8. Before stating the Theorem, we introduce some notation. Let $\phi^* \equiv \left(\frac{1}{\theta^*} + \frac{1}{1-\theta^*}\right)$. Let $Q_{h,t}$ be the measure defined by the following Radon Nykodim derivative

$$\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)} = \frac{\sqrt{2\pi} f^h(\theta^*)}{B_t^{1/2}(s)} e^{l_t(\hat{\theta}_t(s))}, \quad (33)$$

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

Theorem 17 (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 $\lim_{t \rightarrow \infty} B_t(\theta) \rightarrow \infty$, $P^\theta - a.s.$
- (C3) $\lim_{t \rightarrow \infty} \frac{l_t^{(2)}(\theta)}{B_t(\theta)} + 1 = 0$ $P^\theta - a.s.$
- (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

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

$\lim_{t \rightarrow \infty} w_t(\theta, \theta') = w_\infty(\theta, \theta')$ $P^{\theta^*} - a.s.$ uniformly for $\theta, \theta' \in N_\delta(\theta^*)$ and $w_\infty(\theta, \theta) = 0$.

(C5) $\lim_{t \rightarrow \infty} \hat{\theta}_t = \theta^*$, $P^{\theta^*} - a.s.$

(C6) For any $\delta > 0$ and $\omega_\delta = \{\theta : |\theta - \theta^*| \geq \delta\}$ we have

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

(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 (33), then

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

Proof of Proposition 8. We need to verify that (C.1) - (C.7) in Theorem 17 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

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

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,

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

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

(C.3) Notice that by the SLLN,

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

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, $w_t(\theta, \theta') \rightarrow w_\infty(\theta, \theta')$ a.s.

(P^{θ^*}) 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 $\widehat{\theta}_t = \frac{n_{1,t}}{t} \rightarrow \theta^* \quad P^{\theta^*} - a.s.$ 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 $\theta \in \omega_\delta$, $\sup_{x \in (\theta^* - \delta/2, \theta^* + \delta/2)} \frac{\theta^x (1-\theta)^{1-x}}{(\theta^*)^x (1-\theta^*)^{1-x}} \leq 1 - \widetilde{\delta}$. Then,

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

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

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

Proof of Proposition 10. We begin with four Lemmas that will be useful to prove the main result. Lemma 18 shows that the set of paths where $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$ has full measure. Lemma 21 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. Lemmas 22 and 23 show that the latter set also has full measure.

Lemma 18 $P - a.s.$, $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$.

Proof of Lemma 18. Consider $\Omega_1 \equiv \left\{ s : \liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1 \right\}$. For each $s \in \Omega_1$ there exists $T_2(s)$ such that

$$\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1 \quad \forall t \geq T_2(s).$$

Since, $P^{\theta^*} - a.s.$, $p_{1,t}(s) \rightarrow \theta^*(s_t)$ 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 $s \in \Omega_1$,

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

and it follows that

$$\Omega_1 \subset \left\{ s : \liminf_{T \rightarrow \infty} \prod_{t=1}^T \frac{m_t^*(s)}{p_{1,t}(s)} > 0 \right\} \quad (34)$$

On the other hand, by the SLLN for uncorrelated random variables with uniformly bounded second moments, $P^{\theta^*} - a.s.$,

$$\frac{1}{T} \sum_{t=1}^T \left(\log \left(\frac{m_t^*(s)}{p_{1,t}(s)} \right) - E^{P^{\theta^*}} \left[\log \frac{m_t^*}{p_{1,t}} \mid \mathcal{F}_{t-1} \right] \right) \rightarrow 0 \text{ as } T \rightarrow \infty.$$

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

$$\frac{1}{T} \sum_{t=1}^T E^{P^{\theta^*}} \left[\log \frac{m_t^*}{p_{1,t}} \mid \mathcal{F}_{t-1} \right] \rightarrow E^{P^{\theta^*}} \left[\log \frac{m_t^*}{\theta_t^*} \right] < 0 \text{ as } T \rightarrow \infty.$$

