

EVENTS CONCERNING KNOWLEDGE

EDWARD J. GREEN

ABSTRACT. Common knowledge of a Borel event is shown to be a co-analytic event, and is therefore universally measurable. An extension of Aumann’s “agreement theorem” regarding common knowledge of posterior probabilities is proved in the framework of a measure space defined on a complete, separable, σ -compact metric space. *JEL classification:* D82, D83, D84

Keywords: information partition, type space, common knowledge, agreement theorem

1. INTRODUCTION

The two main results to be reported here are (a) a proposition that common knowledge of a Borel event is a co-analytic event, and is therefore universally measurable, and (b) an extension of Aumann’s [1976] “agreement theorem” regarding common knowledge of posterior probabilities to the framework of a measure space defined on a complete, separable, σ -compact metric space.

Analysis involving common knowledge is prominent in game theory, and consequently in several fields of applied economics that draw on game theory. Following the research of Brandenburger and Dekel [1987] and Monderer and Samet [1989], many game theorists have shifted from common knowledge to common certainty as the object of analysis.¹ As Brandenburger and Dekel point out, a main motivation for making that shift is to surmount the limitation of Aumann’s [1976] analysis to the countable-information-partition framework. In particular, the universal type space of a game does not fit that framework. Monderer and Samet [1989, p. 179] point out that common knowledge and common certainty of an event differ by an event of prior probability zero, but do not exactly coincide in general. Ben-Porath [1997] shows that the difference between them is crucial to considering justification of backward induction in extensive-form games of perfect information. Thus, the universal-measurability result to be presented here may have significant game-theoretic applications.

Date: 2012.04.17.

An insightful suggestion by Marcin Peski led to the positive results reported here. The author acknowledges with gratitude the hospitality of CEMFI (<http://www.cemfi.es>) where part of the research has been conducted, and the Human Capital Foundation, (<http://www.hcfoundation.ru/en/>) and particularly Andrey P. Vavilov, for research support through the Center for the Study of Auctions, Procurements, and Competition Policy (CAPCP, <http://capcp.psu.edu/>) at the Pennsylvania State University.

¹Certainty is defined analogously to knowledge, but in terms of probability-1 posterior belief rather than in terms of the known event comprehending a block of an information partition.

The connection between the agreement theorem and financial economics comes via the “no trade theorem” of Milgrom and Stokey [1982] and Tirole [1982].² Milgrom and Stokey’s version of that result can be paraphrased as follows.

No-Trade Theorem. *Suppose that*

- *there are finitely many states of the world;*
- *agents’ prior beliefs about the state of nature are identical;*
- *agents maximize subjective expected utility, and are weakly risk averse;*
- *agents agree on the probabilities, conditional on utility-relevant events, of the signals that they receive;*
- *with respect to prior beliefs, the endowment is Pareto efficient; and*
- *it is common knowledge that a particular trade is feasible, and is unanimously weakly preferred to the endowment with respect to posterior expected utilities;*

then, with respect to posterior expected utility, all traders are indifferent between that particular trade and the endowment.

In Tirole’s exposition, the first premise of the theorem is generalized slightly [1982, p. 1166], to specify that every signal has positive probability of being received (so that the set of signals is countable). This assumption corresponds to Aumann’s [1976, p. 1236] assumption. Regarding the role of this assumption, Tirole writes that “It is clear that the result holds for much more general probability spaces.” The study lays the groundwork for attempting to prove a no-trade theorem in a more general setting, a project with a less certain outcome than Tirole’s remark might suggest.

Standard Borel spaces (Such as the unit interval with its Borel sets) are an indispensable ingredient of the formal models of asset pricing and of other models studied in financial economics. These uncountable state spaces are building blocks of continuous asset-return distributions (such as log-normal), diffusion processes, and so forth. Thus, strictly speaking, the no-trade theorem is irrelevant to financial economics. However researchers in that field have universally adopted Tirole’s view that the formulation of the no-trade theorem in terms of a countable state space is a mere mathematical convenience. In fact, many of the most insightful and careful researchers have proceeded in that way. For example, Brennan and Cao [1996] invoke the no-trade theorem to reduce a dynamic trading problem to a static one. They formulate a model in which each agent receives a normally distributed signal [1996, pp. 166-167] (so the state space cannot be countable). They state that “The no-trade result of Milgrom and Stokey (1982) applies, and no further trade occurs as new public information... becomes available.” [1996, pp. 166-167] However, the upshot of the present research is that there is actually a gap at this step of Brennan and Cao’s proof.

