

BAYESIAN GAMES WITH A CONTINUUM OF STATES (R&R DRAFT)

ZIV HELLMAN

Department of Economics, Bar Ilan University, Ramat Gan, Israel

YEHUDA (JOHN) LEVY

*Department of Economics and Nuffield College, University of Oxford,
Oxfordshire, United Kingdom*

ABSTRACT. Negative results on the the existence of Bayesian equilibria when state spaces have the cardinality of the continuum have been attained in recent years. This has led to the natural question: are there conditions that characterise when Bayesian games over continuum state spaces have measurable Bayesian equilibria? We answer this in the affirmative. Assuming that each type has finite or countable support, Bayesian equilibria may fail to exist if and only if the underlying common knowledge σ -algebra is non-separable. Furthermore, anomalous examples with continuum state spaces have been presented in the literature in which common priors exist over entire state spaces but not over common knowledge components. There are also spaces over which players can have no disagreement, but when restricting attention to common knowledge components disagreements can exist. We show that when the common knowledge σ -algebra is separable all these anomalies disappear.

1. INTRODUCTION

What if we lived in a world in which Bayesian games were not guaranteed always to have Bayesian equilibria?

The effects might be felt widely throughout the literature, as it is difficult to exaggerate the importance which the concept of Bayesian games has attained in a wide range of subfields in economics and game theory, with subjects such as incomplete and asymmetric information models, signalling theory, principal-agent models, adverse selection and the provisions of public goods forming only a very partial list. Many papers in these fields start off by assuming the existence of an

The research of the first author was supported in part by the European Research Council under the European Commission's Seventh Framework Programme (FP7/2007 – 2013)/ERC grant agreement no. 249159, and in part by Israel Science Foundation grants 538/11 and 212/09. The research of the second author was supported in part by Israel Science Foundation grant 1596/10.

equilibrium and continuing their analyses from there. It would be challenging to gain significant theoretical traction, for example, in Bayesian truthful implementation and the related concepts of the revelation principle, the revenue equivalence theorem and optimal Bayesian methods, without first assuming that at least one Bayesian equilibrium exists in particular models being studied.

This isn't usually a concern at all, of course, since Harsányi (1967) proved (along with introducing the very concept of a Bayesian game) that every finite Bayesian game has an equilibrium. This positive result can easily be extended to Bayesian games over countably many states; e.g., Simon (2003).

Over a continuous state space, however, negative results have been shown in recent years. Simon (2003) presented an example of a three-player Bayesian game over a continuum state space with no Bayesian equilibrium.¹ Any hopes that positive results could be restored by considering approximate equilibria instead of exact equilibria were dashed when Hellman (2012b) showed an example of a two-player Bayesian game over a continuum state space with no Bayesian ε -equilibrium for $\varepsilon > 0$ small enough.

These negative results are perturbing. One on occasion hears it said that it is sufficient to concentrate on finite games alone because the world itself is finite. However, as pointed out in Cotter (1991), since there are an infinity of continuous random variables, a more accurate statement would be that decision makers observe only a finite number of variables, each to a finite degree of accuracy. To model this as a finite space, however, requires that the modeller know *a priori* the set of variables actually observed and the degree of accuracy of each observed variable. Given this, the use of infinite games reflects the modeller's ignorance of the decision-making environment, just as infinite horizon models are routinely used to reflect ignorance of the life-span of the decision maker.

Furthermore, limiting attention to finite Bayesian games is far from being sufficient for capturing the full range of possible models that need to be studied. Many models in the literature make use of continuous variables. Examples include models in which prices (as in models of auctions or bargaining, such as that of Chatterjee and Samuelson (1983) for example) are the main state variables, or in which the main variables are profits and outputs in market models (for example Radner (1980)), continuous time points, accumulated wealth, accumulated resources, population percentages, share percentages and so forth.

Furthermore, an extensively-used approach to dealing with a Bayesian game with a finite but large number of states is to analyse instead a similar game with a

¹ By the existence of an equilibrium we mean the existence of a measurable equilibrium. There are several reasons for restricting attention to measurable strategies (and hence measurable equilibria); to consider just two reasons, if a strategy is not measurable it cannot be constructed explicitly, and the payoffs of non-measurable strategies haven't got well-defined expected values. Measurability has in fact been included as a basic requirement in the definition of an equilibrium over uncountable spaces since the earliest literature on the subject (see Schmeidler (1973) for one such example). Throughout this paper we will therefore often use the term 'existence of an equilibrium' as synonymous with 'existence of a measurable equilibrium' without further qualification.

continuum of states. Myerson (1997), for example, informs readers of Chapter 2 of his textbook on game theory, when referring to Bayesian games, that ‘it is often easier to analyze examples with infinite type sets than those with large finite type sets’.

Given the negative results mentioned earlier, however, modellers working with continuum state spaces face the perhaps uncomfortable situation in which they may need to check, in each separate model with which they are working, whether or not an equilibrium exists. This motivates our main result here, which is exhibiting conditions that guarantee the existence of Bayesian equilibria in Bayesian games over a continuum of states, restoring the confidence in the existence of equilibria in the class of games satisfying these conditions.

In most of the paper we assume that in the Bayesian games under consideration every atom of each player’s posterior contains only a finite number of elements.² Note that every known example of a Bayesian game with no Bayesian equilibria satisfies this property, hence characterising conditions for the existence of equilibria in games with this property is of importance. Making this assumption is also concordant with some intuitions that although decision makers may *a priori* consider in their minds a continuum of possible states, for the purposes of observing a definite signal and moving to their posterior probabilities in most realistic cases they can only distinguish a finite number of posterior states to which they assign positive probability.

Part (I) of Theorem 2 then shows that if a Bayesian game satisfies the condition that the common knowledge σ -algebra of the underlying knowledge space is separable then there exists a Bayesian equilibrium. Furthermore, this condition is not only sufficient, it is also necessary in the following sense, as shown in part (II) of Theorem 2: if Ω is a standard Borel space and \mathcal{F} is a sub- σ -algebra of the Borel σ -algebra that is not separable and in which each atom is countable (and which can be generated by some beliefs), then there exists a Bayesian game Γ with state space Ω , a common prior, and common knowledge σ -algebra \mathcal{F} that does not possess an ε -MBE for small enough $\varepsilon > 0$.

The condition of separability of the common knowledge σ -algebra also turns out to be the crucial factor in resolving a series of disturbing ‘paradoxes’ in models over continuum state spaces. These are detailed in Section 3: there are Bayesian games over continuum state spaces that have no Bayesian equilibria, yet if these games are restricted to being played over any common knowledge component of the players, Bayesian equilibria do exist; there are type spaces with common priors such that when restricting to any common knowledge component the resulting type space has no common prior; there are type spaces that exclude any possibility of disagreement between the players, but again when restricting to any common knowledge component the resulting type space does admit disagreements.

These sorts of paradoxes are disturbing because they introduce instability in moving between the *ex ante* stage and the interim stage of analyses. Depending

² This condition can be weakened to countably many elements; see Section 10.

on the stage, one can get different answers to the questions of whether or not there exist equilibria, common priors or disagreements. As shown in this paper, however, all of these paradoxes disappear if the underlying knowledge spaces satisfy the condition of separability of the common knowledge σ -algebra.

Finally, we note that our results are attained mainly using results from descriptive set theory. In fact, we have found that there are parallels between concepts used in game theory and descriptive set theory concepts that are surprisingly useful for arriving at conclusions in game theoretic models. At several points in the body of the paper we strive to make these parallels explicit. Hopefully, these sorts of parallels can be deepened in future research, leading to more new results.

2. PRELIMINARIES AND THE MODEL

2.1. A Few Preliminaries.

A *standard Borel* space is a topological space that is homeomorphic to a Borel subset of a Polish space.³ Whenever we refer to a σ -algebra on a standard Borel space, we mean a sub σ -algebra of the Borel σ -algebra.

For a standard Borel space X , let $\Delta(X)$ denote the space of regular Borel probability distributions on X , with the topology of weak convergence of probability measures, and let $\Delta_f(X) \subseteq \Delta(X)$ (resp. $\Delta_a(X) \subseteq \Delta(X)$) denote the subspace of finitely supported (resp. purely atomic) measures. $\Delta_f(X), \Delta_a(X)$ are Borel subsets⁴ of $\Delta(X)$.

If (Ω, \mathcal{B}) is a measurable space and \mathcal{F} is a sub- σ -algebra of \mathcal{B} , then - Blackwell and Ryll-Nardzewski (1963) - a proper regular conditional distribution (henceforth, proper RCD) given \mathcal{F} , is a mapping $t : \Omega \times \mathcal{B} \rightarrow [0, 1]$ such that:

$$\mu(B) = \int_{\Omega} t(x)(B) d\mu(x), \text{ for all } B \in \mathcal{B}$$

and

$$t(\omega)(A) = 1, \text{ if } \omega \in A \in \mathcal{F}$$

Note that in particular,

$$t(\omega)(T) = E_{\mu}[1_T \mid \mathcal{F}](\omega), \mu\text{-a.e. } \omega \in \Omega$$

In terms that may be more familiar for game theorists, a proper RCD t of a probability measure μ may be thought of as the posterior t of a prior μ with respect to a knowledge structure \mathcal{F} .

³ Equivalently, a measurable space (X, \mathcal{B}) is standard Borel if there exists a metric on X that makes it a complete separable metric space in such a way that \mathcal{B} is then the Borel sigma-algebra, i.e., the smallest σ -algebra containing the open sets.

⁴ $\Delta_f(X)$ can be viewed as $\cup_{n \in \mathbb{N}} \Delta_n(X)$, where $\Delta_n(X)$ consists of the probability measures supported on at most n points. $\Delta_n(X)$ can be viewed as the image in $\Delta(X)$ of $X^n \times \Delta_n$, where Δ_n is the n -simplex, under a finite-to-one map. Similarly, $\Delta_a(X)$ can be viewed as the image of $\{X \in X^{\mathbb{N}} \mid \forall n \neq m, x_n \neq x_m\} \times \{x \in \mathbb{R}^{\mathbb{N}} \mid x \geq 0, \sum_{n=1}^{\infty} x_n = 1\}$ under a countable-to-one map.

A very central concept in this paper is:

Definition 1. A σ -algebra \mathcal{F} on a Borel space Ω is *separable*⁵ if there is a countable⁶ collection of Borel subsets $\{B_n\}_{n \in \mathbb{N}}$ of Ω that generates \mathcal{F} ; that is, \mathcal{F} is the smallest σ -algebra such that $\{B_n\}_{n \in \mathbb{N}} \subseteq \mathcal{F}$.

2.2. Knowledge Spaces.

A *knowledge space* for a nonempty, finite set of players \mathcal{P} is given by a triple $(\Omega, \mathcal{P}, (\mathcal{F}^p)_{p \in \mathcal{P}})$, where Ω is a standard Borel space of *states*, and \mathcal{F}^p for each $p \in \mathcal{P}$ is a σ -algebra over Ω , called *p's knowledge σ -algebra*. Intuitively, the elements in \mathcal{F}^p represent the events that player p can identify, hence the name knowledge σ -algebra.

Let $\mathcal{F} := \bigcap_{p \in \mathcal{P}} \mathcal{F}^p$; that is, \mathcal{F} is the finest σ -algebra contained in all the players' knowledge σ -algebras. \mathcal{F} is called the *common knowledge σ -algebra* of the knowledge space. The elements of \mathcal{F} intuitively represent the events of which all the players can have common knowledge.

2.3. Type Spaces.

Fix a knowledge space $(\Omega, \mathcal{P}, (\mathcal{F}^p)_{p \in \mathcal{P}})$. For each $p \in \mathcal{P}$, a *type function* t^p is mapping $t^p : \Omega \rightarrow \Delta(\Omega)$ which is \mathcal{F}^p -measurable, satisfying $t^p(\omega)(A) = 1$ whenever $\omega \in A \in \mathcal{F}^p$. A triple $(\Omega, \mathcal{P}, (t^p)_{p \in \mathcal{P}})$ (with $(\mathcal{F}^p)_{p \in \mathcal{P}}$ understood implicitly, \mathcal{F}^p is the σ -algebra generated⁷ by t^p) is called a *type space*.

ASSUMPTION: Unless otherwise specified we will assume that $t^p(\omega)(\cdot) \in \Delta_f(\Omega)$,⁸ for all $p \in \mathcal{P}$ and all $\omega \in \Omega$. This assumption will only be relaxed in Section 10 and for one example in Section 3.2.

ASSUMPTION: Unless otherwise specified, we will henceforth assume⁹ that type spaces satisfy *positivity*, i.e., that $t^p(\omega)[\omega] > 0$ for all $p \in \mathcal{P}$ and $\omega \in \Omega$.