It follows that, $P^{\theta^*} - a.s.$,

$$\frac{1}{T} \sum_{t=1}^T \log \left(\frac{m_t^*(s)}{p_{1,t}(s)} \right) \rightarrow E^{P^{\theta^*}} \left[\log \frac{m_t^*}{\theta_t^*} \right] < 0 \text{ as } T \rightarrow \infty,$$

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

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

But then, (34) implies that Ω_1 lies in a zero measure subset of Ω , as desired. ■

We continue with two results that we will need to prove Lemma 21. The first is Levy's conditional form of the Second Borel-Cantelli Lemma which follows from a more general result due to Freedman [10, Proposition 39] and is stated without proof as Lemma 19. The second result, stated in Lemma 20, 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 an event $E \in \mathcal{F}$, 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

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

Lemma 19 (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*

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

Lemma 20 *Let $\{\Omega_t\}_{t=0}^\infty$ be a sequence of events adapted to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. Then*

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

Proof of Lemma 20. Notice that

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

where we use the convention that $\Omega_{1,t}^0 = \Omega$ to handle the case where $N = 1$.

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 $k = 1, 2, \dots$. Since $\Omega_{1,t}^0 = \Omega$, $s \in \Omega_{(t_k+1)-1} \cap \Omega_{1,(t_k+1)-1}^{1-1}$ and therefore

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

and it follows by Lemma 19 that $\sum_{t=1}^{\infty} 1_{\Omega_{t-1} \cap \Omega_{1,t}^1}(s) = +\infty \quad P^{\theta^*} - \text{a.s. } s \in \{\Omega_t \text{ i.o.}\}$.

Suppose that the result holds for $N - 1$. So, for P^{θ^*} -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

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

and it follows by Lemma 19 that $\sum_{t=1}^{\infty} 1_{\Omega_{t-N} \cap \Omega_{1,t}^N}(s) = +\infty \quad P^{\theta^*} - \text{a.s. } s \in \{\Omega_t \text{ i.o.}\}$.

That completes the induction argument and the proof. ■

Lemma 21 $P^{\theta^*} - a.s. s \in \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$.

Proof of Lemma 21. Let $s \in \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$ and $a > 1$. Since $p_{1,t}(s) \rightarrow \theta^*(s)$, there exists $T(s)$ such that for every $t \geq T(s)$, $\theta^*(s) - \frac{\varepsilon}{2} \leq p_{1,t}(s) \leq \theta^*(s) + \frac{\varepsilon}{2}$. Without loss of generality, suppose $m^* > \theta^* + \frac{\varepsilon}{2}$. Let T^a be the smallest integer such that $\left(\frac{m^*}{\theta^* + \frac{\varepsilon}{2}} \right)^{T^a} > a$. Consider the event

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

By Lemma 20 it follows that $s \in \left\{ \Omega_{1,t}^{T^a} \text{ i.o.} \right\} P^{\theta^*} - a.s. s \in \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$. Therefore, $P^{\theta^*} - a.s. s \in \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, there exists a subsequence $\{t_k\}_{k=0}^{\infty}$ such that $s \in \Omega_{1,t_k}^{T^a}$ and so

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

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

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

Since a was arbitrarily chosen, it follows that

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

as desired. ■

Lemma 22 $P^{\theta^*} - a.s. s \in \left\{ \tilde{s} : \limsup \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} \leq 1 \right\}$, there exists $T(s)$ such that $\frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \forall t \geq T(s)$

Proof of Lemma 22. Let $\Omega_1 \equiv \left\{ \tilde{s} : \limsup \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} \leq 1 \text{ and } \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$. Let $s \in \Omega_1$. Since $\Omega_1 \subset \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, then by Lemma 21

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

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

Lemma 23 $\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ i.o. P^{θ^*} - a.s.