The typical context for citation of the no-trade theorem is the literature regarding speculative trade. Virtually every author on that topic begins by noting that speculative trade appears sometimes to occur but (purportedly) cannot be modeled within the Bayesian rational expectations framework because the no-trade theorem

²These two results are proved in superficially different environments, but clearly are conceptually equivalent. Milgrom and Stokey derive their result explicitly as a corollary of the agreement theorem, while Tirole proves his result “from scratch.” It is well understood that, in either case, Aumann’s result is at the core of the no-trade result.

would rule it out, and then proceeds to “spoil” the Bayesian RE model in one respect or another, in order to produce the phenomenon that is to be analyzed. For example, Easley et al. [1998] formulate a model that they analyze by taking a derivative with respect to a price [1998, p. 441], by which they implicitly assume that the state space must be uncountable. They write that “[U]ninformed... trade [arises] for... reasons that are exogenous to the model. This ‘noise trader’ assumption is standard in microstructure models and it reflects the difficulty noted by Milgrom and Stokey (1982)...” [1998, p. 436] But the present research makes it clear that, in fact, it is an open question whether or not the Bayesian RE environment needs to be modified in any way (such as positing noise traders, non-EU preferences, and so forth) in order to accommodate speculative trade.

2. A MEASURABILITY PROBLEM WITH KNOWLEDGE

2.1. Formalizing knowledge in terms of a partition. $(\Omega, \mathcal{B}, \mu)$ is a probability space.³ Throughout this article, it is assumed that⁴

- (1) Ω is uncountable, \mathcal{B} is the σ -algebra generated by the open sets of some Polish topology Ω , and $\mu: \mathcal{B} \rightarrow [0, 1]$ is a countably additive probability measure.

$\mathfrak{P} \subseteq \mathcal{B}$ is the agent’s information partition of Ω . If $\{\psi, \omega\} \subseteq \pi \in \mathfrak{P}$, then the agent’s information cannot distinguish between ψ and ω having occurred.

For $A \subseteq \Omega$, the *saturation* of A is defined by

$$(2) \quad [A]_{\mathfrak{P}} = \bigcup \{ \pi \mid \pi \in \mathfrak{P} \text{ and } \pi \cap A \neq \emptyset \}.$$

$[A]_{\mathfrak{P}}$ is the event that what the agent observes is consistent with A having occurred.

The agent *knows* that B has occurred, if his observation is inconsistent with B not having occurred. The event $\kappa(B)$ that the agent knows that B has occurred is formalized as follows.⁵

$$(3) \quad \kappa(B) = \Omega \setminus [\Omega \setminus B]_{\mathfrak{P}}.$$

2.2. A decision-theoretic example. Jurisprudence is a practical context in which knowledge is distinguished from subjective belief. In United States law, for example, the Supreme Court has ruled the distinction to be crucial to determining the admissibility of expert testimony. “The subject of an expert’s testimony must be ‘scientific... knowledge.’... The word ‘knowledge’ connotes more than subjective belief or unsupported speculation.”⁶

Consider the following example of a decision problem in which, due to a jurisprudential requirement that testimony must be based on knowledge, an agent’s payoff from an action would be contingent on whether or not another (non-strategic) agent has knowledge of an event. Suppose that an avaricious entrepreneur, Mr. Durham, faces a decision, whether or not to manufacture a product that *might* harm consumers. Durham only wants to maximize his profit, and is indifferent to consumers’

³Aumann [1976] specifies that Ω is countable and that $\mathcal{B} = 2^{\Omega}$ (or, tantamount to that assumption, that \mathcal{B} is generated by a countable partition of Ω). The present formalization of common knowledge coincides in essential respects with the formalization provided in Aumann [1999a,b].

⁴A topology induced by a complete, separable metric is called *Polish*.

⁵This definition is equivalent to the definition introduced by Aumann [1976], which parallels the definition of the modal operator *necessity* formulated by Kripke [1959].

⁶*Daubert v. Merrell Dow Pharmaceuticals*, 509 U.S. 579 (1993).

welfare. If the product is brought to market, then it is inevitable that a consumer will sue for compensation of alleged harm. A reknowned scientist, Dr. Schliemann, has already stated publicly that he is certain that the product would be harmful, that he is confident of his ability to prove the harm scientifically, and that he would be eager to testify as an expert witness. Durham’s decision whether or not to go forward turns on the question: can Schliemann really acquire scientific knowledge that the product is harmful?