Type spaces that do not satisfy this condition are *non-positive*. We will offer justification in Proposition 8 as to why we restrict attention to positive type spaces. We note here that from the assumption that $t^p(\omega)(\cdot)$ is a measure with finite support for all $p \in \mathcal{P}$ and all $\omega \in \Omega$, and that t^p is positive, it follows that each atom¹⁰ of each player's knowledge σ -algebra \mathcal{F}^p , and hence also of the common knowledge σ -algebra \mathcal{F} , is countable. The atom of the common knowledge σ -algebra \mathcal{F} containing ω is called the *common knowledge component* containing ω , and is denoted $K(\omega)$.

⁵'Separable' is synonymous with 'countably generated'.

⁶In this paper, 'countable' refers both to finite cardinalities and to countably infinite cardinalities.

⁷The σ -algebra generated by a mapping $f : X \rightarrow Y$ between standard Borel spaces is $\{f^{-1}(B) \mid B \subseteq Y \text{ is Borel}\}$.

⁸Recall that $\Delta_f(\Omega)$ is the set of finitely supported measures over Ω .

⁹This assumption also appears in Samet (1998).

¹⁰An *atom* of a σ -algebra is an element of it which is non-empty and is not strictly contained in any other element.

A measure $\mu^p \in \Delta(X)$ such that t^p is a proper RCD for μ^p given \mathcal{F}^p is a *prior* for t^p . A *common prior* is a measure μ that is a prior for the type functions of all the players $p \in \mathcal{P}$.

Most game theory models¹¹ work with a special case of type spaces that are *partitionally generated*. In such models, each player p has a partition Π^p of Ω . (In our case, Π^p is the collection of atoms of \mathcal{F}^p .) That player's knowledge σ -algebra \mathcal{F}^p is the σ -algebra generated by Π^p .¹² A type function t^p is then defined by a Borel mapping $t^p : \Omega \times \mathcal{F} \rightarrow \mathbb{R}$ such that

- (a) for each $\omega \in \Omega$, $t^p(\omega)(\cdot) \in \Delta(\Omega)$,
- (b) if $\omega' \in \Pi^p(\omega)$ then $t^p(\omega')(\cdot) = t^p(\omega)(\cdot)$.

Intuitively, a type function t^p represents the probability distribution that player p ascribes to the states conditional on receiving a signal that ω is a possible true state.

2.4. Bayesian Games & Bayesian Equilibrium.

A Bayesian game $\Gamma = (\Omega, \mathcal{P}, (t^p)_{p \in \mathcal{P}}, (I^p)_{p \in \mathcal{P}}, (r^p)_{p \in \mathcal{P}})$ consists of the following components:

- $(\Omega, \mathcal{P}, (t^p)_{p \in \mathcal{P}})$ forms a type space.
- I^p is a finite action set for each Player $p \in \mathcal{P}$.
- $r : \Omega \times \prod_{p \in \mathcal{P}} I^p \rightarrow \mathbb{R}^{\mathcal{P}}$ is a bounded measurable payoff function.

As usual, we extend r multi-linearly to $r : \Omega \times \prod_{p \in \mathcal{P}} \Delta(I^p) \rightarrow \mathbb{R}^{\mathcal{P}}$: That is,

$$r(\omega, (x^p)_{p \in \mathcal{P}}) = \sum_{(i^p)_{p \in \mathcal{P}} \in \prod_{p \in \mathcal{P}} I^p} \left(\prod_{p \in \mathcal{P}} x^p[i^p] \right) r(\omega, (i^p)_{p \in \mathcal{P}})$$

A *strategy* of a player $p \in \mathcal{P}$ is a mapping $\Omega \rightarrow \Delta(I^p)$ that is \mathcal{F}^p -measurable. A *measurable Bayesian ε -equilibrium* (ε -MBE), with $\varepsilon \geq 0$, is a profile of strategies $\sigma = (\sigma^p)_{p \in \mathcal{P}}$ such that for each $p \in \mathcal{P}$, each atom A of \mathcal{F}^p , and each $x \in \Delta(I^p)$,

$$\int_A r^p(\omega, \sigma(\omega)) dt^p(\omega) + \varepsilon \geq \int_A r^p(\omega, x, \sigma^{-p}(\omega)) dt^p(\omega)$$

When $\varepsilon = 0$ we will refer simply to an MBE instead of a 0-MBE.

3. THREE PARADOXES & EXAMPLES

The main motivation for the results of this paper is exhibiting conditions only on the common knowledge structure equivalent to the existence of Bayesian equilibria in games over continuum many states. As further motivation, in this section we present ‘three paradoxes’ related to games and type spaces over continuum many

¹¹ This can be broadened to: nearly all models in the economics, game theory and decision theory literature.

¹² By the σ -algebra generated by Π^p , we mean the collection of Borel sets which contain all those elements of Π^p that they intersect.

states, and several more examples for contrast. The results in this paper characterise when these paradoxes may hold and when they are guaranteed not to exist.

3.1. Paradoxes. The first two paradoxes, on Bayesian games and common priors in spaces over continuum many states, have been well-known in the literature for about a decade. The third paradox, on no betting, is fairly new new.

The “Now You See It, Now You Don’t” Bayesian Equilibrium.

Simon (2003) and Hellman (2012b) present examples of Bayesian games that have no Bayesian equilibria. In greater detail, let Γ be one of these Bayesian games, with state space Ω . Then there exists no vector of measurable strategies $(\varphi_1, \dots, \varphi_n)$, one per player, that forms a Bayesian equilibrium.

However, in both cases, one can choose any $\omega \in \Omega$ and consider the common knowledge component of $K(\omega)$, the atom of \mathcal{F} containing ω . (as determined by the partitions of the players). Let $\Gamma|_{K(\omega)}$ be the Bayesian game derived by restricting¹³ Γ to the states in $K(\omega)$. Then there *is* a Bayesian equilibrium of $\Gamma|_{K(\omega)}$, since this component is countable.

The “Now You See It, Now You Don’t” Common Prior.

This paradox was first noted in Simon (2000). We present here a slight variation of a version appearing in Lehrer and Samet (2011).

Consider the following type space over a state space Ω , as depicted in Figure 3.1. Ω is constructed out of four disjoint subsets of \mathbb{R}^2 , labelled A_j for $j \in \{1, 2, 3, 4\}$:

- $A_1 = \{(x, x+1) \mid -1 \leq x < 0\}$
- $A_2 = \{(x, x) \mid -1 \leq x < 0\}$
- $A_3 = \{(x, x-1) \mid 0 \leq x \leq 1\}$
- $A_4 = \{(x, \psi(x)) \mid 0 \leq x \leq 1\}$, where $\psi(x) = x - c \pmod{1}$ for a fixed irrational c in $(0, 1)$.

The knowledge space is partitionally generated, with Π_1 and Π_2 respectively the partitions of the two players. Player 1 is informed of the first coordinate of the state and player 2 is informed of the second coordinate. Thus, each element of $\Pi_1(\omega)$ is composed of the two points on the vertical line that contains the state ω . Similarly, $\Pi_2(\omega)$ contains the two points on the horizontal line that includes the state ω .

The posterior $t^p\omega$ for each of the two points in $\Pi_i(\omega)$ is $\frac{1}{2}$. Furthermore, let μ be the probability measure $\frac{1}{4} \sum_{j=1}^4 \psi_j$, where ψ_j is the Lebesgue measure over A_j . Lehrer and Samet (2011) show that measurability conditions are satisfied by the posteriors and that μ is a common prior for $t^p\omega$.

However, although the entire space Ω has a well-defined common prior, if we again concentrate on the common knowledge component $K(\omega_0)$ of any arbitrary state ω_0 (fixing the posteriors) then there is *no* common prior¹⁴ over $K(\omega_0)$. The

¹³If one restricts the type functions and payoffs to a set which is common knowledge - that is, in the common-knowledge σ -algebra \mathcal{F} , then the resulting game is well-defined.

¹⁴There may, however, be a *common improper prior* over $K(\omega_0)$. An improper prior allows for the possibility that the total measure it defines over a space diverges.




FIGURE 1. The state space consists of the three diagonals A_1 , A_2 , A_3 and of A_4 . The latter is obtained by a rightward shift of the top-right diagonal by an irrational number c .

reason for this is that $K(\omega_0)$ is a doubly infinite countable sequence

$$\dots, \omega_{-(k+1)}, \omega_{-k}, \dots, \omega_{-1}, \omega_0, \omega_1, \dots, \omega_k, \omega_{k+1}, \dots \quad (3.1)$$

such that $(\{\omega_k, \omega_{k+1}\}) \subseteq \Pi_1$ for all odd $k \geq 1$, $(\{\omega_k, \omega_{k-1}\}) \subseteq \Pi_1$ for all even $k \leq 0$, $(\{\omega_k, \omega_{k+1}\}) \subseteq \Pi_2$ for all odd $k \geq 0$, and $(\{\omega_k, \omega_{k-1}\}) \subseteq \Pi_2$ for all even $k \leq -1$. Any common prior ν over $K(\omega_0)$ must satisfy the condition that $\nu(\omega_k) = \nu(\omega_k + 1)$ for all k . Thus all the countably many states in $K(\omega_0)$ must have the same probability, which is impossible.

The “Now You See It, Now You Don’t” Acceptable Bet

For a type space with type spaces $(\Omega, \mathcal{P}, (t^p)_{p \in \mathcal{P}})$ a *bet* is a list of $(f_p)_{p \in \mathcal{P}}$ of bounded¹⁵ random variables $f_p : \Omega \rightarrow \mathbb{R}$. An *acceptable bet* is a bet that satisfies the condition that

$$E^p[f_p \mid \omega] := \int_{\Omega} f_p(s) dt^p(\omega)[s] > 0 \text{ for all } \omega \in \Omega. \quad (3.2)$$

By summing the integrals and integrating over the entire space, it’s clear that the existence of a common prior implies that there is no acceptable bet (see at the argument, e.g., in Hellman (2012a)). Since there is a common prior over the entire compact space in the example depicted in Figure 3.1, there can be no acceptable

¹⁵We assume boundedness to avoid anomalies; see Feinberg (2000).

bet over the entire space. By a result in Feinberg (2000) (see also Heifetz (2006), the converse is true if Ω is compact and we allow only continuous bets.

To show how in this example we can construct acceptable bets on each common knowledge component, once again we concentrate on a particular state ω_0 and the common knowledge component $K(\omega_0)$ containing it. We make use of a variation of a construction from Hellman (2012a) to define the following function $f : K(\omega_0) \rightarrow \mathbb{R}$ with $K(\omega_0)$ as in (3.1):

$$f(\omega_n) = \begin{cases} 1 & \text{if } n = 0 \\ 1 + \sum_{i=1}^n \frac{1}{2^i} & \text{if } n > 0 \text{ is even, or } n < 0 \text{ is odd} \\ -(1 + \sum_{i=1}^n \frac{1}{2^i}) & \text{if } n > 0 \text{ is odd, or } n < 0 \text{ is even} \end{cases}$$

It is easy to check that f is an acceptable bet over $K(\omega_0)$, even though there is no globally acceptable bet over the entire space Ω .

Several researchers have noted that these sorts of paradoxes strike at the heart of major assumptions underpinning research in contemporary economics and game theory, namely those related to the distinction between the *ex ante* stage and the interim stage of analysis. The full state space, over which priors are defined, is usually taken to be the *ex ante* stage while the common knowledge component represents the interim stage after each player receives a signal.

In many presentations, the Bayesian games and type spaces are considered ‘auxiliary constructions’ for the sake of analysis. According to this view, in reality there is no chance move that selects a player’s type; the knowledge and belief of each player determines his or her type. Type spaces and Bayesian games are merely ways to model the incomplete information each player has about the other players’ types. The true situation the players face is the interim stage after the vector of types has been selected. However, incomplete information requires us to consider the *ex ante* stage in order to understand how the players make their choices in the interim stage. Furthermore, in this view, in dynamic situations in which several signals may be received in succession over time, with each such signal refining the information known to the players, what is the interim stage after the receipt of one signal may also be considered the *ex ante* stage with respect to subsequent signals that have not yet been received.

This very standard view is challenged by the paradoxes detailed in this section. They show that when the state space has the cardinality of the continuum there may be a disturbing instability as we move from *ex ante* to interim stages: is there or is there not a common prior? Is there or is there not a Bayesian equilibrium? Is there or is there not a possible disagreement?

Taken together, Theorem 1, Theorem 2, and Theorem 3 in this paper show that all three paradoxes essentially disappear when the underlying common knowledge σ -algebra is separable; in fact, these solutions lead to full characterisations of when these pathologies may occur.