Proof of Lemma 23. Let $\Omega_1 \equiv \left\{ s : \exists T(s) \text{ such that } \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \forall t \geq T(s) \right\}$. 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 definition of agent 2's priors,

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

and it follows that

$$\Omega_1 \subset \left\{ s : \limsup_{T \rightarrow \infty} \prod_{t=1}^T \frac{\theta_t^*(s)}{p_{1,t}(s)} < +\infty \right\} \quad (35)$$

Since A.2 implies that P^{θ^*} is not absolutely continuous with respect to P_1 , it follows by Sandroni [20, Propositions B.1 and B.2.] that

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

and so (35) implies that Ω_1 lies in a zero measure subset of Ω . It follows that, P^{θ^*} - a.s., $\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ i.o. ■

Now we conclude the proof of Proposition 10 arguing that, P -a.s., $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$. Indeed, on the one hand by Lemma 22 and Lemma 23, P^{θ^*} - a.s., $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$. On the other hand, by Lemmas 21 and 23 one concludes that, P^{θ^*} - a.s., $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$. ■

References

- [1] ABREU, D., PEARCE, D. AND STACCHETTI, E. [1990], “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring”, *Econometrica*, **58**, 1041-1063.
- [2] BLACKWELL, D., AND L. DUBINS [1962]: “Merging of Opinions with Increasing Information.” *Ann. Math. Statist.*, pp. 882-86.
- [3] 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. Durlauf. Oxford University Press, Oxford.
- [4] 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.
- [5] BLUME, L. AND D. EASLEY [2006]: “The Market Organism: Long Run Survival in Markets with Heterogeneous Traders.” Mimeo, Cornell University.
- [6] COGLEY, T. AND T. SARGENT [2008]: “The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?” *Journal of Monetary Economics*, Vol. 5, No. 3 (April), 454-476.
- [7] COGLEY, T. AND T. SARGENT [2008]: “"Diverse Beliefs, Survival, and the Market Price of Risk" ” mimeo, New York University.
- [8] EASLEY, D. AND N. M. KIEFER [1988]: “Controlling a Stochastic Process with Unknown Parameters.” *Econometrica*, Vol. 56, No. 5, 1045-1064.
- [9] ESPINO, E. AND T. HINTERMAIER. Forthcoming, 2008. “Asset Trading Volume in a Production Economy.” *Economic Theory*.
- [10] 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.
- [11] FREEDMAN, D. A. [1975]: “On Tail Probabilities for Martingales.” *The Annals of Probability*, vol. 3, No. 1, 100-118.
- [12] HARRIS, M. AND A. RAVIV [1991]: “Differences of Opinion Make a Horserace.” *Review of Financial Studies*, vol. 6 (3), 473-506.

- [13] 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.
- [14] 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.
- [15] LUCAS, R. E. JR. [1978]: “Asset Prices in an Exchange Economy.” *Econometrica*, 46, 1429-1445.
- [16] LUCAS, R. AND N. STOKEY [1984]: “Optimal Growth with Many Consumers.” *Journal of Economic Theory*, 32, 139-171.
- [17] MORRIS, S. (1996): “Speculative Investor Behavior and Learning.” *Quarterly Journal of Economics*, 111, 1111-1133.
- [18] NEGISHI, T. [1960]: “Welfare Economics and the Existence of an Equilibrium for a Competitive Economy.” *Metroeconomica*, 7, 92-97.
- [19] PHILLIPS, P. C. B., AND W. PLOBERGER (1996): “An Asymptotic Theory of Bayesian Inference for Time Series.” *Econometrica*, 64(2), 381–412.
- [20] SANDRONI, A. (2000): “Do Markets Favor Agents Able to Make Accurate Predictions?” *Econometrica*, 68(6), 1303–42.
- [21] SAVAGE, L. (1954): *On the Foundations of Statistics*. Wiley, New York, NY.
- [22] SHIRYAEV, A. (1991): *Probability*. Second Edition, New York: Springer Verlag.
- [23] SCHWARTZ, L. (1965): “On Bayes Procedures.” *Probability Theory and Related Fields*, 4(1), 10-26.