Accepted Manuscript

The dynamics of efficient asset trading with heterogeneous beliefs

Pablo F. Beker, Emilio Espino

PII: S0022-0531(10)00119-5
DOI: [10.1016/j.jet.2010.08.007](https://doi.org/10.1016/j.jet.2010.08.007)
Reference: YJETH 3853

To appear in: *Journal of Economic Theory*

Received date: 23 February 2009
Revised date: 29 May 2010
Accepted date: 30 May 2010

Please cite this article in press as: P.F. Beker, E. Espino, The dynamics of efficient asset trading with heterogeneous beliefs, J. Econ. Theory (2010), doi:10.1016/j.jet.2010.08.007

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

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: August 31, 2010

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 needs some qualifications. We consider an infinite horizon stochastic endowment economy populated by many Bayesian agents with heterogeneous priors over the stochastic process of the states of nature. Our approach hinges on studying the portfolios that decentralize Pareto optimal allocations. Since these allocations are typically history dependent, we propose a methodology to provide a complete recursive characterization when agents believe 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 true process consists of i.i.d. draws from a common distribution and 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. We are also very grateful to one of the referees for very detailed comments that greatly improved the paper and the proof of Proposition 16. 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. Surprisingly, in spite of the emphasis given to efficiency very little is known about the implications of belief heterogeneity on dynamically complete markets. There are some notable exceptions. Sandroni [19] and Blume and Easley [4] provide an analysis of the asymptotic properties of consumption. Cogley and Sargent [8], [9] focus on asset prices. Our paper, instead, focuses 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 homogenous priors.¹ Judd et al. [14] considered a stationary Markovian economy where agents have homogeneous priors but different attitudes towards risk and individual endowments. They show that each investor's equilibrium portfolios 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 stochastic endowment economy populated by infinitely-lived agents who 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. We begin with two examples of dynamically complete markets equilibria that illustrate that the conventional wisdom that belief heterogeneity causes significant trade needs some qualifications. In example 1, the 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 to the truth 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, heterogeneity of priors may generate significant fluctuations in the wealth distribution and the corresponding portfolios in the long run. These examples underscore that the long run behavior of asset trading depends on the nature of the heterogeneity of priors.

This paper establishes the link between the wealth distribution (and the corresponding portfolios) of any dynamically complete markets equilibria and the ratio of the agents' priors about partial histories, i.e. the likelihood ratio of the agents' priors. We argue that belief heterogeneity creates persistent asset trading if and only

¹To avoid any confusion, we use the following terminology. By a prior, we refer to the subjective unconditional probability distribution over infinite sequences of 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 belief.

if the likelihood ratio does not converge. This characterization 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 [18] and references therein.) We consider different classes of priors and ask when the likelihood ratio converges and when it does not.

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 one agent's belief contains the true distribution of the states of nature, as in example 1.

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 [11] that hinges on studying the portfolio that decentralizes a Pareto optimal allocation as a competitive equilibrium. The difficulty is that belief heterogeneity makes Pareto optimal allocations history dependent because optimality requires that the ratio of marginal utilities at any partial history must be inversely proportional to the ratio of the agents' priors about that partial history. 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 of consumption and, consequently, the evolution of the wealth distribution in any dynamically complete markets equilibrium.

In order to obtain a recursive characterization of the set of Pareto optimal allocations in our history dependent framework, we use a strategy similar to Lucas and Stokey's [15].² We argue that the sequential formulation of the planner's problem is equivalent to a recursive dynamic program where the planner, who takes the current state, the agents' beliefs and 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 the transfers needed to decentralize a Pareto optimal allocation as a dynamically complete market equilibrium. Finally, we pin down the welfare weights that makes those transfers equal to zero given the state of nature and agents' beliefs at date 0.

This allows us to establish that the financial wealth distribution (and the corre-

²Lucas and Stokey [15] characterize recursively optimal programs in a deterministic setting where recursive preferences induce the dependence upon histories.

sponding portfolios) converges if both the $\mathcal{B}$ -margin converges and the agents' beliefs over the parameters become homogeneous. When the true distribution is in the support of the agents' priors, the well-known consistency property of Bayesian learning implies that the agents' posterior belief become homogeneous with probability one. To get a thorough understanding of the limiting behavior of portfolios, it remains to explain the asymptotic behavior of the $\mathcal{B}$ -margin. The convergence (with probability one) of the $\mathcal{B}$ -margin follows from Sandroni [19] when the support of the agents' beliefs over the parameters is a countable set and from Phillips and Ploberger [18] when the agents' beliefs have a positive and continuous density. The important message here is that the heterogeneity of priors by itself can generate changes in portfolios but these changes necessarily vanish. Furthermore, we show that portfolios converge to those of a rational expectations equilibrium of an economy where the wealth distribution is determined by the densities of their beliefs evaluated at the true parameter and the date 0 welfare weight.³

To conclude 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. We analyze the exact role played by these assumptions on priors and we argue that it is critical that they are coupled together. In each example we relax one of these assumptions and show that the $\mathcal{B}$ -margin does not converge 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 since they do not have the true parameter in the support of their beliefs, no agent learns. We assume that their beliefs are such that the associated one-period-ahead conditional probabilities have identical entropy, a condition that ensures that the $\mathcal{B}$ -margin fluctuates infinitely often between zero and infinity and, consequently, portfolios fluctuate infinitely often as well. 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 assume only one agent does not know that 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 $\mathcal{B}$ -margin, and so the equilibrium portfolios, fails to converge with probability one.

Our paper relates to the literature on wealth dynamics in general equilibrium.

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.

Sandroni [19] shows that 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.^{4,5} 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 beliefs over the probability of the states of nature have a continuous and positive density.⁶ This is precisely the case that Blume and Easley [4] consider. They prove that among Bayesian learners who have the truth in the support of their beliefs, only those with the lowest dimensional support can have positive consumption in the long run. Cogley and Sargent [8] address the quantitative relevance of belief heterogeneity and learning to explain the market price of risk but they are silent about trading volume.

Our results characterizing the portfolios that support a Pareto optimal allocation are a novel contribution since neither Sandroni [19] 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 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 [19] 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, (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.

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

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

⁵He also considers some cases in which the true distribution is not absolute continuous with respect to any agent's prior. He shows that the entropy of priors determines survival and, therefore, an agent who persistently makes wrong predictions vanishes in the presence of a learner.

⁶In 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.

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. 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 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 belief but actually for all parameters in the support of her belief.⁷ Example 4 also contributes to this literature since it is, to the best of our knowledge, the first of its kind. Since it does not require market incompleteness, it constitutes a robust argument against the market selection hypothesis.

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

2 The Model

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

2.1 The Environment

Time is discrete and indexed by $t = 0, 1, 2, \dots$. The set of possible states of nature is $S \equiv \{1, \dots, K\}$. The state of nature at date 0 is known and denoted by $s_0 \in S$. The set of partial histories up to date $t \geq 1$, S^t , is the t -Cartesian product of S with typical element $s^t = (s_1, \dots, s_t)$. S^∞ is the set of infinite sequences of the states of nature and $s = (s_1, s_2, \dots)$, called a path, is a typical element. For every partial history s^t , $t \geq 1$, a *cylinder* with base on s^t is the set $C(s^t) \equiv \{\tilde{s} \in S^\infty : \tilde{s} = (s^t, \tilde{s}_{t+1}, \dots)\}$ of all paths whose t initial elements coincide with s^t . Let $\mathcal{F}_t$ be the σ -algebra that consists of all finite unions of the sets $C(s^t)$. The σ -algebras $\mathcal{F}_t$ define a filtration $\mathcal{F}_0 \subset \dots \subset \mathcal{F}_t \subset \dots \subset \mathcal{F}$ where $\mathcal{F}_0 \equiv \{\emptyset, S^\infty\}$ is the trivial σ -algebra and

⁷Blume 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. This 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.

$\mathcal{F}$ is the σ -algebra generated by the algebra $\bigcup_{t=1}^{\infty} \mathcal{F}_t$.

For any probability measure $P : \mathcal{F} \rightarrow [0, 1]$ on $(S^\infty, \mathcal{F})$, $P_{s^t} : \mathcal{F} \rightarrow [0, 1]$ denotes its posterior distribution after observing s^t .⁸ Let $P_t(s)$ be the probability of the partial history s^t , i.e. the $\mathcal{F}_t$ -measurable function defined by $P_t(s) \equiv P(C(s^t))$ for all $t \geq 1$ and $P_0 \equiv 1$. Let p_t be the $\mathcal{F}_t$ -measurable function defined by $p_t(s) \equiv \frac{P_t(s)}{P_{t-1}(s)}$.⁹ That is, given the partial history s^{t-1} up to date $t-1$, p_t is the one-period-ahead conditional probability of the states at date t and $p(\xi | s^{t-1})$ denotes its realization at $s_t = \xi$ after the partial history s^{t-1} .¹⁰ For any random variable $x : S^\infty \rightarrow \mathbb{R}$, $E^P(x)$ denotes its mathematical expectation with respect to P . P^θ denotes the probability measure on $(S^\infty, \mathcal{F})$ uniquely induced by i.i.d. draws from a common distribution $\theta : 2^K \rightarrow [0, 1]$, where $\theta(\xi) > 0$ for all $\xi \in S$. Let Δ^{K-1} be the $K-1$ dimensional unit simplex in $\mathbb{R}^K$, $\mathcal{B}(\Delta^{K-1})$ be its Borel sets and $\mathcal{P}(\Delta^{K-1})$ be the set of probability measures on $(\Delta^{K-1}, \mathcal{B}(\Delta^{K-1}))$.

For the examples in sections 3, 7 and for the characterization of the portfolio dynamics in section 6, we need to be explicit about the true data generating process (henceforth, dgp). We will assume that

A0 The true dgp is given by P^{θ^*} for some $\theta^* \gg 0$.

2.2 The Economy

There is a single perishable consumption good every period. The economy is populated by I (types of) infinitely-lived agents where $i \in \mathcal{I} = \{1, \dots, I\}$ denotes an agent's name. A consumption plan is a sequence of functions $\{c_t\}_{t=0}^\infty$ such that $c_0 \in \mathbb{R}_+$ and $c_t : S^\infty \rightarrow \mathbb{R}_+$ is $\mathcal{F}_t$ -measurable for all $t \geq 1$ and $\sup_{(t,s)} c_t(s) < \infty$. Given s_0 , the agent's consumption set, $\mathbb{C}(s_0)$, is the set of all consumption plans.

2.2.1 Beliefs

P_i is the probability measure on $(S^\infty, \mathcal{F})$ that represents agent i 's prior. Throughout this paper, we assume that each agent i assigns positive probability to every partial history s^t ; i.e., $P_i(C(s^t)) > 0$ for all s^t . We say that agent i believes the dgp consists of i.i.d. draws from a common distribution if for every event $A \in \mathcal{F}$

$$P_i(A) = \int_{\Delta^{K-1}} P^\theta(A) \mu_{i,0}(d\theta), \quad (1)$$

⁸Formally, $P_{s^t}(A) \equiv \frac{P(A_{s^t})}{P(C(s^t))}$ for every $A \in \mathcal{F}$, where $A_{s^t} \equiv \{\tilde{s} \in S^\infty : \tilde{s} = (s^t, s'), s' \in A\}$. If $P(C(s^t)) = 0$, then P_{s^t} can be defined as an arbitrary probability measure.