3.2. More Examples. EXAMPLE #1: The state space is $\Omega = \mathbb{R}$ - corresponding to an amount - chosen via some common prior - say, a normal distribution.

Player 1 is told the absolutely value of the amount ω of the asset - but is not told whether it is positive or negative amounts, e.g., whether it is amount they are due or that they owe - while Player 2 is told the absolute value of $|\omega + 1|$ - e.g., if the agent in charge of informing Player 2 what the amount is, in addition to forgetting to specify whether it is an amount due or owed, always overestimates the amount by 1 unit.

It is easy to verify that the atoms of the common knowledge σ -algebra are of the form $\{x + n \mid n \in \mathbb{Z}\}$. Indeed, if τ_1 denotes the generator of Player 1's uncertainty ($\tau_1(x) = -x$) and τ_2 denotes the generator of Player 2's uncertainty ($\tau_2(x) = 1 - x$), then $\tau_1 \circ \tau_2(x) = x - 1$ so $T := \tau_1 \circ \tau_2$ is the shift. The common knowledge σ -algebra, in this case, consists of all those Borel sets F which satisfy $x \in F \rightarrow x + n \in F$ for all $n \in \mathbb{Z}$, or equivalently, $T(F) = F$.

We claim first that this σ -algebra is separable: As the useful Proposition 3 below shows, this is equivalent to the existence of a Borel set $B \subseteq \Omega$ which intersects each atom of \mathcal{F} in exactly one point. In our case, simply take $B = [0, 1)$.

Theorems 1, 2, and 3 then guarantee us that common priors exist on all components, Bayesian equilibria exist regardless of the payoffs, and there are no agreeable bets on any components. Finding the common priors and Bayesian equilibria, however, can be quite cumbersome, and hence the advantage of possessing general existence theorems such as those we present here.

In fact, the last example can be generalised to show that our results apply in many situations, via Theorem 4, which gives a criterion for separability.

EXAMPLE #2: Now, to the game in the previous example, add a 3rd player. We will allow ourselves to deviate slightly from the usual framework allow this player's belief to be supported on a infinite (but countable) number of points; as we explain later in Section 10.2, this more general framework can also be handled by our results. This 3rd player is informed of some amount but he believes that the actual number may be some integer multiple or divisor of what he is told (for example, he may not be sure if this is the total amount, or the amount per person in a group of unknown size after the amount has been divided). Hence, the atoms of his knowledge are generated by multiplication by a positive integer - and in particular, the atoms of his knowledge are of the form $\{\dots, \frac{1}{3}x, \frac{1}{2}x, x, 2x, 3x, \dots\}$. When combined with the uncertainty of Players 1, 2 - which, as we saw, are generated by addition and subtraction by an integer - we see that in this common knowledge σ -algebra \mathcal{G} , x, y are in the same atom iff $x - y$ is rational. It is well known that there cannot exist a set B which intersects each atom in only finitely many points, e.g., Chapter 2 Rudin (1986).

Hence, this common knowledge structure is not separable, and as a result, common knowledge components will not possess common priors, Bayesian equilibria will not exist for certain priors and payoffs, and for certain type structures - even those induced by a common prior - we will find acceptable bets on components.

4. RESULTS

4.1. Common Priors over Components. Let $\tau = (\Omega, \mathcal{P}, (t^p)_{p \in \mathcal{P}})$ be a type space with common knowledge σ -algebra \mathcal{F} . If $K \in \mathcal{F}$ then $\tau_K := (K, \mathcal{P}, (t^p|_K)_{p \in \mathcal{P}})$, consisting of the state space K and the type functions restricted to K , is a well-defined type space. This is true in particular if K is an atom of \mathcal{F} . Furthermore, if μ is a common prior, then we say that a property holds for *almost every common knowledge component* if the set of components for which it does not hold are all contained in a set of measure zero.

Theorem 1 essentially states that given a type space τ with a common prior, the type space τ_K for any common knowledge component K is guaranteed also to have a common prior if and only if the underlying common knowledge σ -algebra is separable almost everywhere.

Theorem 1. *Let τ be a type space with a common prior μ . The following conditions are equivalent:*

- (1) *There exists $X \in \mathcal{F}$ with $\mu(X) = 1$ such that $\mathcal{F}|_X$ is separable.*
- (2) *For almost every common knowledge component K , the type space τ_K has a common prior.*
- (3) *There is a proper regular conditional probability t of μ given \mathcal{F} such that for almost every common knowledge component K and each $x \in K$, $t(x)$ is a common prior for τ_K .*

Remark 2. In particular, it follows that the common knowledge σ -algebra \mathcal{F} generated in Figure 3.1 is not separable. This, however, could be seen by more elementary means: the restriction of \mathcal{F} to any one of the sets A_1, A_2, A_3, A_4 is easily seen to be induced by the equivalence relation induced by an irrational rotational of the circle - i.e., $x \rightarrow x - c \bmod 1$, c being irrational – and this σ -algebra is well-known to be non-separable.

The proof can be found in Section 6

4.2. Bayesian Equilibria. Theorem 2 essentially states that given a type space τ is guaranteed to have a Bayesian equilibrium if and only if the underlying common knowledge σ -algebra is separable almost everywhere.

Theorem 2.

- I. *Let Γ be a Bayesian game in which the common knowledge σ -algebra is separable. Then there exists an MBE for Γ .*
- II. *Let Ω be a standard Borel space, and let \mathcal{F} be sub- σ -algebra of the Borel σ -algebra that is not separable and which is belief induced. Then there exists a Bayesian game Γ with state space Ω , a common prior, and common knowledge σ -algebra \mathcal{F} that does not possess an ε -MBE for small enough $\varepsilon > 0$, and in particular does not possess an MBE.*

To prove Theorem 1.I, we proceed in three steps. First, we will develop a notion of the space of (not necessarily positive) Bayesian games with countably

many states S , player set \mathcal{P} and action sets $(I^p)_{p \in \mathcal{P}}$, which we will denote by $\mathfrak{B}(S, \mathcal{P}, (I^p)_{p \in \mathcal{P}})$ (or just \mathfrak{B} for short). Afterwards, we will prove the existence of a Bayesian equilibrium selection for this class of games. Then we will show how one can measurably map the games induced on each common knowledge component of a general game into the space of games on countably many states S ; the composition of this mapping and the Bayesian equilibrium selection from the second step will give us the required global Bayesian equilibrium. We can construct such a mapping because the separability, it turns out, allows us to measurably enumerate the elements of each atom, and once we have this enumeration we can map the game on each atom to its appropriate game in the space \mathfrak{B} ; when we lack such an enumerate, this cannot be done because we have no canonical way to select the mapping. Details are given in Section 7.

4.3. No Betting. Theorem 3 essentially states that given a type space τ , we are guaranteed that the common knowledge components possess no acceptable bets - effectively, no acceptable bets at the interim stage - if and only if the underlying common knowledge σ -algebra is separable almost everywhere.

Theorem 3.

- I. *Let Γ be a Bayesian game with a common prior in which the common knowledge σ -algebra is separable. Then for almost every common knowledge component K , there are no acceptable bets on K .*
- II. *Let Ω be a standard Borel space, and let \mathcal{F} be sub- σ -algebra of the Borel σ -algebra that is not separable and which is belief induced. Then there exists a Bayesian game Γ with state space Ω , a common prior, and common knowledge σ -algebra \mathcal{F} such that on almost every common knowledge component, there exists an acceptable bet.*

The proof of Theorem 3.I is given in Section 6; the proof of Theorem 3.II is given in Section 8.

4.4. A Condition for Separability.

Theorem 4. *Suppose Ω is a standard Borel space with metric d , and \mathcal{F} is a sub- σ -algebra of the Borel σ -algebra, and suppose that for each atom A in \mathcal{F} , A is countable and $\inf_{x, y \in A} d(x, y) > 0$. Then \mathcal{F} is separable.*

In other words, as long as in each atom the elements 'keep their distance' and don't get 'bunched up', the σ -algebra is separable. The proof appears in Section 5.4.

5. PRELIMINARIES TO PROOFS

5.1. Descriptive Set Theory Preliminaries. A relation \mathcal{E} on a standard Borel space Ω is said to be *Borel* if it is Borel as a subset of $\Omega \times \Omega$; i.e., if the set $\{(x, y) \in \Omega \mid x \mathcal{E} y\}$ is Borel. A Borel equivalence relationship is said to be *countable* if each equivalence class is countable.

If \mathcal{E} is a Borel equivalence relationship on a space Ω , then for each $\omega \in \Omega$, $[\omega]_{\mathcal{E}}$ (or just $[\omega]$ when it is clear to which relationship we are referring) denotes the equivalence class of ω .

Given a σ -algebra \mathcal{F} , there is an *induced equivalence relation*, denoted $\mathcal{E}_{\mathcal{F}}$, defined by

$$[\omega]_{\mathcal{E}_{\mathcal{F}}} := [\omega]_{\mathcal{F}} := \bigcap_{A \in \mathcal{F}} A$$

Given a Borel equivalence relationship \mathcal{E} on Ω and a set $T \subseteq \Omega$, the *saturation* $[T]_{\mathcal{E}}$ of T w.r.t. \mathcal{E} is $[T]_{\mathcal{E}} = \bigcup_{\omega \in T} [\omega]_{\mathcal{E}}$. We will sometimes write $[T]_{\mathcal{F}}$ instead of $[T]_{\mathcal{E}_{\mathcal{F}}}$. Conversely, if \mathcal{E} is a Borel equivalence relationship, the induced σ -algebra $\mathcal{F}_{\mathcal{E}}$ is the collection of saturated Borel sets.

In terms that may be more familiar to game theorists used to working with finite atomic partitions as bases for σ -algebras, $[\omega]_{\mathcal{F}}$ is the atom containing ω and for an event T , $[T]_{\mathcal{F}}$ is the union of the atoms intersecting T .

If \mathcal{F} is a σ -algebra on a standard Borel space such that each atom is countable, then it follows from the Lusin-Novikov theorem (see, for example, Theorem 18.10 of Kechris (1995) or Proposition 10 ahead) that the induced equivalence relationship $\mathcal{E}_{\mathcal{F}}$ is Borel and that the saturation of each Borel set is Borel.

A *transversal* of an equivalence relationship is a set that intersects each equivalence class in exactly one point. A Borel equivalence relationship \mathcal{E} is said to be *smooth* if there is a Borel mapping $\psi : \Omega \rightarrow X$, where X is some standard Borel space, such that $\psi(x) = \psi(y) \leftrightarrow x \sim y$.

Given a standard Borel space Ω and a sub- σ -algebra \mathcal{F} of the Borel σ -algebra, we let Ω/\mathcal{F} denote the quotient space whose elements are the equivalence classes induced by \mathcal{F} and the induced σ -algebra consists of precisely the images of the sets in \mathcal{F} under the quotient map.

We will make repeated use of the following proposition:

Proposition 3. *The following conditions are equivalent for a countable Borel equivalence relationship $\mathcal{E}_{\mathcal{F}}$ induced on Ω by a σ -algebra \mathcal{F} :*

- (a) \mathcal{F} is separable.
- (b) There is a Borel transversal for $\mathcal{E}_{\mathcal{F}}$.
- (c) The quotient space Ω/\mathcal{F} is standard Borel.
- (d) The equivalence relationship $\mathcal{E}_{\mathcal{F}}$ is smooth.

Proof. The equivalence (b) \iff (c) \iff (d) is stated in Propositions 6.3 and 6.4 of Kechris and Miller (2004). If (c) holds and Λ is a countable collection of Borel sets generating the Borel structure on Ω/\mathcal{F} , the collection $\{q^{-1}(U) \mid U \in \Lambda\}$, where $q : \Omega \rightarrow \Omega/\mathcal{F}$ is the quotient map, generates \mathcal{F} , and hence (a) holds.