Let $\Omega = [0, 1]$ be the space of states of the world. Suppose that the the product is harmful in event $H = [0, 1/2)$ and safe in event $S = [1/2, 1]$.

Consider an event $X \subseteq H$ and its translation $Y = \{\omega + 1/2 \mid \omega \in X\} \subseteq S$. Let $Z = X \cup Y$ be the event that Schliemann will be unable to produce scientific proof (that is, to attain knowledge) to discriminate between some pair of states of the world $\omega \in H$ and $\omega + 1/2 \in S$. In that case, Durham would win the lawsuit. Schliemann must give honest testimony, so Durham also wins if Schliemann knows that the product is safe. In either of those events, Schliemann does not know that the product is dangerous.

Schliemann’s information partition is:

$$(4) \quad \pi \in \mathfrak{P} \iff \begin{cases} \pi = \{\omega\} \text{ and } \omega \notin Z \\ \pi = \{\omega, \omega + 1/2\} \text{ and } \omega \in X \end{cases}$$

The event that Schliemann does not actually know that Durham’s product is harmful is $[S]_{\mathfrak{P}} = S \cup X$.

If Durham makes and sells the product, then

- He wins the lawsuit, and subsequently makes huge profit, in event $[S]_{\mathfrak{P}}$.
- He loses the lawsuit, must pay an enormous settlement to the plaintiff, and is subsequently enjoined from selling the product, in event $\kappa(H) = \Omega \setminus [\Omega \setminus H]_{\mathfrak{P}} = \Omega \setminus [S]_{\mathfrak{P}}$.

Let s and f be Durham’s utility outcomes of his success or failure in the lawsuit. Normalize the utility of not selling the product to be 0. Suppose that $f < 0 < s$. Durham “should” sell if his expected utility from it, U , is positive. Let μ be a nonatomic measure on Ω that represents Durham’s beliefs. Then $U = \mu([S]_{\mathfrak{P}}) \cdot s + (1 - \mu([S]_{\mathfrak{P}})) \cdot f$. He should make and sell the product, if $\mu([S]_{\mathfrak{P}})$ is close to 1.

Since $[S]_{\mathfrak{P}} = S \cup X$, $\mu([S]_{\mathfrak{P}}) = \mu(S) + \mu(X)$. But, what is $\mu(X)$? Nothing has been assumed about X . In fact, X might be an event that is not measurable with respect to μ .⁷ In that case, U is undefined, and Durham’s decision problem falls outside the scope of Bayesian decision theory.

The implication of this example for Aumann’s formalism is disturbing. The partition \mathfrak{P} , consisting entirely of singletons and pairs, seems to be as tame as can be. Likewise for the event S , a closed interval, from which the non-measurable epistemic event in the example is derived. There is no sentence in the language of Aumann’s theory, with which to state an axiom that would imply that X must be measurable. Rather, the implicit hope is that a sufficiently ‘nice’ partition and

⁷An event B is *measurable* with respect to μ if there are events A and C in \mathcal{B} such that $A \subseteq B \subseteq C$ and $\mu(A) = \mu(C)$. An event that is not measurable with respect to any nonatomic measure exists, by Oxtoby [1980, Theorems 5.3, 5.4], if the axiom of choice is satisfied. The axiom of choice will be assumed throughout the present article.

event will always generate a measurable knowledge event. The example dashes that hope.

3. FORMALIZING COMMON KNOWLEDGE IN TERMS OF A PARTITION

Suppose that agents 1 and 2 have information partitions \mathfrak{P}_1 and \mathfrak{P}_2 . Let κ_i denote i 's knowledge operator, defined by (3) with respect to \mathfrak{P}_i . Event B has been defined by Lewis [1969] to be *common knowledge* between the agents in the event that B occurs, each agent knows that it occurs, each knows that the other knows, and so forth.