⁹If $P_{t-1}(s) = 0$, then $p_t(s)$ can be arbitrarily defined.

¹⁰Throughout this paper we use the convention that probability measures over infinite paths are denoted by upper case letters, say P , and its associated one-period-ahead conditional probability are denoted by a lower case letter, say p . The same applies to P_i and p_i .

where $\mu_{i,0} \in \mathcal{P}(\Delta^{K-1})$ is agent i 's *belief* over the unknown parameters. Note that $\mu_{i,0}$ fully characterizes the prior of an agent who believe the true dgp consists of i.i.d. draws from a common distribution.

The following assumption will be used to obtain the recursive characterization of equilibrium portfolios.

A1 Agent i believes the true dgp consists of i.i.d. draws from a common distribution.

We want to emphasize that when assumption A1 holds for agents i and j , these agents agree that the dgp consists of i.i.d. draws from a common distribution but they may disagree about the distribution itself (i.e., $\mu_{i,0} \neq \mu_{j,0}$).

REMARK 1: It is ubiquitous in the learning literature related to asset pricing to assume both (i) every agent knows the dgp and (ii) some agent learns the true conditional probability of the states. In our setting, when A0 holds the latter is guaranteed by strengthening A1 to require that the true parameter, θ^* , is in the support of some agent's prior.¹¹ The case where these assumptions hold for every agent is considered in section 6. Section 7 deals with the cases in which either (i) or (ii) does not hold.

REMARK 2: All our results depending on A0 or A1 hold true if instead of assuming an i.i.d process, we assume a time-homogeneous Markov process of any finite order. The i.i.d. assumption is made to ease the exposition.

The following assumptions, when coupled with A1, impose more structure on the agent's prior.¹²

A2.a $\mu_{i,0}$ has countable support and $\mu_{i,0}(\theta^*) > 0$.

A2.b $\mu_{i,0}$ has a density $f_{i,0}$ that is continuous at θ^* with $f_{i,0}(\theta^*) > 0$.

We say that agent i is *dogmatic* if his belief is a point mass probability measure on some θ ; i.e., $\mu_{i,0} = \mu_i^\theta$ where $\theta \in \Delta^{K-1}$ and $\mu_i^\theta : \mathcal{B}(\Delta^{K-1}) \rightarrow [0, 1]$ is given by

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

2.2.2 Preferences

Agents' preferences have a subjective expected utility representation that is time separable with identical discount factors; i.e., for every $c_i \in \mathbb{C}(s_0)$ her preferences are

¹¹The finiteness of S is essential for this learning result.

¹²We adopt the convention of writing $\mu_{i,0}(\{\theta\})$ as $\mu_{i,0}(\theta)$ for all θ .

represented by

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

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

2.2.3 Endowments

Agent i 's endowment at date t is a time-homogeneous function of the state of nature, that is $y_i(s_t) > 0$ for all s_t and the aggregate endowment is $y(s_t) \equiv \sum_{i=1}^I y_i(s_t) \leq \bar{y} < \infty$. An allocation $\{c_i\}_{i=1}^I \in \mathbb{C}(s_0)^I$ is *feasible* if $\sum_{i=1}^I c_{i,t}(s) \leq y(s_t)$ for all s . Let $Y^\infty(s_0)$ be the set of feasible allocations given s_0 . An allocation $\{c_i^*\}_{i=1}^I \in Y^\infty(s_0)$ is *Pareto optimal (PO)* if there is no alternative allocation $\{\hat{c}_i\}_{i=1}^I \in Y^\infty(s_0)$ such that $U_i^{P_i}(\hat{c}_i) > U_i^{P_i}(c_i^*)$ for all i .

3 Heterogeneous Priors and Portfolios: Examples

The purpose of this section is twofold (i) to illustrate the main results of the paper by means of some simple examples and (ii) to explain the difficulties one faces to prove these results in the general setup described in section 2. We consider a dynamically complete market equilibrium and analyze the dynamics of the corresponding portfolios. In order to solve for equilibrium portfolios, we proceed as follows. First, we characterize the set of PO allocations. Then, we construct the date 0 transfers needed to decentralize a PO allocation as a competitive equilibrium with transfers. Finally, we implement the Negishi's approach to single out the competitive equilibrium allocation and its associated portfolio.

For simplicity we assume that there are two states each period ($S = \{1, 2\}$) and the dgp consists of i.i.d. draws from a common distribution θ^* where $\theta^*(1) = \theta^*(2) = \frac{1}{2}$. There are two agents ($I = \{1, 2\}$), $u_i(c) = \ln c$ and $y_i(\xi) = \lambda_i y(\xi) > 0$ for all ξ where $\lambda_1 + \lambda_2 = 1$.

We assume that a complete set of Arrow securities in zero net supply can be traded at each date. Arrow security ξ' issued at date t on path s pays 1 unit of the consumption good at $t + 1$ if $s_{t+1} = \xi'$ and 0 otherwise. The price of Arrow security ξ' and agent i 's holdings at date t after partial history s^t are denoted by $q_t^{\xi'}(s)$ and $a_{i,t}^{\xi'}(s)$, respectively. We assume the agents have no endowment of Arrow securities, i.e. they have zero financial wealth at date 0. Agent i 's budget constraint at date t is

$$c_{i,t}(s) + \sum_{\xi'=1}^2 q_t^{\xi'}(s) a_{i,t}^{\xi'}(s) \leq y_i(s_t) + a_{i,t-1}^{s_t}(s).$$

REMARK 3: In Section 2 we assumed that the range of utility functions was $\mathfrak{R}_+$. This lower bound on utility will be used in the proofs of Theorems 2 and 3. We have verified the conclusions of these Theorems directly for all of the examples in this section and in section 7.

Both in these examples as well as in the more general formulation we analyze later, we follow a three-step procedure to characterize the dynamics of portfolios that sustain the competitive equilibrium allocation.

- STEP 1: *Characterizes the solution to the planner's problem.*

In this setting, the set of PO allocations can be parametrized by welfare weights $\alpha \in \Delta^1$ since they solve the planner's problem

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

From the FOC to the planner's problem one can see that

$$\frac{\alpha_1 P_{1,t}(s)}{\alpha_2 P_{2,t}(s)} \frac{\frac{\partial u_1(c_{1,t}(s))}{\partial c_{1,t}}}{\frac{\partial u_2(c_{2,t}(s))}{\partial c_{2,t}}} = 1 \quad \text{for all } t \text{ and } s, \quad (2)$$

$$c_{1,t}(s) + c_{2,t}(s) = y(s_t). \quad (3)$$

and one concludes that consumption at date t depends only on the welfare weights, the likelihood ratio of the agents' priors through date t , $\frac{P_{2,t}}{P_{1,t}}$, and total output at date t . Indeed, in Appendix A we show that the optimal consumption allocation is

$$c_{i,t}(s)(\alpha) = \frac{\alpha_i P_{i,t}(s)}{\alpha_1 P_{1,t}(s) + \alpha_2 P_{2,t}(s)} y_t(s) \quad \text{for all } i, t \text{ and } s. \quad (4)$$

- STEP 2: *Finds the welfare weights that makes date 0 transfers equal to zero and argues that the portfolios that sustain the CE allocation have two properties: (i) they only depend on the welfare weights, the current state and the likelihood ratio (ii) they are continuous in the likelihood ratio.*

Note that for each α , the date 0 transfer is

$$A_{i,0}(\alpha) = E^{P_i} \left(\sum_{t=0}^{\infty} m_{i,t}(\alpha) (c_{i,t}(\alpha) - y_{i,t}) \right),$$

where $m_{i,t}(s)(\alpha) = \beta^t \frac{c_{i,0}(s)(\alpha)}{c_{i,t}(s)(\alpha)}$ is the stochastic discount factor. In Appendix A we show that

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

and so the CE allocation is the PO allocation parametrized by $(\alpha_1, \alpha_2) = (\lambda_1, \lambda_2)$. For this specification of preferences and endowments, properties (i) and (ii) are easy to verify since portfolios have closed-form representation. In Appendix A we show that equilibrium consumption, portfolios and security prices are

$$\begin{aligned} c_{i,t}(s) &= \left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \right)^{-1} \lambda_i y(s_t), \\ a_{i,t}^{\xi'}(s) &= \frac{1}{1-\beta} y(\xi') \lambda_i \left(\left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \frac{p_j(\xi'|s^t)}{p_i(\xi'|s^t)} \right)^{-1} - 1 \right), \quad \xi' \in \{1, 2\}, \quad (5) \\ q_t^{\xi'}(s) &= \beta p_i(\xi' | s^t) \frac{\left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \right)^{-1} y(s_t)}{\left(\lambda_i + \lambda_j \frac{P_{j,t}(s)}{P_{i,t}(s)} \frac{p_j(\xi'|s^t)}{p_i(\xi'|s^t)} \right)^{-1} y(\xi')}. \end{aligned}$$

It follows from (5) that portfolios converge if and only if the likelihood ratio converges. Thus, changes in portfolios are purely determined by the heterogeneity of priors. The margin of heterogeneity described by the likelihood ratios changes as time and uncertainty unfold. Consequently, (5) 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.

- STEP 3: *Analyzes the evolution of the likelihood ratio (and portfolios).*

This last step is accomplished in sections 3.1 - 3.3 for different specifications of priors. It is key to notice that the likelihood ratio at date t , $\frac{P_{j,t}}{P_{i,t}}$, can be written as the product of t terms each of which is the ratio of the one-period-ahead conditional probabilities, $\frac{p_j(s_k | s^{k-1})}{p_i(s_k | s^{k-1})}$, at dates $k = 1, \dots, t$. In what follows, we consider the class of priors for which the agent end up learning and we show that within that class, the likelihood ratio converges for some priors but does not converge for others. We first provide the details of the benchmark case with homogeneous beliefs.

3.1 Benchmark Case: Homogeneous Priors

We begin assuming that the one-period-ahead conditional probabilities of state 1 after observing partial history s^t , $p_i(1 | s^t)$, is the same for all agents. Then, $\frac{P_{j,t}(s)}{P_{i,t}(s)} = 1$ 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,

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

3.2 Example 1: Heterogeneous Priors I

We consider here the case in which 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 ξ has been realized at the partial history s^t . Since we assume A0 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 . Therefore, both agents learn the true one-period-ahead conditional probability.

By Kolmogorov's Extension Theorem (Shiryaev [20, Theorem 3, p. 163]), there exists a unique P_i on $(S^\infty, \mathcal{F})$ associated to agent i 's one-period-ahead conditional probability. Moreover, for every i , P_i satisfies A1 and A2.b and the agents' beliefs over θ have densities $f_{1,0}(\theta) = 1$ and $f_{2,0}(\theta) = 6 \theta(1)\theta(2)$ on Δ^1 , respectively.¹³ The likelihood ratio is

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

where Γ stands for the Gamma function.¹⁴ The Strong Law of Large Numbers can be applied once again to show that

$$\frac{P_{1,t}(s)}{P_{2,t}(s)} \rightarrow \frac{2}{3} = \frac{f_{1,0}(\frac{1}{2})}{f_{2,0}(\frac{1}{2})} \quad P^{\theta^*} - a.s.$$

It follows from (5) 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

¹³That is, agent i 's prior beliefs over θ follow a Beta distribution $B(i, i)$ on Δ^1 , as in Morris [16].