Now, suppose (a) holds; let $B_1, B_2, \dots \in \mathcal{F}$ generate \mathcal{F} . The map $p : \Omega \rightarrow 2^{\mathbb{N}}$ defined coordinate-wise by $p_n(\omega) = 1_{B_n}(\omega)$ is Borel and satisfies $p(x) = p(y)$ iff $x \mathcal{E}_{\mathcal{F}} y$, and hence $\mathcal{E}_{\mathcal{F}}$ is smooth. \square

The following is from Feldman (1977):

Proposition 4. *Let \mathcal{E} be a countable Borel equivalence relationship on a standard Borel space Ω . Then there is a countable group G of Borel bijections $\Omega \rightarrow \Omega$ such that for each $\omega \in \Omega$, $[\omega]_{\mathcal{E}} = \{g(\omega) \mid g \in G\}$.*

5.2. Preliminaries on Knowledge. Let τ be a type space with knowledge σ -algebras $(\mathcal{F}^p)_{p \in \mathcal{P}}$. For each $p \in \mathcal{P}$ and each set $N \subseteq \Omega$, let $K^p(N)$ denote the saturation of N w.r.t. \mathcal{F}^p , i.e., $K^p(N) = [N]_{\mathcal{F}^p}$. If $\omega \in \Omega$, write for short $K^p(\omega) = K^p(\{\omega\})$. Since the saturation of countable Borel sets under a Borel equivalence relationship is also Borel, we have:

Lemma 5. *If N is Borel, then so is $K^p(N)$.*

For each finite sequence $\hat{p} = (p_1, \dots, p_k) \in \mathcal{P}^* := \cup_{n \geq 0} \mathcal{P}^n$ and $N \subseteq \Omega$, let

$$K^{\hat{p}}(N) = K^{p_k} \left(K^{p_{k-1}} \left(\dots (K^{p_1}(N)) \dots \right) \right)$$

and $K^{\hat{p}}(\omega) = K^{\hat{p}}(\{\omega\})$. Then, define

$$K^{\infty}(N) = \cap_{\hat{p} \in \mathcal{P}^*} K^{\hat{p}}(N)$$

We will say that a countable Borel equivalence \mathcal{E} (or, equivalently, the σ -algebra it induces) is *belief induced* if there are finitely many equivalence relationships $\mathcal{E}_1, \dots, \mathcal{E}_n$ refining it ($\mathcal{E}_k \subseteq \mathcal{E}$ for $k = 1, \dots, n$) such that each \mathcal{E}_k has finite equivalence classes and \mathcal{E} is the finest common coarsening of $\mathcal{E}_1, \dots, \mathcal{E}_n$, i.e., \mathcal{E} is the transitive closure of $\cup_{k=1}^n \mathcal{E}_k$. This is equivalent to saying that \mathcal{E} is the common knowledge equivalence relationship induced by some type space (with finitely supported types). Not all countable Borel equivalence relationships are belief induced; we elaborate in Appendix A.¹⁶

Lemma 6. *Let Ω be a standard Borel space, let \mathcal{G} be a σ -algebra with countable atoms, let $\mu \in \Delta(\Omega)$ and let t be an RCD for μ given \mathcal{G} that satisfies $t(\omega)[\omega] > 0$ for all $\omega \in \Omega$. Let $N \subseteq \Omega$ be a μ -measurable set satisfying $\mu(N) = 0$. Then there is $K \in \mathcal{G}$ with $N \subseteq K$ and $\mu(K) = 0$.*

Proof. For each $n \in \mathbb{N}$, define

$$N_n = \{\omega \in N \mid t(\omega)[\omega] \geq \frac{1}{n}\}$$

Let $K_n = [N_n]_{\mathcal{G}} \in \mathcal{G}$; in other words, $K_n = \cup_{\omega \in N_n} [\omega]$. For all $\omega \in \Omega$, if $\omega \notin K_n$ then

$$t(\omega)(K_n) = 0,$$

while if $\omega \in K_n$ then

$$t(\omega)(K_n) = t(\omega)[\omega] \leq n \cdot t(\omega)[\omega] \leq n \cdot t(\omega)(N_n),$$

Therefore

$$\mu(K_n) = \int_{\Omega} t(\omega)(K_n) d\mu(\omega) \leq n \cdot \int_{\Omega} t(\omega)(N_n) d\mu(\omega) = n \cdot \mu(N_n) = 0$$

and we can take $K = \cup_{n \in \mathbb{N}} K_n$. □

¹⁶ We are grateful to Benjamin Weiss for pointing this out to us.

Corollary 7. *Let τ be a positive type space with a common prior μ . Let $N \subseteq \Omega$ be a μ -measurable set satisfying $\mu(N) = 0$. Then there is $K \in \mathcal{F}$, the common-knowledge σ -algebra, with $N \subseteq K$ and $\mu(K) = 0$.*

To prove Corollary 7, one applies Lemma 6 inductively to show that for each $n \in \mathbb{N}$ and each $\hat{p} \in \mathcal{P}^n$, $K^{\hat{p}}(N)$ is Borel. Corollary 7 will often be used implicitly; in many proofs, when useful, we will automatically assume that some null set we are discarding is common knowledge – or, equivalently, and more to the point, that its complement is common knowledge.

Finally, we justify our concentration on positive type spaces. Proposition 8 essentially states that if a space has a common prior (whether or not positive) then under that prior the event containing the states to which any player p assigns zero probability in the posterior is a null event:

Proposition 8. *Let τ be a type space (not necessarily positive) with a common prior μ . Denote, for each $p \in \mathcal{P}$,*

$$N^p := \{\omega \in \Omega \mid t^p(\omega)[\omega] = 0\}$$

Then $\mu(N^p) = 0$ for all $p \in \mathcal{P}$.

Proof. Recall that each type function is \mathcal{F}^p -measurable. Let $x \in N^p$ and $\omega \in \Omega$. If ω is not in the same atom of x then $t^p(x)[\omega] = 0$ since ω is not in the support of $t^p(x)$. Otherwise, since $t^p(x)[x] = 0$ and $t^p(x)[\omega] = t^p(x)[x]$, we again conclude that $t^p(x)[\omega] = 0$. The proposition follows from the definition of an RCD. \square

We will also need:

Proposition 9. *Let $\mathcal{F}^1, \dots, \mathcal{F}^p$ be knowledge σ -algebras with finite atoms on a standard Borel space Ω , and let $\mu \in \Delta(\Omega)$. Let $\varepsilon > 0$. Then there exists a positive type space τ with common prior ν with these knowledge structures $\mathcal{F}^1, \dots, \mathcal{F}^p$, and such that $\|\nu - \mu\| \leq \varepsilon$, the distance being the total variation norm.*

The proof appears in Appendix A.

5.3. Some More Descriptive Set Theory Theorems. The following variants are slight strengthening of the Lusin-Novikov theorem (see, for example, Theorem 18.10 of Kechris (1995)):

Proposition 10. *Given a σ -algebra \mathcal{F} that induces a countable Borel equivalence relationship on a standard Borel space Ω , there exist:*

- I. Partial¹⁷ Borel mappings g_1, g_2, \dots , from Ω/\mathcal{F} to Ω such that for all $q \in \Omega/\mathcal{F}$, $q = \cup_{\omega \in \text{dom}(g_n)} \{g_n(\omega)\}$, $g_n(q) \neq g_m(q)$ whenever $n \neq m$ and $\omega \in \text{dom}(g_n) \cap \text{dom}(g_m)$, and if $q \notin \text{dom}(g_k)$, then $q \notin \text{dom}(g_n)$ for all $n > k$.*
- II. Partial Borel mappings f_1, f_2, \dots , from Ω to Ω such that for all $\omega \in \Omega$, $[\omega]_{\mathcal{F}} = \cup_{\omega \in \text{dom}(f_n)} \{f_n(\omega)\}$, $f_n(\omega) \neq f_m(\omega)$ whenever $n \neq m$ and $\omega \in \text{dom}(f_n) \cap$*

¹⁷ A mapping from a certain domain is called partial if it is defined only on a subset of the domain; it follows that the domain, as in the inverse image of the entire range space, is Borel.

$\text{dom}(f_m)$, and if $\omega \notin \text{dom}(g_k)$, then $q \notin \text{dom}(g_n)$ for all $n > k$.

The second part follows from the first by taking $f_n(\cdot) = g_n([\cdot]_{\mathcal{F}})$, since the mapping $\Omega \rightarrow \Omega/\mathcal{F}$ assigning each element to its equivalence class is easily seen to be Borel when \mathcal{F} is smooth.

If (Ω, \mathcal{E}) , (Λ, \mathcal{D}) are standard Borel spaces with Borel equivalence relations \mathcal{E} and \mathcal{D} induced on them, (Ω, \mathcal{E}) is said to *embeddable* into (Λ, \mathcal{D}) if there is an injective Borel mapping $\psi : \Omega \rightarrow \Lambda$ such that for all $\omega, \eta \in \Omega$, $\omega \mathcal{E} \eta \iff \psi(\omega) \mathcal{D} \psi(\eta)$; in this case, we denote $(\Omega, \mathcal{E}) \sqsubset (\Lambda, \mathcal{D})$.

A countable Borel equivalence relationship is said to be *hyperfinite* (Dougherty et al. (1994)) if it is induced by the action of a Borel \mathbb{Z} -action on Ω ; i.e., if there is a bijective¹⁸ Borel mapping $T : \Omega \rightarrow \Omega$ such that $x \mathcal{E} y \iff \exists n \in \mathbb{Z}, T^n(x) = y$.

Proposition 11. *Let $\mathcal{E}_1, \mathcal{E}_2$ be non-smooth countable Borel equivalence relationships on standard Borel spaces Ω_1, Ω_2 , with \mathcal{E}_1 being hyperfinite. Then $(\Omega_1, \mathcal{E}_1) \sqsubset (\Omega_2, \mathcal{E}_2)$.*

Proof. Let \mathcal{E}_t be the tail equivalence relationship on $C = 2^{\mathbb{N}}$; i.e., if $S : C \rightarrow C$ is defined by $(Sx)_n = x_{n+1}$, then $x \mathcal{E}_t y$ iff $\exists k, m \in \mathbb{N}$ such that $S^k(x) = S^m(y)$; \mathcal{E}_t is non-smooth and hyperfinite, see (Dougherty et al., 1994, Sec. 6). By the Glimm-Effros dichotomy for countable Borel equivalence relationships, Harrington et al. (1990), since \mathcal{E}_2 is not smooth,¹⁹ $(C, \mathcal{E}_t) \sqsubset (\Omega_2, \mathcal{E}_2)$; denote such an embedding by θ . By Theorem 7.1 of Dougherty et al. (1994), any two non-smooth hyperfinite equivalence relationships can be embedded into each other, and (C, \mathcal{E}_t) is known to be hyperfinite (Dougherty et al., 1994, Ch.6), hence $(\Omega_1, \mathcal{E}_1) \sqsubset (C, \mathcal{E}_t)$; denote such an embedding φ . This yields $\psi = \varphi \circ \theta$ as the required embedding. \square

5.4. Proof of Theorem 4. Let $(V_n)_{n \in \mathbb{N}}$ be a countable basis for Ω , and let G be a countable group which induces the equivalence relationship $\mathcal{E} = \mathcal{E}_{\mathcal{F}}$ induced by the σ -algebra \mathcal{F} as in Proposition 4. For each $n \in \mathbb{N}$, define $f_n : \Omega \rightarrow \{0, 1, 2, \dots, \infty\}$ by

$$f_n(\omega) = |\{V_n \cap [\omega]_{\mathcal{F}}\}|$$

We will show that f_n is Borel: Fix some well-ordered $>$ on G , and define

$$\psi_g(\omega) = 1_{g(V_n)}(\omega) \prod_{g' < g} (1 - 1_{g^{-1}(\cdot)=g'^{-1}(\cdot)}(\omega))$$

where 1_A is the indicator function of A ; that is, $\psi_g(\omega)$ indicates whether $g^{-1}(\omega) \in V_n$ and, in addition, that this point has not appeared before for previous elements of G acting on ω . Hence, $f_n = \sum_{g \in G} \psi_g$. So each $n \in \mathbb{N}$, f_n is Borel and

¹⁸ If a Borel mapping between standard Borel spaces is injective, a theorem by Kuratowski states that its image is standard Borel and that its inverse is Borel.

¹⁹ The Glimm-Effros dichotomy is usually stated for a state space Ω that is Polish; however, a Borel space can always be endowed with a Polish topology inducing the same Borel structure, since all standard Borel spaces are Borel isomorphic.

our assumption of $\inf_{x,y \in A} d(x,y) > 0$ for each atom A of \mathcal{F} is easily seen to imply that for each ω , $E_n(\omega) = \{n \mid f_n(\omega) = 1\}$ is non-empty, and the function $m(\cdot) = \min E_n(\cdot)$ is measurable. Hence, the correspondence

$$\Psi(\omega) = \{g(\omega) \mid g \in G\} \cap V_{m(\omega)} = \bigcup_{g \in G} (\{g(\omega)\} \cap V_{m(\omega)})$$

has a Borel graph and, by the definition of $m(\cdot)$, satisfies $|\Psi(\omega)| = 1$ for all ω ; hence, Ψ is a measurable function. Furthermore, Ψ is constant on each atom, and therefore \mathcal{F} is smooth.

6. PROOFS OF THEOREM 1 AND THEOREM 3.I

Lemma 12. *The mapping $\Delta(\Omega) \times \Omega \rightarrow \mathbb{R}$ given by $(\nu, \omega) \rightarrow \nu(\{\omega\})$ is Borel.²⁰ (REFERENCE NEEDED.)*

Lemma 13. *The correspondence*

$$\Psi(\omega) = \{\nu \in \Delta_a(\Omega) \mid \nu(K(\omega)) = 1 \text{ and } \nu|_{K(\omega)} \text{ is a common prior for } \tau_{K(\omega)}\}$$

has a Borel graph, and $|\Psi(\omega)| \leq 1$ for all $\omega \in \Omega$; hence, Ψ is a partial function.

Proof. The fact that $|\Psi(\omega)| \leq 1$ (i.e., that on a countable space in which no proper non-empty subset is common knowledge there exists at most one common prior) follows from Proposition 3 of Hellman and Samet (2011).

Let $(f_n)_{n \in \mathbb{N}}$ be as in Proposition 10. Define for each $n \in \mathbb{N}$, $g_n : \Delta_a(\Omega) \times \Omega \rightarrow [0, 1]$ by:

$$g_n(\nu, \omega) = \begin{cases} \nu(\{f_n(\omega)\}) & \text{if } \omega \in \text{Dom}(f_n) \\ 0 & \text{if } \omega \notin \text{Dom}(f_n) \end{cases}$$

and for each $m, n \in \mathbb{N}$ and $p \in \mathcal{P}$, define

$$D_{n,m} = \Delta_a(\Omega) \times (\text{dom}(f_n) \cap \text{dom}(f_m))$$

and

$$H_{n,m}^p = \{\omega \in D_{n,m} \mid t^p(f_n(\omega))[f_m(\omega)] \cdot \nu(K^p(f_m(\omega))) = \nu(\{f_m(\omega)\})\}$$

Each $D_{n,m}$ and $H_{n,m}$ is Borel - to see this, note that the mapping $\omega \rightarrow t^p(\omega)$ is Borel, and

$$\nu(K^p(f_n(\omega))) = \sum_{m \text{ s.t. } f_m(\omega) \in K^p(f_n(\omega))} \nu(\{f_m(\omega)\})$$

Finally,

$$\begin{aligned} \Psi(\omega) = \{ \nu \in \Delta_a(\Omega) \mid & \left(\sum_{n=1}^{\infty} g_n(\nu, \omega) \in \{0, 1\} \right) \\ & \wedge_{p \in \mathcal{P}} \wedge_{n,m \in \mathbb{N}} ((\nu, \omega) \notin D_{n,m} \vee (\nu, \omega) \in H_{n,m}^p) \} \end{aligned}$$

□

²⁰Recall that $\Delta(\Omega)$ is endowed with the topology of narrow convergence of probability measures.

Proof. (of Theorem 1.) Clearly, property (3) implies property (2). Suppose (2) holds; then, for Ψ as in Lemma 13, $\Psi(\omega) = 1$ for μ -a.e. $\omega \in \Omega$. Hence, after restricting Ψ to some $X \in \mathcal{F}$ of full μ -measure, the graph of Ψ defines a Borel function $\psi : X \rightarrow \Delta_a(\Omega)$, which clearly satisfies $K(x) = K(y) \iff \psi(x) = \psi(y)$; hence $\mathcal{F}|_X$ is smooth.

Finally, assume property (1) holds, and assume w.l.o.g., $\Omega = X$. By²¹ Theorem 1 of Blackwell and Ryll-Nardzewski (1963), there is a μ -a.e. proper RCD t for μ given \mathcal{F} . The claim that t is a common prior on μ -a.e. component follows now from Proposition 14 below. \square

Proof. (Proof of Theorem 3.I) By Theorem 1, almost every common knowledge component K has a common prior. This is sufficient, by Theorem 1.a. in Hellman (2012a), to conclude that there can be no acceptable bet over τ_K . \square

Proposition 14. *Let $\mathcal{E}, \mathcal{E}'$ be smooth countable Borel equivalence relationships on a standard Borel space Ω , with \mathcal{E}' refining \mathcal{E} (that is, $\mathcal{E}' \subseteq \mathcal{E}$) let μ be a regular Borel probability measure on Ω , and let t, t' be proper RCD's of μ w.r.t. the σ -algebras $\mathcal{F}, \mathcal{F}'$ induced by $\mathcal{E}, \mathcal{E}'$, respectively. Then for μ -a.e. $\omega \in \Omega$ and \mathcal{E}' -equivalence class C' with $\omega \in C'$,*

$$t(\omega)(\cdot \mid C') = t'(\omega)(\cdot) \quad (6.1)$$

Proof. It suffices to show that for μ -a.e. $\omega \in \Omega$ and each \mathcal{E}' -equivalence class C' such that $\omega \in C'$,

$$t(\omega)(\{\omega\} \mid C') = t'(\omega)[\omega]$$

Indeed, this suffices since both t, t' are constant in each \mathcal{E}' -equivalence class, and both sides of (6.1) vanish for sets supported outside of C' . Since $\omega \in C'$, this is equivalent to showing that for μ -a.e. $\omega \in \Omega$ and such C' ,

$$t'(\omega)[\omega] \cdot t(\omega)(C') = t(\omega)[\omega] \quad (6.2)$$

Note that since $\mathcal{E}, \mathcal{E}'$ are smooth, the induced quotient spaces $\Omega/\mathcal{F}, \Omega/\mathcal{F}'$ are standard Borel by Proposition 3 and μ induces measures on these quotient spaces. Throughout this proof, it will be convenient to view t, t' as functions on $\Omega/\mathcal{F}, \Omega/\mathcal{F}'$. - i.e., to view the RCD's as a function of the equivalence class, not of its elements.

Lemma 15. *For any bounded real-valued random variable X on (Ω, μ) ,*

$$\int_{\Omega} X(\omega) d\mu(\omega) = \int_{\Omega/\mathcal{F}} \left(\sum_{\omega \in C} X(\omega) \cdot t(C)[\omega] \right) d\mu(C) \quad (6.3)$$

Proof. It suffices to verify (6.3) in the case $X = 1_A$, A being Borel, and then to use an approximation argument. In this case, the left-hand side of Equation (6.3)

²¹ The condition given there for the existence of proper RCD's is easily seen to follow from the existence of a Borel transversal, which – by Proposition 3 – follows from separability.

is just $\mu(A)$, while the other side is

$$\begin{aligned} \int_{\Omega/\mathcal{F}} \left(\sum_{\omega \in C} 1_A(\omega) \cdot t(C)[\omega] \right) d\mu(C) &= \int_{\Omega/\mathcal{F}} t(C)(A \cap C) d\mu(C) \\ &= \int_{\Omega/\mathcal{F}} t(C)(A) d\mu(C) \end{aligned}$$

In general, for an \mathcal{F} -measurable function $f : \Omega \rightarrow \mathbb{R}$ – which induces a measurable function $f : \Omega/\mathcal{F} \rightarrow \mathbb{R}$ – we have

$$\int_{\Omega/\mathcal{F}} f(C) d\mu(C) = \int_{\Omega} f(\omega) d\mu(\omega)$$

(again, one checks it first for simple \mathcal{F} -measurable functions) and in particular for $f(\cdot) = t(\cdot)(A)$. Hence,

$$\int_{\Omega/\mathcal{F}} t(C)(A) d\mu(C) = \int_{\Omega} t(\omega)(A) d\mu(\omega) = \mu(A)$$

as required. \square

Now, note that on Ω/\mathcal{F}' there is the equivalence relationship \mathcal{E}^* induced by \mathcal{E} ; that is, two elements of Ω/\mathcal{F}' are \mathcal{E}^* equivalent if they are subsets of the same equivalence class of \mathcal{E} . \mathcal{E}^* is easily seen to be Borel and smooth as well; denote its induced σ -algebra on Ω/\mathcal{F}' as \mathcal{F}^* . Let t^* denote the proper RCD of μ (as a measure on Ω/\mathcal{F}) w.r.t \mathcal{F}^* , which exists by²² Theorem 1 of Blackwell and Ryll-Nardzewski (1963).

Lemma 16. *For μ -a.e. $C \in \Omega/\mathcal{F}$ and each \mathcal{E}' -equivalence class $C' \subseteq C$,*

$$t^*(C)[C'] = t(C)(C')$$

Proof. For any bounded real-valued random variable X on $(\Omega/\mathcal{F}', \mu)$ (by abuse of notation, we let X also denote the induced \mathcal{F}' -measurable random variable defined on Ω), by repeated use of Lemma 15,

$$\begin{aligned} \int_{\Omega/\mathcal{F}} \left(\sum_{\omega \in C} X(\omega) \cdot t(C)[\omega] \right) d\mu(C) &= \int_{\Omega} X(\omega) d\mu(\omega) = \int_{\Omega/\mathcal{F}'} X(C') d\mu(C') \\ &= \int_{\Omega/\mathcal{F}} \left(\sum_{C' \in C} X(C') \cdot t^*(C)[C'] \right) d\mu(C) \end{aligned}$$

where the sum over $C' \in C$ is taken over \mathcal{E}' -equivalence classes. (The middle equality follows by definition for indicator functions.) However, for μ -a.e. $\omega \in \Omega$,

$$\sum_{\omega \in C} X(\omega) \cdot t(C)[\omega] = \sum_{C' \in C} X(C') \cdot t(C)(C')$$

²² See explanation and footnote when this result is used in the proof of Theorem 1.

Hence,

$$\int_{\Omega/\mathcal{F}} \left(\sum_{C' \in C} X(C') \cdot t(C)(C') \right) d\mu(C) = \int_{\Omega/\mathcal{F}} \left(\sum_{C' \in C} X(C') \cdot t^*(C)[C'] \right) d\mu(C)$$

and this holds for any bounded real-valued random variable X . \square

We now complete the proof. For any bounded real-valued random variable X on Ω , by Lemma 15 (applied first to the equivalence relationship \mathcal{E}' on Ω , and then to the equivalence relationship \mathcal{E}^* on Ω/\mathcal{F}'), and by Lemma 16,

$$\begin{aligned} \int_{\Omega} X(\omega) d\mu(\omega) &= \int_{\Omega/\mathcal{F}'} \left(\sum_{\omega \in C'} X(\omega) \cdot t'(C')[\omega] \right) d\mu(C') \\ &= \int_{\Omega/\mathcal{F}} \left(\sum_{\omega \in C'} X(\omega) \cdot t'(C')[\omega] \right) \sum_{C' \subseteq C} t^*(C)[C'] d\mu(C) \\ &= \int_{\Omega/\mathcal{F}} \left(\sum_{C' \subseteq C} \sum_{\omega \in C'} X(\omega) \cdot t(C)(C') \cdot t'(C')[\omega] \right) d\mu(C) \\ &= \int_{\Omega/\mathcal{F}} \left(\sum_{\omega \in C} X(\omega) \cdot t(C)([\omega]_{\mathcal{E}'}) \cdot t'([\omega]_{\mathcal{E}'})([\omega]) \right) d\mu(C) \end{aligned}$$

Comparing this to Equation (6.3), we see that for μ -a.e. $\omega \in \Omega$,

$$t([\omega]_{\mathcal{E}})([\omega]_{\mathcal{E}'}) \cdot t'([\omega]_{\mathcal{E}'})([\omega]) = t([\omega]_{\mathcal{E}})([\omega])$$

or, denoting $C' = [\omega]_{\mathcal{E}'}$ and recalling $t([\omega]_{\mathcal{E}}) = t(\omega)$, and similarly for \mathcal{E}' , t' , we deduce Equation (6.2). \square

7. PROOF OF THEOREM 2

Fix a countable set S and an element $s_0 \in S$. Let \mathfrak{B} denote the collection of all \mathcal{P} -tuples $(s^p, g^p)_{p \in \mathcal{P}}$ for which $(S, \mathcal{P}, (I^p)_{p \in \mathcal{P}}, (s^p, g^p)_{p \in \mathcal{P}})$ constitutes a Bayesian game, where (s^p) being the types, (g^p) being the payoffs. \mathfrak{B} is endowed with the topology of point-wise convergence²³: $(s_{\alpha}^p, g_{\alpha}^p)_{p \in \mathcal{P}} = \Upsilon_{\alpha} \rightarrow \Upsilon = (s^p, g^p)_{p \in \mathcal{P}}$ in \mathfrak{B} if for every player $p \in \mathcal{P}$, every $\omega \in S$, and every pure action profile $a \in \prod_{p \in \mathcal{P}} I^p$ in $g_{\alpha}^p(\omega, a) \rightarrow g^p(\omega, a)$ and $s_{\alpha}^p(\omega) \rightarrow s^p(\omega)$:

Proposition 17. *\mathfrak{B} is homeomorphic to a Borel subset $\Xi := ((S \times [0, 1])^* \times \mathbb{R}^{\prod_{p \in \mathcal{P}} I^p})^{S \times \mathcal{P}}$ and hence is standard Borel (where for a set A , $A^* = \bigcup_{n=0}^{\infty} A^n$ with each A^n being both closed and open).*

The simple intuition is that for each player and state pair $(\omega, p) \in S \times \mathcal{P}$, we need to specify both an element in $(S \times [0, 1])^*$ – a finite list of states that are in the same element of the knowledge partition as ω , and the probabilities themselves to

²³ We define the topology in terms of nets.

these states – as well as an element of $\mathbb{R}^{\prod_{p \in \mathcal{P}} I^p}$, which specifies what payoff that player will receive as a result of each possible action profile.