¹⁴Recall that if n is an integer, then $\Gamma(n) = (n-1)!$

homogeneous priors. Indeed,

$$\begin{aligned} c_{1,t}(s) &\rightarrow \frac{\lambda_1}{\lambda_1 + \lambda_2} y(s_t) < \lambda_1 y(s_t), \\ q_t^{\xi'}(s) &\rightarrow \beta \frac{1}{2} \frac{y(s_t)}{y(\xi')}. \end{aligned}$$

The reason is that in the economy that starts with homogenous beliefs the agents' financial wealth is zero while in the limit economy 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 beliefs because logarithmic preferences make intertemporal marginal rates of substitution, and thus asset prices, independent of the wealth distribution. In general, however, asset prices 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 the likelihood ratio might not converge in spite that agents learn.

3.3 Example 2: Heterogeneous Priors II

Now, consider the case in which

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

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

$$E^{P^{\theta^*}}(\log p_{1,t+1} | \mathcal{F}_t)(s) = E^{P^{\theta^*}}(\log p_{2,t+1} | \mathcal{F}_t)(s). \quad (6)$$

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 [13, Proposition 4.5 (a)] that

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

and, therefore,

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

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

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

Since each agent's debt attains his natural debt limit infinitely often, individual portfolios are highly volatile. A priori, this result 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.¹⁵□

3.4 The Main Lesson and Challenges Ahead

Whether portfolios converge depends on the convergence of the likelihood ratio. In these two examples, the agents learn the true one-period-ahead conditional probability. Example 2, however, makes evident that learning does not guarantee the convergence of the likelihood ratio. Why does learning ensure the convergence of the likelihood ratio in example 1 but not in example 2? The main difference is that while in example 1 both A1 and A2.b hold, neither of these assumptions hold in example 2. Those two assumptions together not only imply that the true dgp, P^{θ^*} , is in the

¹⁵It is worth mentioning that even though this example is generic in θ^* , condition (6) makes it non-generic in the space of beliefs. See example 4 below for an example in which not only some agent does not know that the dgp consists in i.i.d. draws from a common distribution (her prior does not satisfy A1) and portfolios do not converge but also it is robust in the space of beliefs.

support of the agents' prior but also that the probability that both agents assign to long histories converges to zero at a comparable speed. The latter is guaranteed by the continuity of the density at the true parameter.

To apply these three-step procedure to the more general setup described in section 2, one faces two difficulties that we avoid in the examples by carefully choosing preferences, individual endowments and priors. First, equilibrium portfolios do not always have closed-form representation. Closed-form representation for equilibrium asset holdings as in (5) was very useful to establish properties (i) and (ii) in step 2 described in section 3. However, they are a particular feature derived from logarithmic preferences and constant individual endowment shares. Consequently, one needs to use a different approach to argue that properties (i) and (ii) hold. 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 chose.

We tackle these difficulties as follows. We begin with a recursive characterization of PO allocations parametrized by welfare weights under the assumption that A1 holds. In section 4, we use this characterization to show that the evolution of any PO 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 step 1 of the examples. In section 5, we compute recursively the transfers needed to decentralize each PO allocation as a competitive equilibrium with zero transfers. Afterwards, we tackle the difficulties associated with the lack of closed form representation for the likelihood ratios. In section 6, we consider a broad class of priors satisfying A1 and A.2 (a or b). We apply recent results in probability theory to prove that the likelihood ratios converge with probability one, as in step 3 of example 1. Finally, in section 7 we explain the exact role played by assumptions A1 and A2 (a or b) and we argue that it is critical that they are coupled together for every agent. We do so by providing two additional examples, each of which relax one of these assumptions for at least one agent, in which the likelihood ratio does not converge and consequently portfolios change infinitely often as in example 2.

4 A Recursive Approach to Pareto Optimality

This section provides a recursive characterization of the set of PO allocations and a version of the Principle of Optimality for economies with heterogeneous priors.

In this section we assume that A1 holds for every agent. It is well known that

under A1, Bayes' rule implies that for all $t \geq 1$, *posterior beliefs* evolve according to

$$\mu_{i,s^t}(d\theta) = \frac{\theta(s_t) \mu_{i,s^{t-1}}(d\theta)}{\int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta)}, \quad (7)$$

where $\mu_{i,s^0} = \mu_{i,0}$.

Lemma 1 *Suppose agent i 's prior satisfies A1. Then, for every $B \in \mathcal{F}$*

$$P_{i,s^t}(B) = \int_{\Delta^{K-1}} P^\theta(B) \mu_{i,s^t}(d\theta). \quad (8)$$

4.1 Pareto Optimal Allocations

Let $\mu_0 = (\mu_{0,1}, \dots, \mu_{0,I})$. 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

$$v^*(s_0, \mu_0, \alpha) \equiv \sup_{\{c_i\}_{i=1}^I \in Y^\infty(s_0)} \sum_{i=1}^I \alpha_i E^{P_i} \left(\sum_t \beta^t u_i(c_{i,t}) \right). \quad (9)$$

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:

$$\frac{\alpha_i P_{i,t}(s)}{\alpha_j P_{j,t}(s)} \frac{\frac{\partial u_i(c_{i,t}(s))}{\partial c_{i,t}}}{\frac{\partial u_j(c_{j,t}(s))}{\partial c_{j,t}}} = 1 \quad \text{for all } i, j \in \mathcal{I}, \text{ for all } t \text{ and all } s, \quad (10)$$

$$\sum_{i=1}^I c_{i,t}(s) = y(s). \quad (11)$$

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 at date t , 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* at date t . 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. Since (8) implies that condition (10) is equivalent to

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

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, $\mu_{i,s^{t-1}}(d\theta)$. Moreover, since the ratio of marginal utilities at date t equals the $\mathcal{B}$ -margin times the ratio of welfare weights at date 0, $\frac{\alpha_j P_{j,t}(s)}{\alpha_i P_{i,t}(s)}$, 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' posterior 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 [10]). 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

$$\alpha_{i,t}(s) \propto \alpha_{i,t-1}(s) \int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta) \quad \text{where } \alpha_{i,0}(s) = \alpha_i. \quad (12)$$

In Section 4.2 below we present a formal exposition of this result.

4.2 Recursive Characterization of PO Allocations

Given that PO allocations are typically history dependent, standard recursive methods cannot be applied. We tackle this issue by adapting the method developed by Lucas and Stokey [15].

Let $\mu = (\mu_1, \dots, \mu_I)$. In Appendix B we show that v^* solves the functional equa-

tion¹⁶

$$v(\xi, \alpha, \mu) = \max_{(c, w'(\xi')) \in \mathbb{R}_+^I \times \mathbb{R}_+^I} \sum_{i=1}^I \alpha_i \left\{ u_i(c_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu)(d\theta) w'_i(\xi') \right\}, \quad (13)$$

subject to

$$\sum_{i=1}^I c_i = y(\xi) \quad \text{for all } \xi, \quad (14)$$

$$\alpha'(\xi') \equiv \arg \min_{\tilde{\alpha} \in \Delta^{I-1}} \left[v(\xi', \tilde{\alpha}, \mu'(\xi, \mu)) - \sum_{i=1}^I \tilde{\alpha}_i w'_i(\xi') \right] \geq 0 \quad \text{for all } \xi', \quad (15)$$

where

$$\mu'_i(\xi, \mu)(B) = \frac{\int_B \theta(\xi) \mu_i(d\theta)}{\int_{\Delta^{K-1}} \theta(\xi) \mu_i(d\theta)} \quad \text{for any } B \in \mathcal{B}(\Delta^{K-1}). \quad (16)$$

In the recursive dynamic program defined by (13) - (16), the current state, ξ , captures the impact of changes in aggregate output while (α, μ) summarizes and isolates the $\mathcal{B}$ -margin of heterogeneity and its law of motion. 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 (15) characterizes the set of continuation utility levels attainable at $(\xi', \mu'(\xi, \mu))$ (see Lemma 14 in Appendix B).¹⁷ The weights $\alpha'(\xi')$ that attain the minimum in (15) 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

$$\alpha'_i(\xi, \alpha, \mu)(\xi') = \frac{\alpha_i \int \theta(\xi') \mu'_i(\xi, \mu)(d\theta)}{\sum_h \alpha_h \int \theta(\xi') \mu'_h(\xi, \mu)(d\theta)}. \quad (17)$$

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

¹⁶In sections 4.2 and 5, we abuse notation and let c be a non-negative vector and c_i its i^{th} component.

¹⁷To understand condition (15) 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.

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

$$c_i(\xi, \alpha) + \sum_{h \neq i} \left(\frac{\partial u_h}{\partial c_h} \right)^{-1} \left(\frac{\alpha_i}{\alpha_h} \frac{\partial u_i(c_i(\xi, \alpha))}{\partial c_i} \right) = y(\xi). \quad (18)$$

for each 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 $\hat{c} \in \mathbb{C}(s_0)^I$ if

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

for all i and all $t \geq 0$ and s where $\alpha_0(s) = \alpha_0$ and $\mu_{i,s^0} = \mu_{i,0}$. 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 (13) - (16).

Theorem 2 (The Principle of Optimality) *Suppose A1 holds for all agents. 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 (13) - (16).*

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 $\mathcal{B}$ -margin 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 2 shows that the date 0 optimal plan is consistent in the sense that the continuation plan is indeed the solution from date t onwards.

4.2.1 Discussion: An Alternative Approach

There is an alternative approach to state the dynamic program defined by (13) - (16): 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.¹⁸

¹⁸Since 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.

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 welfare weights, a property that is satisfied by v^* (see Lemma 13 in Appendix B).¹⁹ 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.²⁰

5 Determinants of the Financial Wealth Distribution

In this section we study the determinants of the financial wealth distribution of a dynamically complete markets equilibrium. In section 5.1 we define a competitive equilibrium with sequential trading (CEST). In section 5.2 we construct recursively the date zero-transfers needed to decentralize each PO allocation as a CEST and we employ a properly adapted version of the Negishi's approach to pin down the PO allocation that can be decentralized as a CEST with zero transfers.

5.1 Competitive Equilibrium with Sequential Trading

Every period t , after having observed s^t , agents trade the consumption good and a complete set of Arrow securities in zero net supply. Arrow security ξ' issued at date t on path s pays one unit of consumption if next period's state is ξ' and 0 otherwise. Let $q_t^{\xi'}(s)$ and $a_{i,t}^{\xi'}(s)$ denote the price of security ξ' issued at date t and agent i 's holdings of it at date t , respectively, $a_{i,t} \equiv (a_{i,t}^1, \dots, a_{i,t}^K)$ and $a_i \equiv \{a_{i,t}\}_{t=0}^\infty$

¹⁹Lucas and Stokey [15] do not make this restriction and so it is unclear whether discounting holds in their approach.

²⁰There 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 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.

for all i . All prices are in units of the date- t consumption good, a price system is denoted by $q \equiv \{q_t^{\xi'}\}_{t=0}^{\infty}$. We say that a_i supports $c_i \in \mathbb{C}(s_0)$ if $a_{i,-1} = 0$ and $c_{i,t}(s) + \sum_{\xi'} q_t^{\xi'}(s) a_{i,t}^{\xi'}(s) = y_i(\xi) + a_{i,t-1}(s)$ for all t and s . The budget set of agent i is $\mathbb{B}(q, y_i) \equiv \{c \in \mathbb{C}(s_0) : \exists a_i \text{ that supports } c_i \text{ and } \inf_{(t,s)} a_{i,t}(s) > -\infty\}$.