Although we will not need it, the proof shows this mapping can be chosen to be natural up to a choice of a well-ordering on S . Henceforth, we will identify \mathfrak{B} with some such fixed subset of Ξ .

Proof. Write $\mathfrak{B} = \prod_{p \in \mathcal{P}} (\mathfrak{B}_s^p \times \mathfrak{B}_g^p)$, where \mathfrak{B}_s^p (resp. \mathfrak{B}_g^p) denotes the projection of \mathfrak{B} to the space of types (resp. payoffs) for Player p , with the induced topologies. It's enough to show that \mathfrak{B}_s^p is homeomorphic to a Borel subset of $((S \times [0, 1])^*)^S$ and that \mathfrak{B}_g^p is homeomorphic to Borel subseteq of $\mathbb{R}^{S \times \prod_{p \in \mathcal{P}} I^p}$.

The latter claim is trivial once one notices that for any countable set C , the set of bounded functions in \mathbb{R}^C is Borel, as it can be written

$$\bigcup_{n \in \mathbb{N}} \bigcap_{c \in C} \{a \in \mathbb{R}^C \mid |a_c| \leq n\}.$$

and that the Tychonoff topology is indeed the required topology of point-wise convergence. We turn to the former claim. Fix some well-ordering $<$ on S . As mentioned above, the intuition describing the map from \mathfrak{B}_s^p to $((S \times [0, 1])^*)^S$ is the following: for each $\omega \in S$, the player has to specify the finite list of states he believes he could be in and the weight each one receives. Finite lists of states are ordered by $>$. Hence, the image of \mathfrak{B}_s^p under such a map is given by the subset of Ξ defined by three conditions: Being supported on finite sets, they have total mass of unity, and they are constant on the set they are supported on. Mathematically:

$$\begin{aligned} & \bigcap_{\omega \in S} \bigcup_{F \subseteq S, |F| < \infty} \bigcap_{x \notin F} \{s^p \in ((S \times [0, 1])^*)^S \mid s^p(\omega)(x) = 0\} \\ & \bigcap_{\omega \in S} \{s^p \in ((S \times [0, 1])^*)^S \mid \sum_{x \in S} s^p(\omega)[x] = 1\} \\ & \bigcap_{\omega, \eta, \zeta \in S} \{s^p \in ((S \times [0, 1])^*)^S \mid s^p(\omega)[\eta] > 0 \rightarrow s^p(\omega)(\zeta) = s^p(\eta)(\zeta)\} \end{aligned}$$

and, again the topology is the topology of point-wise convergence. \square

The space Σ^p of strategies for Player p on a countable space is clearly a compact subspace of $(\Delta(I^p))^S$, hence the space of strategy profiles $\Sigma = \prod_{p \in \mathcal{P}} \Sigma^p$ is a compact space.

Proposition 18. *The Bayesian equilibrium correspondence $BE : \mathfrak{B} \rightarrow \Sigma$ has a Borel graph and takes on compact non-empty values.*

Proof. The fact that every Bayesian game with a countable state space has at least one Bayesian equilibrium follows from standard fixed point arguments; see, e.g., Simon (2003). The fact that the set of Bayesian equilibrium is compact also follows by standard arguments. To show that the graph G of the BE correspondence is Borel, note that

$$\begin{aligned} G &= \{((s^p, g^p)_{p \in \mathcal{P}}, \sigma) \in \mathfrak{B} \times \Sigma \mid \forall \omega \in S, \forall p \in \mathcal{P}, \forall x \in \Delta_{\mathbb{Q}}(I^p), \\ & \quad \sum_{v \in S} 1_{v \in K^p(\omega)} \cdot g^p(v, \sigma(w)) s^p(\omega)[v] \geq \sum_{v \in S} 1_{v \in K^p(\omega)} \cdot g^p(v, x, \sigma^{-p}(w)) s^p(\omega)[v]\} \end{aligned}$$

where for a finite set A , $\Delta_{\mathbb{Q}}(A)$ denotes the probability distributions on a A which give rational weights to all points. \square

The following corollary then results from Proposition 18 and the selection theorem of Kuratowski and Ryll-Nardzewski (1965) (see also Himmelberg (1975)):

Corollary 19. *There exists a Borel mapping $\psi : \mathfrak{B} \rightarrow \Sigma$ such that for all $\Lambda \in \mathfrak{B}$, $\psi(\Lambda)$ is a Bayesian equilibrium of Λ .*

Given two Bayesian games

$$(S, \mathcal{P}, (I^p)_{p \in \mathcal{P}}, (s_S^p)_{p \in \mathcal{P}}, (g_S^p)_{p \in \mathcal{P}})$$

and

$$(T, \mathcal{P}, (I^p)_{p \in \mathcal{P}}, (s_T^p)_{p \in \mathcal{P}}, (g_T^p)_{p \in \mathcal{P}})$$

with countable state spaces and the same player and action sets, an *isomorphism* from S to T is a bijective mapping $\phi : S \rightarrow T$ such that:

- For all $\omega \in S$ and pure action profile x , $g_S(\omega, x) = g_T(\phi(\omega), x)$.
- For all $\omega, \eta \in S$ and $p \in \mathcal{P}$, $s_S^p(\omega)[\eta] = s_T^p(\phi(\omega))[\phi(\eta)]$.

Proposition 20. *Let $\Gamma = (\Omega, \mathcal{P}, (I^p)_{p \in \mathcal{P}}, (t_S^p)_{p \in \mathcal{P}}, (r_S^p)_{p \in \mathcal{P}})$ be a Bayesian game such that the common knowledge σ -algebra \mathcal{F} is separable and aperiodic,²⁴ and let $\mathfrak{B} = \mathfrak{B}(S, \mathcal{P}, (I^p)_{p \in \mathcal{P}})$ be the set of Bayesian games with countable state space S with the same player and action space as Γ . Then Ω/\mathcal{F} is standard Borel and there is a Borel map $\Phi : \Omega \rightarrow S$ which is \mathcal{F} -measurable and a Borel map $\Lambda : \Omega/\mathcal{F} \rightarrow \mathfrak{B}$ such that for each $\omega \in \Omega$, if we denote*

$$\Gamma_\omega = (K(\omega), \mathcal{P}, (I^p)_{p \in \mathcal{P}}, (t^p|_{K(\omega)})_{p \in \mathcal{P}}, (r^p|_{K(\omega)})_{p \in \mathcal{P}})$$

then $\Theta|_{K(\omega)}$ is an isomorphism of Γ_ω to $\Lambda(K(\omega))$.

Proof. Let ζ_1, ζ_2, \dots be an enumeration of S , and let g_1, g_2, \dots be as in Proposition 10 w.r.t. \mathcal{F} . Define $\Phi : \Omega \rightarrow S$ by $\Phi(\omega) = \zeta_{n(\omega)}$, where $n(\omega)$ is the unique n such that $g_n(K(\omega)) = \omega$. We can then define $\Lambda(q) = (g_q^p, s_q^p)_{p \in \mathcal{P}}$ by

$$g_q^p(\Phi(\omega), x) = r^p(\omega, x)$$

and

$$s_q^p(\Phi(\omega))[\Phi(\eta)] = t^p(\omega)[\eta]$$

It is straightforward to check that Φ and Λ so defined satisfy the requirements. \square

Proof. (of Theorem 2.I) For simplicity, take the case that the common knowledge equivalence relationship is aperiodic. Otherwise, partition the space into the common knowledge components of each size, and on each use a modified version of Proposition 20 with S being of a fixed countable or finite size.

²⁴ An equivalence relationship is aperiodic if each equivalence class is infinite. So we will say that a σ -algebra is aperiodic if each atom is infinite.

Let $\psi : \mathfrak{B} \rightarrow (\prod_{p \in \mathcal{P}} \Delta(I^p))$ be a Bayesian equilibrium selection as in Corollary 19. Let Φ, Λ be as in Proposition 20 for some countable set S . For each $\omega \in \Omega$, define

$$\sigma(\omega) = \psi(\Lambda(K(\omega)))(\Phi(\omega))$$

Such σ is then an MBE. \square

Proof. (of Theorem 2.II) Let $C = 2^{\mathbb{N}}$ denote the Cantor space and let \mathcal{E}_t be the tail equivalence relationship; i.e., if $S : C \rightarrow C$ is defined by $(Sx)_n = x_{n+1}$, then by $x \mathcal{E}_t y$ iff $\exists k, m \geq 0, S^k(x) = S^m(y)$. This is a countable Borel equivalence relationship which is non-smooth and hyperfinite, see (Dougherty et al., 1994, Sec. 6). Now, let $X = \{-1, 1\} \times C$, and define $S_X : X \rightarrow X$ by $S_X(x_0, x_1, x_2, \dots) = (-x_0, S(x_1, x_2, \dots))$. Let \mathcal{E}_X be the equivalence relationship on X given by

$$\mathcal{E}_X = \{(x, y) \mid \exists k, m \geq 0, S_X^k(x) = S_X^m(y)\}$$

This relationship is hyperfinite as the product of hyperfinite relationships (see Proposition 5.2 of Dougherty et al. (1994)) and hence by Proposition 11, $(X, \mathcal{E}) \sqsubset (C, \mathcal{E}_t)$; let ψ denote such an embedding.

Let $\Gamma_X = (X, \{1, 2\}, \{L, R\} \times \{L, R\}, t_X^1, t_X^2, r_X^1, r_X^2)$ be the two-player game presented in Hellman (2012b) with state space X as above which does not possess an ε -MBE for small enough ε ; the common knowledge equivalence relationship of that game is indeed \mathcal{E}_X . The type space in that game is deduced from a common prior μ_X , and it also easy to see there that there is $\delta > 0$ such that if ν is a different common prior, but satisfies $\|\nu - \mu_X\| \leq \delta$ (in total-variation distance), then the game with the induced type space also does not possess a ε -MBE for ε small enough. Fix some such δ, ε .

Let $\mathcal{H}^1, \mathcal{H}^2$ be the image on the space $\psi(X)$ under ψ of the knowledge σ -algebras for players 1, 2 in Γ_X , and extend these to σ -algebras on Ω by adding all Borel subsets of $\Omega \setminus \psi(X)$ and observing the generated σ -algebra. Since \mathcal{F} is belief-induced, there are knowledge σ -algebras for Players 3, \dots, n for some $n \geq 3$, $\mathcal{H}^3, \dots, \mathcal{H}^n$, such that \mathcal{F} is the common knowledge σ -algebra of the knowledge σ -algebras $\mathcal{H}^1, \dots, \mathcal{H}^n$; i.e., $\mathcal{F} = \cap_{p=1}^n \mathcal{H}^p$.

By Proposition 9, we deduce that there is a common prior ν satisfying $\|\nu - \mu\| \leq \delta$ such that the induced type space induced by ν , $\tau = (\Omega, \{1, \dots, n\}, t^1, \dots, t^n)$, is positive, where $\mu = \psi_*(\mu_X) = \mu_X \circ \psi^{-1}$. Now, define the payoffs by

$$r^j(\omega, x) = \begin{cases} r_X^j(\psi^{-1}(\omega, x)) & \text{if } j = 1, 2 \text{ and } \omega \in \psi(X) \\ 0 & \text{otherwise} \end{cases}$$

By the properties of Γ_X listed above and for ε chosen above, the game

$$\Gamma = (\Omega, \{L, R\}^n, \{1, 2, 3, \dots, n\}, t^1, t^2, t^3, \dots, t^n, r^1, r^2, r^3, \dots, r^n)$$

which has common prior ν , does not possess an ε -MBE \square

8. PROOF OF THEOREM 3

The proof of Theorem 3.I was already given in Section 6. To prove Theorem 3.II, we first note that it holds on the example Γ_X given in Section 3, as explained there, and in fact using only two players. The common knowledge equivalence relationship \mathcal{E}_X is easily seen to be hyperfinite - it is clearly induced by a \mathbb{Z} -action - and $(X, \mathcal{E}_X), (X, \mathcal{E})$ are both non-smooth, where $\mathcal{E}_X, \mathcal{E}$ denote the respectively common knowledge equivalence relationships. Hence, by Proposition 11, $(X, \mathcal{E}_X) \sqsubset (\Omega, \mathcal{E})$ via an embedding ψ , and let μ be the induced measure - $\mu = \psi_*(\mu) = \mu \circ \psi^{-1}$. Last K be a common knowledge component in Γ such that $\psi^{-1}(K)$ is one of those components in Γ_X on which there is an acceptable bet (f_1, f_2) (with $f_1 = -f_2$). By assumption, this is true for μ -a.e. component K .