Definition. A *Competitive Equilibrium with Sequential Trading (CEST)* is an allocation $\{\hat{c}_i\}_{i \in I}$, supporting portfolios $\{\hat{a}_i\}_{i \in I}$ and a price system $\hat{q}$ such that
(CE 1) For every i , $\hat{c}_i \in B(\hat{q}, y_i)$ and $U_i^{P_i}(\hat{c}_i) \geq U_i^{P_i}(c_i)$ for all $c_i \in B(\hat{q}, y_i)$.
(CE 2) For every s and t , $\sum_{i=1}^I \hat{c}_{i,t}(s) = y(\xi)$ and $\sum_{i=1}^I \hat{a}_{i,t+1}^{\xi'}(s) = 0$ for all ξ' .

5.2 Competitive Decentralization

We begin by defining $A_i(\xi, \alpha, \mu)$ as the solution to the functional equation

$$A_i(\xi, \alpha, \mu) = c_i(\xi, \alpha) - y_i(\xi) + \sum_{\xi'} Q(\xi, \alpha, \mu)(\xi') A_i(\xi', \alpha', \mu'), \quad (19)$$

where $\mu'(\xi, \mu)$ and $\alpha'(\xi, \alpha, \mu)(\xi')$ are given by (16) and (17), respectively, and we define the state price²¹

$$Q(\xi, \alpha, \mu)(\xi') = \beta \int \theta(\xi') \mu'_i(\xi, \mu) (d\theta) \frac{\partial u_1(c_1(\xi', \alpha'(\xi, \alpha, \mu)(\xi')) / \partial c_1}{\partial u_1(c_1(\xi, \alpha)) / \partial c_1}. \quad (20)$$

Expression (19) computes recursively the present discounted value of agent i 's excess demand at the PO allocation priced by (20). Intuitively, $A_i(\xi, \alpha, \mu)$ is the transfer needed to support as a competitive equilibrium the PO allocation parametrized by α given (ξ, μ) (see Espino and Hintermaier [11] for further discussion.)

In Theorem 3, we show that A_i is well-defined. Furthermore, we show that there exist a welfare weight α_0 such that A_i is zero for every i . The allocation parametrized by α_0 is the natural candidate to be decentralized as a competitive equilibrium.

Theorem 3 Suppose A1 holds for all agents. Then, there is a unique continuous function A_i solving (19). 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 .

We follow the Negishi's approach to decentralize the PO allocation parametrized by α_0 as a CEST. Define $\mu_{s^{t-1}} = (\mu_{1,s^{t-1}}, \dots, \mu_{I,s^{t-1}})$ and for each s and t let

$$\hat{a}_{i,t}^{\xi'}(s) = A_i(\xi', \alpha'(s_t, \alpha_t(s), \mu_{s^{t-1}})(\xi'), \mu_{s^t}) \quad (21)$$

$$\hat{q}_t^{\xi'}(s) = Q(s_t, \alpha_t(s), \mu_{s^{t-1}})(\xi') \quad (22)$$

with $\mu_{s^{-1}} = \mu_0$, α_t is generated by (17) and $\alpha_0 = \alpha(s_0, \mu_0)$.

Theorem 4 The unique PO allocation $\hat{c}$ solving (9) for $\alpha_0 = \alpha(s_0, \mu_0)$ can be decentralized as a CEST with supporting portfolios $\hat{a}$ and prices $\hat{q}$.

²¹The choice of agent 1 to define Q is without loss of generality since Pareto optimality implies that the intertemporal marginal rates of substitution are equalized across agents.

In a decentralized competitive setting with sequential trading, $A_i(s_t, \alpha_t(s), \mu_{s^{t-1}})$ can be interpreted as the financial wealth that agent i needs to have at date t on path s to afford the consumption bundle corresponding to the PO allocation parametrized by $\alpha_t(s)$ given $(s_t, \mu_{s^{t-1}})$.

REMARK 4: Observe that if agents have both homogeneous and dogmatic beliefs (i.e., $\mu_{i,0} = \mu_i^\theta$ for all i and for some $\theta \in \Delta^{K-1}$), it follows immediately that for each i , $\alpha_{i,t+1}(s) = \alpha_{i,0}$ and $\mu_{i,s^t} = \mu_i^\theta$ for all s and all $t \geq 0$. Therefore, $c_{i,t}(s) = c_i(s_t, \alpha_0)$ for all s and all $t \geq 0$. Very importantly, if we define $\mu^\theta \equiv (\mu_1^\theta \dots \mu_I^\theta)$, we note that $(A_i(1, \alpha_0, \mu^\theta), \dots, A_i(K, \alpha_0, \mu^\theta)) \in \mathbb{R}^K$ and satisfies $A_i(s_0, \alpha_0, \mu^\theta) = 0$ for all i .

5.3 The Fixed Equilibrium Portfolio Property

Definition: Let $A_{i,t}(s) \equiv A_i(s_t, \alpha_t(s), \mu_{s^{t-1}})$ be agent i 's financial wealth at date t on path s . We say that the fixed equilibrium portfolio (FEP hereafter) property holds on s if for each i , there exists $\{\tilde{a}_i(1), \dots, \tilde{a}_i(K)\} \in \mathbb{R}^K$ such that $A_{i,t}(s) = \tilde{a}_i(s_t)$ for all $t \geq 0$. The FEP property holds asymptotically on s if for each i , there exists $\{\tilde{a}_i(1), \dots, \tilde{a}_i(K)\} \in \mathbb{R}^K$ such that $A_{i,t}(s) \rightarrow \tilde{a}_i(s_t)$.

If the FEP property holds on s , any portfolio that decentralizes a PO allocation with a fixed set of non-redundant assets is constant over time.

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 $\mathcal{B}$ -margin of heterogeneity is constant) and degenerate beliefs. Specifically, Remark 4 implies that $\tilde{a}_i(\xi) = A_i(\xi, \alpha_0, \mu^\theta)$ for all ξ since in their setting $\mu_{i,0} = \mu_i^\theta$ 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 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 financial wealth distribution, $A_{i,t}(s) = A_i(s_t, \alpha_0, \mu_{s^{t-1}})$, 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 financial 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' belief is that state 1 is very likely (i.e. $\mu_{i,0}$ 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, $A_{i,t}(s) > 0$ 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 along time because the evolution of the $\mathcal{B}$ -margin of heterogeneity affects the financial wealth distribution through the welfare weights. Therefore, the FEP property does not hold in a dynamically complete markets equilibrium when priors are heterogeneous implying that the result in Judd *et al.* [14] is not robust to the introduction of this margin of heterogeneity. Since agents observe the same data and update their priors in a Bayesian fashion, a much deeper question is whether this trading activity fades out as the agents learn. Our recursive approach permits to study this issue directly.

We say that the $\mathcal{B}$ -margin of heterogeneity vanishes on a path s if $\alpha_t(s)$ converge on s . The following proposition is a consequence of the continuity of A_i and that the convergence of the $\mathcal{B}$ -margin implies learning. It underscores that if a PO allocation can be decentralized as a *CEST*, the associated financial wealth distribution converges to a fixed vector depending only in the current state s_t whenever the $\mathcal{B}$ -margin of heterogeneity vanishes. Consequently, asset trading reduces to the minimum.

Proposition 5 *Suppose A1 holds for all agents. If the $\mathcal{B}$ -margin of heterogeneity vanishes on a path s , then the FEP property holds asymptotically on s .*

6 Limiting Welfare Weights

In this section, we analyze how the $\mathcal{B}$ -margin of heterogeneity evolves over time.

Since we are interested in the dynamic of portfolios on paths that actually occur, we need to introduce the true dgp. In the remainder of this section we assume A0 and that A1 and A2 (a or b) holds for every agent. These three assumptions coupled together imply that every agent has the true dgp in the support of their priors.

From condition (17) and Theorem 2, the ratio of welfare weights is

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\int_{\Delta^{K-1}} \theta(s_t) \mu_{i,s^{t-1}}(d\theta) \alpha_{i,t-1}(s)}{\int_{\Delta^{K-1}} \theta(s_t) \mu_{j,s^{t-1}}(d\theta) \alpha_{j,t-1}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{P_{i,t}(s)}{P_{j,t}(s)}, \quad (23)$$

and, therefore, the asymptotic behavior of $\alpha_t(s)$ depends on the limit behavior of the likelihood ratios.

We show that when every agent has the true dgp in the support of their priors, the likelihood ratios converge and so do welfare weights. The main challenge one faces is

that both the numerator and the denominator of the likelihood ratio are vanishing. So, it is crucial to understand their relative rate of convergence. Our approach consists in finding a common normalization that makes the rates of convergence of the numerator and denominator of the normalized likelihood ratio comparable. While looking for that normalization, we found some technical difficulties that forced us to treat separately the cases with countable and uncountable support.

6.1 Countable Support

We first consider the case where the support of every agent's 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). If the support of agent i 's belief is countable and contains θ^* , it is well-known that agent i learns θ^* . Moreover, the true probability distribution is absolutely continuous with respect to agent i 's prior. As Sandroni [19] shows, this condition is equivalent to the convergence to a positive (finite) number of the ratio of agent i 's prior through date t to the true probability distribution of the first t states. Then, the rate of convergence to zero of $P_{i,t}$ is similar to that of $P_t^{\theta^*}$. In fact, in Proposition 6 we show that

$$\frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} \rightarrow \mu_{i,0}(\theta^*) \quad P^{\theta^*} - a.s. \quad (24)$$

Proposition 6 *Suppose A0 holds. If agent i satisfies A1 and A2.a, then (24) holds.*

Proposition 6 implies that one can use $P_t^{\theta^*}$ to normalize every $P_{i,t}$ to show that the agents' normalized likelihood ratios, $\frac{P_{i,t}(s)/P_t^{\theta^*}(s)}{P_{j,t}(s)/P_t^{\theta^*}(s)}$, also have a finite positive limit. This implies that

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

Since A_i is homogeneous of degree zero in α and $\mu_{i,s^{t-1}}$ converges weakly to $\mu_i^{\theta^*}$ for every agent i , it follows by Proposition 5 that for every state ξ ,

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

where $\alpha_i^* = \alpha_{i,0} \mu_{i,0}(\theta^*)$ and $\alpha_{i,0}$ is the welfare weight defined in Theorem 3. Therefore, we obtain the following result which completely characterizes the limiting properties of the economy.

Theorem 7 *Suppose A0 holds. If A1 and A2.a hold for all agents, then, $P^{\theta^*} - a.s.$, every PO allocation converges to the PO allocation parametrized by $(\alpha^*, \mu^{\theta^*})$. Furthermore, $P^{\theta^*} - a.s.$, the FEP property holds asymptotically with $\tilde{a}_i(\xi) = A_i(\xi, \alpha^*, \mu^{\theta^*})$ for all ξ and i .*

6.2 Uncountable Support

Now we turn to the case where every agent has beliefs with uncountable support that satisfy A2.b. Under our assumptions, the true dgp is *never* absolutely continuous with respect to agent i 's prior and so Sandroni [19, Proposition B.1 and B.2] implies that the ratio of agent i 's prior through date t to the true probability distribution of the first t states converges to zero, i.e. $\frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} \rightarrow 0$, P^{θ^*} - a.s. Therefore, normalizing with $P_t^{\theta^*}$ leaves unanswered the question about the convergence of the likelihood ratios since both the normalized numerator and denominator still converge to zero. The proper common normalization is provided by a result in Phillips and Ploberger [18, Theorem 4.1] (stated as Theorem 18 in the appendix for completeness) that let us establish that under assumptions A0, A1 and A2.b

$$\frac{1}{\frac{(2\pi)^{(K-1)/2} f_{i,0}(\theta^*)}{(t \phi^*)^{1/2}}} \frac{P_{i,t}(s)}{P_t^{\hat{\theta}_t}(s)} \rightarrow 1 \quad P^{\theta^*} - a.s., \quad (25)$$

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 A0 holds. If agent i satisfies A1 and A2.b, then (25) holds.*

Notice that A2.b embodies two assumptions on priors: (i) the truth is in the support of agent i 's prior, i.e. $f_{i,0}(\theta^*) > 0$, and (ii) the density is continuous. The latter assumption is not necessary to argue that agent i learns θ^* but it is necessary to apply Phillips and Ploberger to show the convergence of the likelihood ratios.

Proposition 8 implies that if A1 and A2.b holds for every agent, one can use $P_t^{\hat{\theta}_t}$ to normalize every $P_{i,t}$ to show that the agents' normalized likelihood ratios, $\frac{P_{i,t}(s)/P_t^{\hat{\theta}_t}(s)}{P_{j,t}(s)/P_t^{\hat{\theta}_t}(s)}$, also have a finite positive limit. This implies that

$$\frac{\alpha_{i,t}(s)}{\alpha_{j,t}(s)} = \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{P_{i,t}(s)/P_t^{\hat{\theta}_t}(s)}{P_{j,t}(s)/P_t^{\hat{\theta}_t}(s)} \rightarrow \frac{\alpha_{i,0}}{\alpha_{j,0}} \frac{f_{i,0}(\theta^*)}{f_{j,0}(\theta^*)} \quad P^{\theta^*} - a.s.$$

By a reasoning analogous to the one we used in the countable case, it follows that for every state ξ ,

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

where $\alpha_i^* = \alpha_{i,0} f_{i,0}(\theta^*)$ and $\alpha_{i,0}$ is the welfare weight defined in Theorem 3. We summarize all these results in the following theorem.

Theorem 9 *Suppose A0 holds. If A1 and A2.b hold for all agents, then every efficient allocation converges to the PO allocation parametrized by $(\alpha^*, \mu^{\theta^*})$, P^{θ^*} - a.s. Furthermore, P^{θ^*} - a.s., the FEP property holds asymptotically with $\tilde{a}_i(\xi) = A_i(\xi, \alpha^*, \mu^{\theta^*})$ for all ξ and i .*

6.3 Discussion

Theorems 7 and 9 argue forcefully that when the agents know the data is generated by i.i.d. draws from a common distribution (i.e., $A1$ holds) and the true parameter is in the support of every agent's belief (i.e., $A2$ 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 $\{\alpha_i^*\}_{i \in \mathcal{I}}$. That is, the density of the agents' 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 allows to compute the limiting allocation without solving for the equilibrium, using only exogenous parameters describing the economy at date 0.

The mechanics to obtain our results is to exploit A_i '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. Note that agent i 's welfare weight, α_i , is the planner's current valuation of an additional unit of agent i 's utility. By consistency, then, $\alpha_i \beta \int \theta(\xi') \mu'_i(\xi, \alpha, \mu)(d\theta)$ is the planner's current valuation of an additional unit of agent i 's next period utility at state ξ' . This is the economics behind the law of motion (23), before the normalization. Secondly, since the evolution of these weights is driven by the behavior of the likelihood ratios, we are lead to study their dynamics.

So far we have implicitly made two critical assumptions regarding the support of the agent's belief, namely, (i) it contains the true parameter and (ii) it has the same dimension for all agents. The logic behind these two assumptions is as follows. As Blume and Easley [4] and Sandroni [19] 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, for those agents whose welfare weights converge to zero consumption vanishes and their wealth approaches the so-called natural debt limit (see condition (19)). The limiting economy, thus, mimics the economy where those agents' property rights on their individual endowments have been redistributed among the other agents. But then those agents are irrelevant to understand the long-run properties of the portfolios that decentralize PO allocations.

7 Persistent Trade

In this section we argue, by means of two examples, that to show the *FEP* property holds asymptotically it is essential to assume that some agent learns the true conditional probability of the states (section 7.1) and that every agent knows

the dgp (section 7.2).

In the examples below we assume that $A0$ holds and, for simplicity, that there are two states of nature, that is $K = 2$.

7.1 Example 3: Dogmatic Priors

Here we show that the assumption that agents learn the true conditional probability of the states is necessary in the sense that when it is not satisfied, the *FEP* property may not hold even if agents' priors satisfy $A1$, no matter how close they are to the truth and with respect to each other.

We assume there are two agents whose priors beliefs are point masses on θ and $\tilde{\theta}$, respectively, where $\theta \neq \tilde{\theta}$ and $\theta^*(1) \ln \frac{\theta(1)}{\tilde{\theta}(1)} + \theta^*(2) \ln \frac{1-\theta(1)}{1-\tilde{\theta}(1)} = 0$. Since agents have heterogeneous "dogmatic" priors with the same entropy, then it can be shown that both agents survive.^{22, 23} The ratio of one-period-ahead conditional probabilities, $\frac{p_{1,t}}{p_{2,t}}$, is a simple random variable that takes values in $\left\{ \frac{\theta(1)}{\tilde{\theta}(1)}, \frac{1-\theta(1)}{1-\tilde{\theta}(1)} \right\}$. The logarithm of the likelihood ratio can be written as the sum of conditional zero mean random variables as follows:

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

where $E^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = 0$ and $\text{var}^{P^{\theta^*}}(x_k | \mathcal{F}_{k-1})(s) = E^{P^{\theta^*}}(x_k^2 | \mathcal{F}_{k-1})(s) = E^{P^{\theta^*}}(x_k^2) > 0$. So, the log likelihood ratio is the sum of uniformly bounded random variables with zero conditional mean and conditional variance bounded away from zero. Once again, it follows by Freedman [13, 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,

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

²²See Blume and Easley [5] for a general analysis of optimal consumption paths in i.i.d. economies where agents have degenerate prior beliefs.

²³The assumption that priors are dogmatic is not necessary for the result to be true. The necessary condition is that the limits of the agent's posterior beliefs are different but have identical entropy. The latter condition, that holds in our example, implies that even though the result is generic in θ^* , it is not generic in the space of priors satisfying $A1$.

Inspecting condition (23), it is evident that welfare weights do not converge in this example. Since beliefs are degenerate at θ_i , posterior beliefs are also degenerated at θ_i and, therefore, agent i 's financial wealth is $A_i(\xi, \alpha_t(s), (\mu_1^\theta, \mu_2^{\tilde{\theta}}))$. We can conclude that, $P^{\theta^*} - a.s.$, the *FEP* property does not necessarily hold asymptotically.

7.2 Example 4: Different Likelihood Functions

In example 2 we show that the *FEP* property may not hold asymptotically when no agent satisfies *A1*. To underscore the importance of assuming that *A1* holds for every agent, here we consider, instead, the case in which *A1* does not hold for one agent while it holds for the other. One agent, on the one hand, has a prior satisfying *A1* and *A2.b* 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. Moreover, unlike example 2, it is generic in the space of beliefs.

Agent 1's prior satisfies *A1* and *A2.b*. It is well-known that the set of parameters $\theta \in \Delta^1$ for which μ_{1,s^t} converges weakly to μ_1^θ , $P^\theta - a.s.$, is a $\mu_{1,0}$ -full measure subset of Δ^1 (see Schwartz [21, Theorems 3.2 and 3.3]). We suppose $\theta^* \in \Delta^1$ is an element of that subset and so agent 1 learns, $P^{\theta^*} - a.s.$, the true one-period-ahead conditional probability.

7.2.1 Agent 2's priors

Now we construct agent 2's priors. Fix $\tilde{\theta} \in \Delta^1$ such that $\tilde{\theta} \neq \theta^*$ and let $p^{\tilde{\theta}}(1|s_0) \in (0, 1)$ and $p^{\tilde{\theta}}(2|s_0) = 1 - p^{\tilde{\theta}}(1|s_0)$. For each partial history s^t , $t \geq 1$, define

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

and $p^{\tilde{\theta}}(2|s^t) = 1 - p^{\tilde{\theta}}(1|s^t)$. That is, $p^{\tilde{\theta}}(\cdot|s^t)$ is given by the true one period-ahead conditional probability whenever the likelihood ratio $\prod_{k=1}^t \frac{p^{\tilde{\theta}}(s_k|s^{k-1})}{p_1(s_k|s^{k-1})}$ is smaller than or equal to one. When that ratio is strictly greater than one, on the other hand, $p^{\tilde{\theta}}(\cdot|s^t)$ is given by $\tilde{\theta}(\cdot)$.

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

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

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

The prior of agent 2 is

$$P_2(B) \equiv \tilde{P}(B),$$

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

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

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 following proposition shows that the likelihood ratio of 2's prior to 1's prior fluctuates at least between 1 and $+\infty$. Its proof follows from two steps

■ First, we argue that $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$ almost surely. Note that on the set of paths where $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$, agent 2's one-period-ahead conditional probability of state 1 is given by $\tilde{\theta}(1)$ and so the set of paths where $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ is a subset of the set of paths where $\frac{P_t^{\tilde{\theta}}(s)}{P_{1,t}(s)}$ is bounded away from zero eventually. But since agent 1's learns almost surely, that set has measure zero.

■ Second, we argue that $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$ almost surely. To show this we define Ω_1 as the set of paths where the likelihood ratio is greater than one infinitely often and argue that: (a) the likelihood ratio exceeds any pre-specified upper bound infinitely many times on Ω_1 and (b) Ω_1 has P^{θ^*} -measure 1. To see why (a) holds note that whenever the likelihood ratio is greater than 1, one must have that (i) the ratio of agent 2's to agent 1's one-period-ahead conditional probabilities is bounded away from one²⁴ and (ii) the conditional probability that the likelihood ratio exceeds a pre-specified upper bound in a fixed number of periods is uniformly bounded away

²⁴This is because only agent 1's one-period-ahead conditional probability converges to the truth on those paths.

from zero. An application of Levy's conditional form of the Second Borel-Cantelli Lemma implies that the likelihood ratio exceeds a pre-specified upper bound infinitely many times on Ω_1 . To prove (b), it suffices to show that $\Omega \setminus \Omega_1$ has measure zero. Note that on that set, agent 2's one-period-ahead conditional probability coincides with the true one except for finitely many periods. Therefore, $\Omega \setminus \Omega_1$ is a subset of the set of paths where $\frac{P_t^{\theta^*}(s)}{P_{1,t}(s)}$ is bounded above. But our assumption that agent 1's prior satisfies A1 and A.2b and Sandroni's results [19, Propositions B.1 and B.2.] implies that the latter set has P^{θ^*} -measure zero.²⁵

Proposition 10 *Suppose A0 holds. If agent 1's prior satisfies A1 and A2.b, 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.*

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

We consider again the economy described in Section 3 where portfolios are given by (5). 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 A2.b 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 A0 holds. If agent 1's prior satisfies A1 and A2.b, 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 .*

7.2.3 Further Remarks

In Sandroni's [19] terminology, agent 1 *eventually makes accurate next period predictions* while agent 2 does not and yet both agent survive. At a first glance,

²⁵If agent 1 had a prior belief with countable support (so that A2.a holds) 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.