We need to show that there is an acceptable bet in K . A modified definition (compare with (3.2)) of an acceptable bet which is helpful is that of a *Dutch book*: In this case, we require $\sum_p f_p < 0$ and $E^p[f_p \mid \cdot] > 0$ at each point for each player. More generally, if $L \subseteq K$, we will say that $(f_p)_{p \in \mathcal{P}}$ is a Dutch book in L if these inequalities hold throughout L . It's easy to show that the existence of an acceptable bet (on a countable space, or any subset of it) is equivalent to the existence of a Dutch book.

In our case, we begin with an acceptable bet for Players 1, 2 on a subset of K : Indeed, it is the image under ψ of the acceptable bet on $\psi^{-1}(K)$. Hence, we also have a Dutch book for these players on that subset - and we need to show that there is a Dutch book on the entire space:

Proposition 21. *Let $\Gamma_K = (K, \mathcal{P}, (t^p)_{p \in \mathcal{P}})$ be a countable positive²⁵ type space such that K does not strictly contain any non-empty common knowledge set, let $\mathcal{Q} \subseteq \mathcal{P}$ and $L \subseteq K$, and suppose there is an Dutch book for the players in \mathcal{Q} on L . Then there is a Dutch book for all the players in \mathcal{P} on all of K .*

Note that the fact that L is not common knowledge for all players is not relevant; in fact, we do not even make use of the fact that it is common knowledge for the players in \mathcal{Q} . Also note that the assumptions imply that $E^p[g_q \mid \cdot] \geq 0$ for all $q \in \mathcal{Q}$ on all K .

Proof. First, we observe that there exists such a Dutch book for all players in \mathcal{P} on all of K : If we define $(g_p)_{p \in \mathcal{P} \setminus \mathcal{Q}}$ to be positive but small enough in L and vanish outside of L , then $(g_p)_{p \in \mathcal{P}}$ is a Dutch book in L and we still have $E^p[g_p \mid \cdot] \geq 0$ for all $p \in \mathcal{P}$ on all K . Fix some $M > \sup_{p \in \mathcal{P}, \omega \in L} |g_p(\omega)|$.

Now, we proceed inductively and keep enlarging L : Let $\omega_0 \in K \setminus L$ and $p_0 \in \mathcal{P}$ be such that $K^{p_0}(\omega_0) \cap L \neq \emptyset$; if there are no such ω_0, p_0 , then we are done by the assumption that there are no subsets which are common knowledge. Fix some

²⁵We could modify the proof to include the non-positive case, but it would be unnecessary and also somewhat more cumbersome.

$\omega_1 \in K^{p_0}(\omega_0) \cap L$. By assumption, $E^{p_0}[g_p \mid \omega_0] = E^{p_0}[g_p \mid \omega_1] > 0$. Let

$$\gamma = \frac{1}{2} \min \left[- \sum_{p \in \mathcal{P}} g_p(\omega_1), t^{p_0}(\omega_0)[\omega_0] \cdot (M - \sup_{p \in \mathcal{P}, \omega \in L} |g_p(\omega)|) \right]$$

By assumption of $(g_p)_{p \in \mathcal{P}}$ being a Dutch bet in L , $\gamma > 0$. Hence, define $(g'_p)_{p \in \mathcal{P}}$ by:

$$g'^p(\omega) = \begin{cases} g^p(\omega) & \text{if } p \neq p_0, \omega \neq \omega_0 \text{ or } p = p_0, p \neq \omega_0, \omega_1 \\ g^p(\omega) + \gamma & \text{if } p = p_0, \omega = \omega_1 \\ g^p(\omega) - \gamma \cdot \frac{t^{p_0}(\omega_0)[\omega_1]}{t^{p_0}(\omega_0)[\omega_0]} & \text{if } p = p_0, \omega = \omega_0 \\ g^p(\omega) + \frac{\gamma}{|\mathcal{P}|} \cdot \frac{t^{p_0}(\omega_0)[\omega_1]}{t^{p_0}(\omega_0)[\omega_0]} & \text{if } p \neq p_0, \omega = \omega_0 \end{cases}$$

It is then easy to check that $(g'^p)_{p \in \mathcal{P}}$ is a Dutch book on $L \cup \{\omega_0\}$ which still satisfies $E^p[g'_p \mid \cdot] \geq 0$ for all $p \in \mathcal{P}$ on all K (note that $E^{p_0}[g_{p_0} \mid \omega_0] = E^{p_0}[g'_{p_0} \mid \omega_0]$, and all others payoffs have decreased nowhere and have increased at ω_0). Furthermore, we still have $M > \sup_{p \in \mathcal{P}, \omega \in L} |g'_p(\omega)|$. Now repeat the procedure with $L \cup \{\omega_0\}$ replacing L ; the resulting Dutch book from this process will also be bounded by M . \square

9. A FURTHER RESULT: AGREEING TO AGREE

In this section we assume that there are only two players. The following definitions are taken from Lehrer and Samet (2011):

Definition 22. Let E be an event in the state space (Ω, \mathcal{B}) with information structure (Π^1, Π^2) and type functions (t^1, t^2) . An agreement on E is an event of the form

$$\{\omega \in \Omega \mid t^1(\omega)(E) = t^2(\omega)(E) = p\}$$

for some $0 < p < 1$. We say that agreeing to agree is possible for E (with μ) if there is a common prior μ for the type functions t^1, t^2 and an agreement A on E such that $\mu(K^\infty(A)) > 0$.

We also define an ignorance operator as in Lehrer and Samet (2011):

Definition 23. The event that Player p is ignorant of event E is

$$I^p(E) = (\Omega \setminus K^i(E)) \cap (\Omega \setminus K^i(\Omega \setminus E))$$

and $I(E) := I^1(E) \cap I^2(E)$.

In addition, for event F , we define the knowledge operator K_F^p , which is the knowledge operator induced by the partition generated Π^p and $\{F, \Omega \setminus F\}$, and we define the associated higher-order knowledge operators, the common knowledge operator K_F^∞ , as well as the operator I_F .

Theorems 1 and 2 from Lehrer and Samet (2011) can be summarized as follows:

Theorem 5. Assume the state space is countable.²⁶ The following conditions are equivalent for an event E :

²⁶ Countable sets are automatically endowed with the discrete σ -algebra.

- (i) Agreeing to agree is possible for E (with some common prior).
- (ii) There exists a non-empty finite event F such that $F \subseteq K^\infty(I(E))$ and $F \subseteq K_F^\infty(I_F(E))$.
- (iii) Agreeing to agree is possible for E with a common prior with finite support.

As remarked in Lehrer and Samet (2011), the implication (ii) \rightarrow (i) holds even if the state space is uncountable; but, by example, the converse direction does not hold. We wish to prove the following, answering an open problem raised in Section 5.1 of Lehrer and Samet (2011).

Theorem 6. *Assume that the state space (Ω, \mathcal{B}) is standard Borel. Assume the player's knowledge structure is such that the common knowledge σ -algebra \mathcal{F} is separable. Let μ be a common prior. The following conditions are equivalent for an event E :*

- (i) Agreeing to agree is possible for E with μ .
- (ii) There exists an event F with $\mu(F) > 0$ such that $F \subseteq K^\infty(I(E))$, $F \subseteq K_F^\infty(I_F(E))$, and such that the intersection of F with any common knowledge component is finite.
- (iii) Agreeing to agree is possible for E with a common prior ν , which is absolutely continuous w.r.t. μ , for which there exists a Borel set G , intersecting each common knowledge component in finitely many points, such that $\nu(G) = 1$.

Proof. Regardless of whether \mathcal{F} is separable or not, by taking $\nu(\cdot) = \mu(\cdot \mid F)$, we see that (ii) implies (iii); also, that (iii) implies (i) is immediate.

To prove that (i) implies (ii), we rely on the countable case. For each common knowledge component C , let $I(\cdot, C)$, $K(\cdot, C)$ and, for each $H \subseteq C$ let $I_H(\cdot, C)$, $K_H(\cdot, C)$ denote the versions of the operators $I(\cdot)$, $K(\cdot)$, $I_H(\cdot)$, $K_H(\cdot)$ restricted to C . Note that if \mathfrak{C} is a collection of common knowledge components and $A, H \subseteq \cup_{C \in \mathfrak{C}} C$ are any sets, then

$$K(A) = \cup_{C \in \mathfrak{C}} K(A \cap C, C)$$

and similarly for K^∞ and I , while

$$K_H(A) = \cup_{C \in \mathfrak{C}} K_{H \cap C}(A \cap C, C)$$

and similarly for K_H^∞ and I_H . (Intuitively, these operators apply independently on each common knowledge component.) Define a correspondence from the standard Borel space Ω/\mathcal{F} to the standard Borel space Z of non-empty finite subsets of Ω (identified with Ω^* modulo the appropriate permutations):

$$\Theta(C) = \{F \in Z \mid F \subseteq K^\infty(I(E \cap C, C)) \text{ and } F \subseteq K_F^\infty(I_F(E \cap C, C), C)\}$$

We note that since the common knowledge σ -algebra is separable, there is by Theorem 1 a mapping ρ from Ω/\mathcal{F} to probability distributions, assigning to each common knowledge component a common prior on it. We contend that there is $\mathfrak{C} \subseteq \Omega/\mathcal{F}$ satisfying $\mu(\mathfrak{C}) > 0$ (where μ also denotes the measure induced on

Ω/\mathcal{F}) such that, for all $C \in \mathfrak{C}$, agreeing to agree is possible for $E \cap C$ with $\rho(C)$. Indeed,

$$0 < \mu(K^\infty(E)) = \int_{\Omega} \rho([\omega])(E \cap C) d\mu(\omega) = \int_{\Omega/\mathcal{F}} \rho(C)(E \cap C) d\mu(C)$$

Hence, by Theorem 5, $\Theta(C) \neq \emptyset$ for all $C \in \mathfrak{C}$. By the Aumann selection theorem (e.g., Himmelberg (1975)), up to the need discard a set of measure zero, there is a Borel mapping $\theta : \mathfrak{C} \rightarrow Z$ such that $\theta(C) \in \Theta(C)$ for all $C \in \mathfrak{C}$. Since this map is clearly injective ($\theta(C) \cap \theta(C') = \emptyset$ in fact if $C \neq C'$) its image is Borel,²⁷ and hence it's easy to see that so is $F = \cup_{C \in \mathfrak{C}} \theta(C)$. This F is the required set, since

$$F = \cup_{C \in \mathfrak{C}} \theta(C) \subseteq \cup_{C \in \mathfrak{C}} K^\infty(I(E \cap C, C)) = K^\infty(I(E))$$

and similarly

$$F = \cup_{C \in \mathfrak{C}} \theta(C) \subseteq \cup_{C \in \mathfrak{C}} K_{F \cap C}^\infty(I_{F \cap C}(E \cap C, C)) = K_F^\infty(I_F(E))$$

□

10. EXTENSIONS AND VARIATIONS

10.1. Countable Partitions. Consider the following natural model. Let the continuum states be represented by the real numbers in the interval $[0,1]$ and suppose that following receipt of a signal each player gives positive support to a sub-interval of $[0,1]$. Further assume that players are limited to some finite accuracy in their measurements and therefore the end-points of the sub-intervals in their posteriors are limited to rational numbers. In that case there can only be a countable number of distinct of partition elements in the posteriors.

Limiting the partitions to countable cardinalities is sufficient to guarantee the existence of Bayesian equilibria, even when the cardinality of the support of every posterior element is the continuum. This follows from Theorem 1 of Milgrom and Weber (1985), because the countable cardinality of the partition elements guarantees that the game has absolutely continuous information, as defined in that paper.

10.2. Types with Countable Support. None of the results of this paper would change if we allow for type spaces with countable support; that is, for each $\omega \in \Omega$ and each Player p , $t^p(\omega)$ is a purely atomic (but not necessarily finitely supported) Borel measure. The proofs all remain largely the same, with only minor alterations. The condition of belief induced also remains unaltered:

Proposition 24. *A countable Borel equivalence relationship \mathcal{E} is belief induced iff there are smooth countable Borel equivalence relationships $\mathcal{E}_1, \dots, \mathcal{E}_n$ which generate \mathcal{E} .*

Proof. Since clearly the countable Borel equivalence relationship generated by a single player's type is smooth, it suffices to show that if \mathcal{E} is a smooth countable Borel equivalence relationship, then there are Borel equivalence relationships

²⁷This follows from Kuratowski's theorem.

$\mathcal{E}_1, \mathcal{E}_2$ with finite equivalence classes which generate \mathcal{E} . Let g_1, g_2, \dots correspond to \mathcal{F} - the σ -algebra induced by \mathcal{E} - as in Proposition 10. For convenience, we write $\{g_m(\omega), g_n(\omega)\}$ even if ω is not in one or both of the domains of g_m, g_n ; in these cases, this set is either empty (if in neither domain) or consists of a single element (if belonging to one domain). Then set,

$$\mathcal{E}_1 = \{(x, y) \in \Omega \times \Omega \mid \exists q \in \Omega/\mathcal{F}, k \in \mathbb{N}, (x, y) \subseteq \{g_{2k-1}(q), g_{2k}(q)\}\}$$

$$\mathcal{E}_2 = \{(x, y) \in \Omega \times \Omega \mid \exists q \in \Omega/\mathcal{F}, k \in \mathbb{N}, (x, y) \subseteq \{g_{2k}(q), g_{2k+1}(q)\}\}$$

It is easy to see that the equivalence classes of $\mathcal{E}_1, \mathcal{E}_2$ are all of size at most 2, and that \mathcal{E} is generated by $\mathcal{E}_1, \mathcal{E}_2$. \square

11. APPENDIX: ON BELIEF INDUCED RELATIONSHIPS

As we have mentioned, not all countable Borel equivalence relationships are belief induced. This can be shown using the concept of the *cost* of a countable Borel equivalence relationship \mathcal{E} with an invariant²⁸ measure μ . We briefly recall this concept; for a more comprehensive treatment, see Kechris and Miller (2004).

A Borel graph G on a standard Borel space Ω is a Borel relation on Ω (i.e., a Borel subset of $\Omega \times \Omega$) that is irreflexive and symmetric. A Borel graph G induces a Borel equivalence relationship \mathcal{E} on Ω : \mathcal{E} is the reflexive and transitive closure of G . We say that G is a *graphing* of \mathcal{E} . Given such a graph, for each $v \in \Omega$, let $d_G(v) \in \{0, 1, 2, \dots, \infty\}$ denote the cardinality of the set $\{w \in \Omega \mid (v, w) \in G\}$. Clearly, if $d_G(v)$ is countable for all $v \in \Omega$, then so is the induced equivalence relationship \mathcal{E} ; conversely, if \mathcal{E} is a countable Borel equivalence relationship, then it is induced by some Borel graph with vertices of countable degree (this follows easily, e.g., from Proposition 4).

The cost of a countable Borel equivalence relationship \mathcal{E} (with respect to an invariant measure μ) is defined as:

$$C_\mu(\mathcal{E}) := \inf \left\{ \frac{1}{2} \int_\Omega d_G(\omega) d\mu(\omega) \mid G \text{ spans } \mathcal{E} \right\}$$

A result of Levitt, e.g. (Kechris and Miller, 2004, Ch. 20), is that if T is a Borel transversal for a countable Borel equivalence relationship \mathcal{E} with an \mathcal{E} -invariant measure μ , then $C_\mu(\mathcal{E}) = \mu(\Omega \setminus T)$; in particular, if μ is finite, then so is $C_\mu(\mathcal{E})$.

Suppose that \mathcal{E} is a countable Borel equivalence relationship, and is the equivalence relationship generated by $\mathcal{E}_1, \dots, \mathcal{E}_n$ (that is, the coarsest equivalence relationship that each \mathcal{E}_k refines), and suppose that μ is \mathcal{E} -invariant. It is then clearly also \mathcal{E}_k invariant for each $k = 1, \dots, n$, and it's easy to see that²⁹

$$C_\mu(\mathcal{E}) \leq \sum_{k=1}^n C_\mu(\mathcal{E}_k)$$

²⁸ A (not-necessarily finite) measure is \mathcal{E} -invariant if for every Borel bijection $f : \Omega \rightarrow \Omega$ satisfying $f(\omega) \sim_{\mathcal{E}} \omega$ for all $\omega \in \Omega$, it holds that for all Borel $A \subseteq \Omega$, $\mu(f^{-1}(A)) = \mu(A)$.

²⁹ Note that μ is \mathcal{E}_k -invariant for each $k = 1, \dots, n$; hence, for any graphings G_1, \dots, G_n of $\mathcal{E}_1, \dots, \mathcal{E}_n$, respectively, $G = \cup_{k=1}^n G_k$ spans \mathcal{E} .

Combining this observation with the result of Levitt – and the fact that finite³⁰ Borel equivalence relationships are clearly always smooth³¹ – we see that if $\mathcal{E}_1, \dots, \mathcal{E}_n$ are finite, $C_\mu(\mathcal{E})$ is finite.

Hence, to show a non-belief induced countable Borel equivalence relationship, it suffices to find one with infinite cost w.r.t. some invariant measure \mathcal{E} on it. A result of Gaboriau, e.g. (Kechris and Miller, 2004, Cor 27.10), states that if \mathcal{E} is a countable Borel equivalence relationship with finite invariant measure μ , and \mathcal{T} is a Borel tree³² that is a graphing of \mathcal{E} , then $C_\mu(\mathcal{E}) = \frac{1}{2} \int_\Omega d_{\mathcal{T}}(\omega) d\mu(\omega)$.

Now, let F_∞ denote the free (non-abelian) group with countably many generations. This group acts on 2^{F_∞} via $(f(x))(g) = x(f \cdot g)$ for $x \in 2^{F_\infty}$, $f, g \in F_\infty$, and induces a countable Borel equivalence relationship by $x \sim y$ iff $\exists g \in F_\infty$ with $g \cdot x = y$. From this, one deduces easily that if $\mu = \prod_{f \in F_\infty} (\frac{1}{2}, \frac{1}{2})$ (which is clearly \mathcal{E} -invariant) it holds by Gaboriau's result that $C_\mu(\mathcal{E}) = \infty$.

We also need to complete the following:

Proof. (Proof of Proposition 9) It's enough to find some common prior ν with positive types which is not necessarily close to μ , because then we can replace ν with $\epsilon\nu + (1 - \epsilon)\mu$ and similarly mix the types we had constructed previously with the types resulting from the proper RCD's of μ w.r.t. the knowledge structures.

Let G^p be a countable group which generates the knowledge structure \mathcal{F}^p , as in Proposition 4. Let G be the countable group generated by G^1, \dots, G^p . Let $(\alpha_g)_{g \in G}$ be some collection of positive real numbers which sum to unity, and let

$$\nu = \sum_{g \in G} \alpha_g g_*(\mu)$$

where $g_*(\mu) = \mu \circ g^{-1}$, let τ^p be a proper RCD of ν w.r.t. to \mathcal{F}^p for each $p \in \mathcal{P}$. It follows from Proposition 8 that the set N of ω for each $t^p(\omega)[\omega] = 0$ for some $p \in \mathcal{P}$ is of ν -measure 0. We contend that $\nu(K(N)) = 0$: Once we have this, one can redefine the types in an arbitrary measurable and positive way (while still assuring that they generate $\mathcal{F}^1, \dots, \mathcal{F}^n$ on $K(N)$, of course; for example, for $\omega \in K(N)$ and $p \in \mathcal{P}$, let $t^p(\omega)$ be uniform on $K(\omega)$); ν remains a common prior for the altered types.

In fact, we contend that for each $g \in G$, g sends ν -null sets to ν -null sets; this will suffice, since the group G generates the common knowledge equivalence relationship (in the sense of Proposition 4).

³⁰ An equivalence relationship is called finite if its equivalence classes are all finite.

³¹ One can see that finite Borel equivalence relationships are always smooth, for example, by taking a Borel ordering $<$ on Ω , and choose the $<$ -minimal element in each equivalence class to get a transversal

³² A Borel tree is a Borel graph with no cycles.

So fix $g_0 \in G$, and let N be any null set. Observe that

$$\nu(N) = \sum_{g \in G} \alpha_g g_*(\mu)(N)$$

and hence $g_*(\mu)(N) = 0$ for all $g \in G$; hence,

$$\nu(g(N)) = \sum_{g \in G} \alpha_g g_*(\mu)(g_0(N)) = \sum_{g \in G} \alpha_g (g \cdot g_0^{-1})_*(\mu)(N) = \sum_{g \in G} \alpha_{g \cdot g_0} (g)_*(\mu)(N) = 0$$

□

REFERENCES

- BECKER, H. AND KECHRIS, A.S. (1994), *The Descriptive Set Theory of Polish Group Actions*, Cambridge University Press.
- BLACKWELL, D. AND RYLL-NARDZEWSKI, C. (1963), Non-Existence of Everywhere Proper Conditional Distributions, *Ann. Math. Statist.*, 34, 223–225.
- BLACKWELL, D. AND DUBINS, L. (1975), On Existence and Non-Existence of Proper, Regular, Conditional Distributions, *Ann. Prob.*, 3, 741–752.
- CHATTERJEE, K., AND W. SAMUELSON (1973), Bargaining Under Incomplete Information, *Operations Research*, 5, 835–851.
- COTTER, K.D. (1991), Similarity of Games with Incomplete Information, *J. of Math. Econ.*, 20, 501–520.
- DOUGHERTY, R., JACKSON, S., KECHRIS, A.S. (1994), The Structure of Hyperfinite Borel Equivalence Relations, *Trans. Amer. Math. Soc.*, 341, 193–225.
- FEINBERG, Y. (2000), Characterizing Common Priors in the Form of Posteriors, *Journal of Economic Theory* 91, 127–179.
- FELDMAN, J. AND MOORE, C.C. (1977), Ergodic Equivalence Relations, Cohomology and von Neumann Algebras, I, *Trans. Amer. Math. Soc.*, 234, 289–324.
- HARSÁNYI, J. (1967), Games of Incomplete Information Played by Bayesian Players, Part I: The Basic Model, *Management Sci.*, 14, 159–182.
- HARRINGTON, L., KECHRIS A.S., LOUVEAU, A. (1990), A Glimm-Effros Dichotomy for Borel Equivalence Relations, *J. Amer. Math. Soc.*, 3, 903–228.
- HEIFETZ, A. (2006), The Positive Foundation of the Common Prior Assumption, *Games and Economic Behavior*, 56, 105–120.
- HELLMAN, Z. (2012A), Countable Spaces and Common Priors Equilibria. DP #604, Center for the Study of Rationality, Hebrew University, Jerusalem.
- HELLMAN, Z. (2012B), A Game with No Bayesian Approximate Equilibria. DP #615, Center for the Study of Rationality, Hebrew University, Jerusalem.
- HELLMAN, Z. AND SAMET, D. (2011), How Common are Common Priors?, *Games and Econ. Behav.*, 74, 517–525.
- HERRLICH, H. (2006), *Axiom of Choice*. Lecture Notes in Mathematics, Springer.
- HIMMELBERG, C. J. (1975), Measurable relations, *Fund. Math.*, 87, 53–72.
- KECHRIS, A. S. (1995), *Classical Descriptive Set Theory*, Graduate Texts in Mathematics 156. Springer.
- KECHRIS, A.S. AND MILLER, B.D. (2004), *Topics in Orbit Equivalence*, Lecture Notes in Mathematics 1852. Springer-Verlag, Berlin.

- KURATOWSKI, K. AND RYLL-NARDZEWSKI, C. (1965), A General Theorem on Selectors, *Bull. Pol. Acad. Sci. Math (Ser. Math.)*, 13, 379-403.
- LEHRER, E. AND SAMET, D. (2011), Agreeing to Agree, *Theoretical Economics*, 6, 269–287.
- MILGROM, P. AND R. WEBER (1985), Distributional Strategies for Games with Incomplete Information, *Mathematics of Operations Research*, 11, 619–631.
- MYERSON, R. (1997), *Game Theory: Analysis of Conflict*, Cambridge, Harvard University Press.
- RADNER, R. (1980), Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives, *Journal of Economic Theory*, 22, 136–154.
- RUDIN, W. (1986), *Real and Complex Analysis*, 3rd Ed., McGraw-Hill Book Company.
- SAMET, D. (1998), Iterated Expectations and Common Priors, *Games and Economic Behavior*, 24, 131–141.
- SCHMEIDLER, D. (1973), Equilibrium Points of Nonatomic Games, *Journal of Statistical Physics*, 7 (4), 295–300.
- SIMON, R. (2000), The Common Prior Assumption in Belief Spaces: An Example, DP #228, Center for the Study of Rationality, Hebrew University, Jerusalem.
- SIMON, R. (2003), Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria. *Israel J. Math.*, 138, 1, 73–92.