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, $A0$, $A1$ and $A2.b$ 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 $A1$.²⁶ 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 $A1$, 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 $A1$ shuts down a margin of heterogeneity that might be critical not only for survival but also for asset pricing and trading volume.

8 Concluding Remarks

If every agent knows that the data is generated by i.i.d. draws from a common distribution and some agent has the true probability distribution over states of nature in the support of her 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.

²⁶They do have an example in which agent 1 satisfies $A1$ 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, the likelihood ratios converge with probability one.

9 Appendix A

In this Appendix we show that (5), used throughout Examples 1 - 4, denotes the equilibrium Arrow security holdings.

First, observe that the planner's problem is

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

The first order conditions imply that

$$\alpha_i \beta^t P_{i,t}(s) \frac{1}{c_{i,t}(s)(\alpha)} = \eta_t(s)(\alpha) \quad \text{for all } i, t \text{ and } s, \quad (26)$$

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

$$c_{i,t}(s)(\alpha) = \frac{\alpha_i P_{i,t}(s)}{\alpha_1 P_{1,t}(s) + \alpha_2 P_{2,t}(s)} y_t(s) \quad \text{for all } i, t \text{ and } s. \quad (27)$$

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

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

Using (26) and (27) 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

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

Thus, equilibrium portfolios at date t are

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

10 Appendix B

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

Proof of Lemma 1. Available as supplementary material. ■

Before proving Theorem 2, we need some definitions and preliminary results. We begin by 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 (13) - (15). In Proposition 16 we argue that the operator has a unique fixed point and in Proposition 17 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 beliefs at date 0, $s_0 = \xi$ and $\mu_0 \equiv (\mu_{1,0}, \dots, \mu_{I,0}) = \mu$, define the *utility possibility correspondence* by

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

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. Available as supplementary material. ■

The following result characterizes the value function v^* and hints how to restrict the set of functions to consider in (13)-(16). For any $\alpha \in \mathbb{R}_+^I$, define $\|\alpha\| = \sum_{i=1}^I \alpha_i$ and let F_H be the linear space of functions $f : S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1}) \rightarrow \mathbb{R}_+$ that are homogeneous of degree one with respect to α , bounded in the norm

$$\|f\|_\infty = \sup_{(\xi, \alpha, \mu) \in S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1})} \{f(\xi, \alpha, \mu) : \|\alpha\| = 1\},$$

and continuous with respect to the metric induced by $\|\cdot\|_\infty$. Note that with this metric F_H is a complete metric space.

Lemma 13 The value function $v^* \in F_H$ and it is increasing in α .

Proof of Lemma 13. Available as supplementary material. ■

Lemma 12 implies that the set of PO allocations can be parametrized by welfare weights α . Indeed, it is straightforward to prove that

$$v^*(\xi, \alpha, \mu) = \sup_{w \in \mathcal{U}(\xi, \mu)} \sum_{i=1}^I \alpha_i w_i, \quad (28)$$

Notice that the maximum in (28) 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 (15) is satisfied.

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

Proof of Lemma 14. Available as supplementary material. ■

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

$$(Tv)(\xi, \alpha, \mu) = \max_{(c, w'(\xi'))} \sum_{i \in \mathcal{I}} \alpha_i \left\{ u_i(c_i) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu'_i(\xi, \mu) (d\theta) w'_i(\xi') \right\} \quad (29)$$

subject to (14) - (16).

In Proposition 16 we apply Boyd's [6] Weighted Contraction Mapping Theorem to show that the operator T has a unique fixed point on F_H . Furthermore, Proposition 17 below shows that the value function (9) is the unique solution to the functional equation defined by (13)-(16).

Let φ be a positive continuous function on $S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1})$ and let C_φ be the space of all continuous functions f on $S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1})$ which have bounded φ -norm (i.e., $\sup \left\{ \frac{|f(x)|}{\varphi(x)} \right\}$ is finite.) Here we state the version of Boyd's theorem that we use in Proposition 16:

Theorem 15 (Boyd's Weighted Contraction Mapping Theorem.) *Let $T : C_\varphi \rightarrow C_\varphi$ such that*

- (1) *T is non-decreasing;*
- (2) *$T(0) \in C_\varphi$;*
- (3) *$T(f + a\varphi)(\xi, \alpha, \mu) \leq Tf(\xi, \alpha, \mu) + a\beta\varphi$ for all (ξ, α, μ) .*

Then T has a unique fixed point.

In this paper, we let $\varphi(\xi, \alpha, \mu) = \|\alpha\|$ for all (ξ, α, μ) and we further restrict C_φ by requiring homogeneity with respect to α , i.e. we set $C_\varphi = F_H$. It is immediate that $\|\cdot\| \in F_H$ as required. The key to be able to apply a contraction mapping argument is that F_H is a complete metric space.

Proposition 16 *The operator T has a unique fixed point $v \in F_H$ and the corresponding policy functions are continuous.*

Proof of Proposition 16. 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 $0 \leq w'_i(\xi') \leq \|f\|_\infty$ for all i and all ξ' , it follows that $\|Tf\|_\infty < \infty$. Since $\mu'(\xi, \mu)$ is weakly continuous in μ for all ξ (Easley and Kiefer [10, Theorem 1]), it follows by the Maximum Theorem that $(Tf)(\xi, \alpha, \mu)$ is continuous in (α, μ) for all ξ (Easley and Kiefer [10, 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, standard arguments show 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 hypothesis (1) - (3) of Boyd's [6] Weighted Contraction Mapping Theorem to conclude that there exists a unique $v \in F_H$ such that $v = Tv$.

(1) 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, $(Tf)(\xi, \alpha, \mu) \leq (Tg)(\xi, \alpha, \mu)$ for all (ξ, α, μ) .

(2) Let $f(\xi, \alpha, \mu) = 0$ for all (ξ, α, μ) . Since f is HOD 1 with respect to α and $T : F_H \rightarrow F_H$, then $T(0) \in F_H$.

(3) Consider any arbitrary $a > 0$ and note that $f(\xi, \alpha, \mu) + \|\alpha\| a$ is HOD 1 with respect to α . Let $(\hat{c}, \hat{w}'(\xi'))(\xi, \alpha, \mu)$ attain $T(f + \|\alpha\| a)(\xi, \alpha, \mu)$. Fix $(\xi', \mu'(\xi, \mu))$, let $\tilde{f}(\alpha) = f(\xi', \alpha, \mu')$ and notice that $\|\alpha\| = 1$ if $\alpha \in \Delta^{I-1}$. Define

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

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

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 $\alpha \in \Delta^{I-1}$, $\alpha \cdot w' = \alpha \cdot (w - a) = \alpha \cdot w - a \leq \tilde{f}(\alpha)$.

(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 and define

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

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

$$\begin{aligned} \bar{w}_1^a(w_2) &= \min_{0 \leq \alpha_1 \leq 1} \left(\frac{\tilde{f}(\alpha_1, 1 - \alpha_1)}{\alpha_1} + \frac{(a - w_2)}{\alpha_1} \right) + w_2 \\ &= \min_{0 \leq \alpha_1 \leq 1} \left(\tilde{f}\left(1, \frac{1}{\alpha_1} - 1\right) + \frac{(a - w_2)}{\alpha_1} \right) + w_2 \\ &= \tilde{f}(1, 0) + (a - w_2) + w_2 = \tilde{f}(1, 0) + a, \end{aligned}$$

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. $\bar{w}_1^a(w_2) = \bar{w}_1^a = \tilde{f}(1, 0) + a$ for all $w_2 \leq a$.

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

$$\tilde{f}(\alpha) = \alpha_1 \tilde{f}\left(1, \frac{1 - \alpha_1}{\alpha_1}\right) \geq \alpha_1 \tilde{f}(1, 0) = \alpha \cdot w',$$

and thus $w' \in U^0$.²⁷ 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.

²⁷We underscore here that without assuming that $\tilde{f}$ is HOD 1 and monotone (i.e., $\tilde{f} \in F_H$), this result does not necessarily hold. More precisely, these assumptions guarantee that $\arg \min \left(\frac{\tilde{f}(\alpha_1, 1 - \alpha_1)}{\alpha_1} + \frac{a - w_2}{\alpha_1} \right) = 1$. 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

$$\bar{w}_1^a = \min \left(\frac{\tilde{f}(\alpha_1, 1 - \alpha_1)}{\alpha_1} + \frac{a - w_2}{\alpha_1} \right) > \min \frac{\tilde{f}(\alpha_1, 1 - \alpha_1)}{\alpha_1} + \min \frac{a - w_2}{\alpha_1} = \bar{w}_1^0 + a - w_2.$$

Notice that if $(\widehat{c}, \widehat{w}'(\xi'))(\xi, \alpha, \mu)$ attain $T(f + \|\alpha\| a)(\xi, \alpha, \mu)$, there is $w'(\xi') \in U^0$ such that $w'(\xi') \geq \widehat{w}'(\xi') - a$ for all ξ' . By monotonicity, $(\widehat{c}, w'(\xi') + a)(\xi, \alpha, \mu)$ also attain $T(f + \|\alpha\| a)(\xi, \alpha, \mu)$. Note that for any (ξ, α, μ) , it follows by definition that

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

and thus

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

Since (ξ, α, μ) was arbitrarily chosen, we can conclude that the operator $T(f + \|\alpha\| a)(\xi, \alpha, \mu) \leq Tf(\xi, \alpha, \mu) + \beta \|\alpha\| a$ for all (ξ, α, μ) as desired. ■

Proposition 17 v^* solves the functional equation defined by (13) - (16).

Proof. Given $s_0 = \xi$ and $c_i \in \mathbb{C}(s_0)$, for each ξ' let ${}_{\xi'}c_i \in \mathbb{C}(\xi')$ be the ξ' -continuation of c_i , defined as

$$\begin{aligned} {}_{\xi'}c_{i,0} &= c_i(\xi') \\ {}_{\xi'}c_i(s^t) &= c_i(\xi', s_2, \dots, s_t) \quad \text{for all } t \geq 2. \end{aligned}$$

Note that

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

subject to $[v^*(\xi', \tilde{\alpha}, \mu'(\xi, \mu)) - \sum_{i=1}^I \tilde{\alpha}_i w'_i(\xi')] \geq 0$ for all $\tilde{\alpha} \in \Delta^{I-1}$. 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 14. Consequently, v^* solves the functional equation (13) - (16). ■

Theorem 2 uses our previous results to conclude that the corresponding version of the principle of optimality holds in our setting.

Proof of Theorem 2. Consider the operator T . Let (c, α', w') be the policy functions associated with its fixed point v^* . $(c_i^*)_{i=1}^I$ is the feasible allocation generated by (c, α', w') . Suppose, on the contrary, that $(c_i^*)_{i=1}^I$ is not PO. Then, there exists an alternative feasible allocation $(\widehat{c}_i)_{i=1}^I$ such that

$$\begin{aligned} & \sum_{i=1}^I \alpha_i \left\{ u_i(c_{i,0}^*) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi' c_i^*) \right\} \\ & < \sum_{i=1}^I \alpha_i \left\{ u_i(\widehat{c}_{i,0}) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi' \widehat{c}_i) \right\} \\ & \leq v^*(\xi, \alpha, \mu) \end{aligned} \quad (30)$$

For each agent i , define $\widehat{w}_i'(\xi') = U_i^{P_i, \xi'}(\xi' \widehat{c}_i)$ for all ξ' . Observe that $\sum_{i=1}^I \widehat{c}_{i,0} = y(\xi)$, $(\widehat{w}_i'(\xi'))_{i=1}^I \in \mathcal{U}(\xi', \mu'(\xi, \mu))$ for all ξ' and that Lemma 14 implies that $v^*(\xi', \widetilde{\alpha}, \mu'(\xi, \mu)) \geq \sum_{i=1}^I \widetilde{\alpha}_i \widehat{w}_i'(\xi')$ for all $\widetilde{\alpha} \in \Delta^{I-1}$ and all ξ' . Then, $(\widehat{c}_{i,0}, \widehat{w}_i')_{i=1}^I$ satisfies (14) - (16).

Moreover, (30) can be rewritten as

$$\begin{aligned} & \sum_{i=1}^I \alpha_i \left\{ u_i(c_i(\xi, \alpha)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) w_i'(\xi') \right\} \\ & < \sum_{i=1}^I \alpha_i \left\{ u_i(\widehat{c}_{i,0}) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) \widehat{w}_i'(\xi') \right\} \end{aligned}$$

This contradicts that (c, α', w') are the policy functions associated with v^* .

Suppose that there is a PO allocation $(c_i^*)_{i=1}^I$ that is not generated by (c, α', w') . Let $(c_i)_{i=1}^I$ be the feasible allocation generated by (c, α', w') . For each agent i , define

$$\begin{aligned} w_i'(\xi') &= U_i^{P_i, \xi'}(\xi' c_i) \text{ and } (w_i^*)'(\xi') = U_i^{P_i, \xi'}(\xi' c_i^*) \text{ for all } \xi'. \text{ Note that} \\ \sum_{i=1}^I \alpha_i U_i^{P_i}(c_i) &= \sum_{i=1}^I \alpha_i \left\{ u_i(c_{i,0}) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi' c_i) \right\} \\ &= \sum_{i=1}^I \alpha_i \left\{ u_i(c_i(\xi, \alpha)) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) w_i'(\xi') \right\} \\ &= v^*(\xi, \alpha, \mu) \\ &> \sum_{i=1}^I \alpha_i \left\{ u_i(c_{i,0}^*) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) (w_i^*)'(\xi') \right\} \\ &= \sum_{i=1}^I \alpha_i \left\{ u_i(c_{i,0}^*) + \beta \sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_i'(\xi, \mu) (d\theta) U_i^{P_i, \xi'}(\xi' c_i^*) \right\} \\ &= \sum_{i=1}^I \alpha_i U_i^{P_i}(c_i^*) \end{aligned}$$

where the inequality follows because $\sum_{i=1}^I c_{i,0}^* = y(\xi)$, $\left((w_i^*)'(\xi')\right)_{i=1}^I \in \mathcal{U}(\xi', \mu'(\xi, \mu))$ for all ξ' ; i.e., $\left(c_{i,0}^*, (w_i^*)'\right)_{i=1}^I$ satisfies (14) - (16). But then $(c_i^*)_{i=1}^I$ is not PO, a contradiction. ■

Proof of Theorem 3. Let F be the linear space of functions $f : S \times \mathbb{R}_+^I \times \mathcal{P}(\Delta^{K-1}) \rightarrow \mathbb{R}_+$ that are bounded in the norm $\|\cdot\|_\infty$ and continuous with respect to the metric induced by $\|\cdot\|_\infty$. Consider the operator $\tilde{T}_i$ on F given by

$$\begin{aligned} (\tilde{T}_i h)(\xi, \alpha, \mu) &= (c_i(\xi, \alpha) - y_i(\xi)) u_1'(c_1(\xi, \alpha)) \\ &\quad + \sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu_1'(\xi, \mu) (d\theta) h(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)). \end{aligned}$$

Step 1. First we check that $\tilde{T}_i : F \rightarrow F$. Suppose that $h \in F$. Consider first

$$\sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu_1'(\xi, \mu) (d\theta) h(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)), \quad (31)$$

and observe that α' and μ' are both continuous. Also, it follows that $\int_{\Delta^{K-1}} \theta(\xi') \mu_1'(\xi, \mu) (d\theta)$ is continuous. Thus, the expression (31) is continuous in (ξ, α, μ) . (31) is bounded because h is bounded and $\sum_{\xi'} \int_{\Delta^{K-1}} \theta(\xi') \mu_1'(\xi, \mu) (d\theta) = 1$.

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 $\tilde{T}_i h \in F$.

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

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

$$\begin{aligned} \tilde{T}_i(h + a)(\xi, \alpha, \mu) &= (c_i(\xi, \alpha) - y_i(\xi)) \frac{\partial u_1(c_1(\xi, \alpha))}{\partial c_1} \\ &\quad + \sum_{\xi'} \beta \int_{\Delta^{K-1}} \theta(\xi') \mu_1'(\xi, \mu) (d\theta) h(\xi', \alpha'(\xi, \alpha, \mu)(\xi'), \mu'(\xi, \mu)) + \beta a. \\ &= (\tilde{T}_i h)(\xi, \alpha, \mu) + \beta a. \end{aligned}$$

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

Therefore, we can apply the contraction mapping theorem to conclude that $\tilde{T}_i$ is a contraction with a unique solution $h_i \in F$ for each i .

To complete the proof, define $A_i(\xi, \alpha, \mu) = h_i(\xi, \alpha, \mu)/u'_1(c_1(\xi, \alpha, \mu))$. It can be checked immediately that A_i is a continuous and it is the unique solution to the functional equation (19). Notice that $\sum_{i=1}^I A_i(\xi, \alpha, \mu)$ is the unique fixed point of the operator $\hat{T} : F \rightarrow F$ defined by

$$\hat{T}z(\xi, \alpha, \mu) = \sum_{\xi'} M(\xi, \alpha, \mu)(\xi') z(\xi', \alpha', \mu'). \quad (32)$$

Since $R(\xi, \alpha, \mu) = 0$ for all (ξ, α, μ) is a fixed point of $\hat{T}$, it follows by uniqueness that

$$\sum_{i=1}^I A_i(\xi, \alpha, \mu) = 0, \quad \text{for all } (\xi, \alpha, \mu).$$

Step 3. Available as supplementary material. ■

Proof of Theorem 4. It follows by standard arguments. ■

11 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 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 [19], $P^{\theta^*} - a.s.$,

$$0 < \lim_{t \rightarrow \infty} \frac{P_{i,t}(s)}{P_t^{\theta^*}(s)} < \infty, \quad (33)$$

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

$$\lim_{t \rightarrow \infty} \frac{\sum_{\theta \neq \theta^*} P_t^\theta(s) \frac{\mu_{i,0}(\theta)}{1 - \mu_{i,0}(\theta^*)}}{P_t^{\theta^*}(s)} = 0. \quad (34)$$

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 (34). ■

Let $Q_{h,t}$, $t \geq 1$, be the measure defined by the following Radon Nykodim derivative

$$\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)} = \frac{(2\pi)^{\frac{K-1}{2}} f_{h,0}(\theta^*)}{|B_t^*(s)|^{\frac{1}{2}}} e^{l_t(\hat{\theta}_t(s))}, \quad (35)$$

where $l_t(\theta) \equiv \ln \frac{P_t^\theta}{P_t^{\theta^*}}$, $\hat{\theta}_t$ is the Maximum Likelihood Estimator (MLE) of θ , $l_t^{(1)}(\theta)$ is the score, B_t^θ is the conditional variance (matrix) process of the martingale $l_t^{(1)}(\theta)$ under $P_t^{\theta^*}$, $B_t^* = B_t^{\theta^*}$ and for any matrix Z , $|Z|$ denotes its determinant. Set $l_0^{(1)}(\theta_\xi) = 0$, let $\varepsilon_t(\theta_\xi) = l_t^{(1)}(\theta_\xi) - l_{t-1}^{(1)}(\theta_\xi)$ and let ε_t be the vector in $\mathbb{R}^K$ with entries $\varepsilon_t(\theta_1), \dots, \varepsilon_t(\theta_K)$. Then, $B_t^\theta \equiv \sum_{k=1}^t E^{P^\theta}(\varepsilon_k \varepsilon_k' | \mathcal{F}_{k-1})$ and it can be shown that $B_t^* = t \phi^*$ where $\phi^* = \left| E^{P^{\theta^*}}(\varepsilon_1 \varepsilon_1' | \mathcal{F}_0) \right|$.

The following Theorem, due to Phillips and Ploberger [18, page 392], considers the case in which $\theta^* \in \Delta^1$ (i.e., $K = 2$) and will be used in the proof of Proposition 8. The Theorem extends to the case of a multivariate θ (see Phillips and Ploberger [18, Theorem 4.1, page 396]).

Theorem 18 (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

$$\lim_{t \rightarrow \infty} \sqrt{B_t^*} \int_{\omega_\delta} f_{h,0}(\theta) \frac{P_t^\theta(s)}{P_t^{\theta^*}(s)} d\theta = 0 \quad P^{\theta^*} - a.s.$$

- (C7) The density of the belief, $f_{h,0}(\theta)$, is continuous at θ^* with $f_{h,0}(\theta^*) > 0$.

If $Q_{h,t}$ is the measure defined by the Radon Nykodim derivative in (35), 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. For simplicity, we consider the case in which $\theta \in \Delta^1$. The case in which $\theta \in \Delta^{K-1}$ can be proved in a similar fashion by verifying the conditions in Theorem 4.1 in Phillips and Ploberger [18].

We need to verify that (C.1) - (C.7) in Theorem 18 hold. Let n_t be the number of times that state 1 has occurred up to date t .

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

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

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

Set $l_0^{(1)}(\theta) = 0$ and let $\varepsilon_k(\theta) = l_k^{(1)}(\theta) - l_{k-1}^{(1)}(\theta)$. Then $\varepsilon_k(\theta)$ takes values $\frac{1}{\theta(1)}$ and $-\frac{1}{1-\theta(1)}$ with probabilities $\theta(1)$ and $1-\theta(1)$. Therefore,

$$\begin{aligned} B_t^\theta &= \sum_{k=1}^t E^{P^\theta} \left[\varepsilon_k(\theta)^2 \middle| \mathcal{F}_{k-1} \right] \\ &= \sum_{k=1}^t \left(\theta(1) \left(\frac{1}{\theta(1)} \right)^2 + (1-\theta(1)) \left(-\frac{1}{1-\theta(1)} \right)^2 \right) \\ &= \sum_{k=1}^t \left(\frac{1}{\theta(1)} + \frac{1}{1-\theta(1)} \right) \\ &= t \left(\frac{1}{\theta(1)} + \frac{1}{1-\theta(1)} \right) = t \frac{1}{\theta(1)(1-\theta(1))} \end{aligned} \tag{36}$$

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

(C.3) Notice that by the SLLN,

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

so the desired result holds.

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

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

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

(C.6) By the SLLN, there is $T(s)$ such that for all $t \geq T(s)$, $P^{\theta^*} - a.s., \frac{n_{1,t}}{t} \in (\theta^*(1) - \delta/2, \theta^*(1) + \delta/2)$. In addition, there is $\tilde{\delta}$ such that for every $\theta \in \omega_\delta$,

$$\sup_{x \in (\theta^*(1) - \delta/2, \theta^*(1) + \delta/2)} \frac{\theta(1)^x (1 - \theta(1))^{1-x}}{(\theta^*(1))^x (1 - \theta^*(1))^{1-x}} \leq 1 - \tilde{\delta}.$$
 Then,

$$\begin{aligned} \sqrt{B_t^*} \int_{\omega_\delta} f_{h,0}(\theta) \frac{P_t^\theta}{P_t^{\theta^*}} d\theta &= \sqrt{B_t^*} \int_{\omega_\delta} f_{h,0}(\theta) \frac{\theta(1)^{n_{1,t}} (1 - \theta(1))^{n_{2,t}}}{(\theta^*(1))^{n_{1,t}} (1 - \theta^*(1))^{n_{2,t}}} d\theta \\ &= \sqrt{B_t^*} \int_{\omega_\delta} f_{h,0}(\theta) \left(\frac{\theta(1)^{\frac{n_{1,t}}{t}} (1 - \theta(1))^{\frac{n_{2,t}}{t}}}{(\theta^*(1))^{\frac{n_{1,t}}{t}} (1 - \theta^*(1))^{\frac{n_{2,t}}{t}}} \right)^t d\theta \\ &\leq \sqrt{B_t^*} (1 - \tilde{\delta})^t = \sqrt{t} \phi^* (1 - \tilde{\delta})^t \end{aligned}$$

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

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

Note that by (36), $B_t^\theta = t \phi^*$ where $\phi^* \equiv \frac{1}{\theta^*(1)(1 - \theta^*(1))}$. Then,

$$\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)} = \frac{\sqrt{2\pi} f_{h,0}(\theta^*) P_t^{\theta_t^*(s)}(s)}{\sqrt{t} \phi^* P_t^{\theta^*}(s)}$$

and so

$$\lim_{t \rightarrow \infty} \frac{1}{\sqrt{2\pi} f_{h,0}(\theta^*)} \frac{P_{h,t}(s)}{P_t^{\theta_t^*(s)}(s)} = \lim_{t \rightarrow \infty} \frac{P_{h,t}(s)}{Q_{h,t}(s)} = \lim_{t \rightarrow \infty} \frac{\frac{P_{h,t}(s)}{P_t^{\theta^*}(s)}}{\frac{Q_{h,t}(s)}{P_t^{\theta^*}(s)}} = 1 \quad P^{\theta^*} - a.s.$$

■

Let $\tilde{\theta}_t$ denote the i.i.d. random variable that takes values $\tilde{\theta}(1)$ and $1 - \tilde{\theta}(1)$ in states 1 and 2, respectively. Let θ_t^* denote the i.i.d. random variable that takes values $\theta^*(1)$ and $1 - \theta^*(1)$ in states 1 and 2, respectively.

Proof of Proposition 10. We begin with four Lemmas that will be useful to prove the main result. Lemma 19 shows that the set of paths where $\liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$ has full measure. Lemma 22 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 23 and 24 show that the latter set also has full measure.

Lemma 19 $P^{\theta^*} - a.s., \liminf \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1$.

Proof of Lemma 19. 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, $P^{\theta^*} - a.s., p_{1,t}(s) \rightarrow \theta^*(s_t)$ and $\theta^*(\xi) \in (0, 1)$ for all ξ , there exists $\varepsilon > 0$ and $T_1(s) > 1$ 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 $s \in \Omega_1$,

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

and it follows that

$$\Omega_1 \subset \left\{ s : \liminf_{T \rightarrow \infty} \prod_{t=1}^T \frac{\tilde{\theta}_t(s)}{p_{1,t}(s)} > 0 \right\} \quad (37)$$

On the other hand, by the SLLN for uncorrelated random variables with uniformly bounded second moments (Chung [9, page 103]), $P^{\theta^*} - a.s.$,

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

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

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

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

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

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

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

But then, (37) 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 22. The first is Levy's conditional form of the Second Borel-Cantelli Lemma which follows from a more general result due to Freedman [12, Proposition 39] and is stated without proof as Lemma 20. The second result, stated in Lemma 21, considers a sequence of adapted events $\{\Omega_t\}_{t=1}^\infty$ and shows that, $P - a.s.$, a path s that lies in infinitely many of the Ω_t , it also lies in infinitely many elements of the adapted sequence of events $\{\Omega_{1,t}^N\}_{t=N+1}^\infty$ where each event $\Omega_{1,t}^N$, $t > N$, consists of those paths s in Ω_{t-N} for which $s_{t-N} = \dots = s_t = 1$.

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 20 (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*

$$\sum_{t=0}^{\infty} 1_{\Omega_t}(s) = +\infty \quad P - \text{a.s. on } \left\{ \tilde{s} : \sum_{t=0}^{\infty} E^P [1_{\Omega_t} | \mathcal{F}_{t-1}] (\tilde{s}) = +\infty \right\}.$$

Lemma 21 *Let $\{\Omega_t\}_{t=0}^\infty$ be a sequence of events adapted to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. Then*

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

Proof of Lemma 21. Available as supplementary material. ■

Lemma 22 *Suppose A0 holds and that agent 1's prior satisfies A1 and A2.b. Then, $P^{\theta^*} - \text{a.s. on } \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$.*

Proof of Lemma 22. 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 $p_{1,t}(s) \rightarrow \theta^*(s_t)$, there exists $T(s)$ such that for every $t \geq T(s)$, $\theta^*(s_t) - \frac{\varepsilon}{2} \leq p_{1,t}(s) \leq \theta^*(s_t) + \frac{\varepsilon}{2}$. Without loss of generality, suppose $\tilde{\theta}(1) > \theta^*(1) + \frac{\varepsilon}{2}$. Let T^a be the smallest integer such that $\left(\frac{\tilde{\theta}(1)}{\theta^*(1) + \frac{\varepsilon}{2}} \right)^{T^a} > a$. 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 21 it follows that $s \in \{\Omega_{1,t}^{T^a} \text{ i.o.}\} \quad P^{\theta^*} - \text{a.s. on } \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$. Therefore, $P^{\theta^*} - \text{a.s. on } \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, there exists a subsequence $\{t_k\}_{k=0}^\infty$ such that $s \in \Omega_{1,t_k}^{T^a}$ and so

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

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

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

Since a was arbitrarily chosen, it follows that

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

as desired. ■

Lemma 23 *Suppose A0 holds and that agent 1's prior satisfies A1 and A2.b. Then, $P^{\theta^*} - a.s.$ on $\left\{ \tilde{s} : \limsup \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} \leq 1 \right\}$, there is $T(s)$ such that $\frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \forall t \geq T(s)$.*

Proof of Lemma 23. Let $\Omega_2 \equiv \left\{ \tilde{s} : \limsup \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} \leq 1 \text{ and } \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$. Let $s \in \Omega_2$. Since $\Omega_2 \subset \left\{ \tilde{s} : \frac{P_{2,t}(\tilde{s})}{P_{1,t}(\tilde{s})} > 1 \text{ i.o.} \right\}$, then by Lemma 22

$$\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty \quad P^{\theta^*} - a.s. \text{ on } \Omega_2,$$

and it follows that $P^{\theta^*}(\Omega_2) = 0$, as desired. ■

Lemma 24 *Suppose A0 holds and that agent 1's prior satisfies A1 and A2.b. Then, $P^{\theta^*} - a.s.$, $\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ i.o.*

Proof of Lemma 24. Let $\Omega_3 \equiv \left\{ s : \exists T(s) \text{ such that } \frac{P_{2,t}(s)}{P_{1,t}(s)} \leq 1 \forall t \geq T(s) \right\}$. Then, for every $s \in \Omega_3$

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

By definition of agent 2's priors,

$$\prod_{k=T(s)}^t \frac{p_{2,k}(s)}{p_{1,k}(s)} = \prod_{k=T(s)}^t \frac{\theta_k^*(s)}{p_{1,k}(s)} \quad \forall s \in \Omega_3.$$

and it follows that

$$\Omega_3 \subset \left\{ s : \limsup_{T \rightarrow \infty} \prod_{t=1}^T \frac{\theta_t^*(s)}{p_{1,t}(s)} < +\infty \right\} \quad (38)$$

Since A.2 implies that P^{θ^*} is not absolutely continuous with respect to P_1 , it follows by Sandroni [19, 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 (38) implies that Ω_3 lies in a zero measure subset of Ω . It follows that, $P^{\theta^*} - a.s.$, $\frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$ 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 23 and Lemma 24, $P^{\theta^*} - a.s.$, $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} > 1$. On the other hand, by Lemmas 22 and 24 one concludes that, $P^{\theta^*} - a.s.$, $\limsup \frac{P_{2,t}(s)}{P_{1,t}(s)} = +\infty$. ■

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 [2009]: “The Market Organism: Long Run Survival in Markets with Heterogeneous Traders.” *Journal of Economic Dynamics and Control*, Vol. 33, N°3 (May), pp. 1023-1035.
- [6] BOYD, J. [1990], “Recursive Utility and the Ramsey Problem”, *Journal of Economic Theory*, 50, pp. 326-345.
- [7] 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.
- [8] COGLEY, T. AND T. SARGENT [2009]: “Diverse Beliefs, Survival, and the Market Price of Risk.” *Economic Journal*, Vol. 119 (March), 354-376.
- [9] CHUNG, K. L. [1974] “A Course in Probability Theory,” London, Acad. Press.
- [10] EASLEY, D. AND N. M. KIEFER [1988]: “Controlling a Stochastic Process with Unknown Parameters.” *Econometrica*, Vol. 56, No. 5, 1045-1064.
- [11] ESPINO, E. AND T. HINTERMAIER [2009] “Asset Trading Volume in a Production Economy.” *Economic Theory*, 39 (2), pp. 231-258.
- [12] 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.

- [13] FREEDMAN, D. A. [1975]: “On Tail Probabilities for Martingales.” *The Annals of Probability*, vol. 3, No. 1, 100-118.
- [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. AND N. STOKEY [1984]: “Optimal Growth with Many Consumers.” *Journal of Economic Theory*, 32, 139-171.
- [16] MORRIS, S. (1996): “Speculative Investor Behavior and Learning.” *Quarterly Journal of Economics*, 111, 1111-1133.
- [17] NEGISHI, T. [1960]: “Welfare Economics and the Existence of an Equilibrium for a Competitive Economy.” *Metroeconomica*, 7, 92-97.
- [18] PHILLIPS, P. C. B., AND W. PLOBERGER (1996): “An Asymptotic Theory of Bayesian Inference for Time Series.” *Econometrica*, 64(2), 381-412.
- [19] SANDRONI, A. (2000): “Do Markets Favor Agents Able to Make Accurate Predictions?” *Econometrica*, 68(6), 1303-42.
- [20] SHIRYAEV, A. (1991): *Probability*. Second Edition, New York: Springer Verlag.
- [21] SCHWARTZ, L. (1965): “On Bayes Procedures.” *Probability Theory and Related Fields*, 4(1), 10-26.