This concept generalizes to common knowledge among a set of agents. If I is a finite set of agents, then define \mathfrak{P}_I , the *common-knowledge partition*, to be the finest partition that each \mathfrak{P}_i (for $i \in I$) refines. Aumann [1976, p. 1237] sketches an argument (recapitulated below in the proof of proposition 5) that the definition of common knowledge in terms of Lewis' iterated-knowledge concept is equivalent to the following definition of common knowledge.

$$(5) \quad \kappa_I(B) = \Omega \setminus [\Omega \setminus B]_{\mathfrak{P}_I}.$$

This definition of a common-knowledge event is parallel to the definition (3) of the event that an individual agent knows the same event, with \mathfrak{P}_I being substituted for the information partition \mathfrak{P} .

Defining common knowledge in this way risks, by the same argument as in section 2.2, the possibility that the common-knowledge partition will have a non-measurable block. In that case, it is foreseeable that the event of agents having common knowledge of some Borel event will be non measurable.

Consider common knowledge between 2 agents. Let \mathfrak{P}_1 be the partition defined in (4) and let $\mathfrak{P}_2 = \{\{\omega\} \mid \omega < 1\} \cup \{[1, 2]\}$. Then $X \cup [1/2, 1]$ is the block of $\mathfrak{P}_{\{1,2\}}$ that includes the block $[1/2, 1]$ in \mathfrak{P}_2 . If X is not measurable, then, a common-knowledge partition corresponding to information partitions built from the simplest closed sets—singletons, pairs, and a closed interval—includes a non-measurable block.

4. MEASURABILITY, EQUIVALENCE RELATIONS, AND AGENTS' TYPES

4.1. Measurable, universally measurable, and analytic sets. A few concepts relating to measurability of an event are needed, in order to amend Aumann's formalism and to study to what degree measurability difficulties are thus avoided.

Event $B \subseteq \Omega$ is *measurable* with respect to μ , a countably additive probability measure, if there are events A and C in \mathcal{B} such that $A \subseteq B \subseteq C$ and $\mu(A) = \mu(C)$. In that case, μ can be extended to a countably additive measure, μ^* , on a σ -algebra containing $\mathcal{B} \cup \{B\}$, and this *measure completion* satisfies

$$(6) \quad \mu^*(B) = \mu^*(C) \text{ and } \forall D \in \mathcal{B} \quad \mu^*(D) = \mu(D).$$

As has been mentioned in footnote 7, there are events that are not measurable with respect to any nonatomic probability measure (in fact, with respect to any measure for which no countable set has probability 1).

On the other hand, there are events that are *universally measurable*, that is, measurable with respect to every probability measure. These include some events that are not Borel (that is, not in \mathcal{B}). It will be shown below that, when agents' information is represented appropriately, the event that an agent knows a Borel event

is a universally measurable event, even though his possession of that knowledge may not be a Borel event.

A set S , together with a σ -algebra \mathcal{S} of (some of) its subsets, is a *standard Borel space* if there is some Polish topology, such that \mathcal{S} is the smallest σ -algebra that contains all of its open sets.

Now the classes $\Sigma_1^1(S)$ of *analytic* subsets and $\Pi_1^1(S)$ of *co-analytic* subsets of a standard Borel space, (S, \mathcal{S}) , are defined. Define $\Sigma_1^1(S)$ to be the class of sets $X \subseteq S$ such that, for some standard Borel space (T, \mathcal{T}) and some set $Y \in \mathcal{S} \times \mathcal{T}$, $X = \{\sigma \mid \exists \omega \in T (\sigma, \omega) \in Y\}$.⁸ Define $\Pi_1^1(S)$ to be the class of sets $X \subseteq S$ such that, for some standard Borel space (T, \mathcal{T}) and some set $Y \in \mathcal{S} \times \mathcal{T}$, $X = \{\sigma \mid \forall \omega \in T (\sigma, \omega) \in Y\}$. Define $\Sigma_2^1(S)$ to be the class of sets $X \subseteq S$ such that, for some standard Borel space (T, \mathcal{T}) and some set $Y \in \Pi_1^1(S \times T)$, $X = \{\sigma \mid \exists \omega \in T (\sigma, \omega) \in Y\}$.

Also define $\Delta_1^1(S) = \Sigma_1^1(S) \cap \Pi_1^1(S)$, and define $\Pi_2^1(S)$ to be the class of sets $X \subseteq S$ such that, for some standard Borel space (T, \mathcal{T}) and some set $Y \in \Sigma_1^1(S \times T)$, $X = \{\sigma \mid \forall \omega \in T (\sigma, \omega) \in Y\}$.

Lemma 1. *The following assertions hold for all standard Borel spaces (S, \mathcal{S}) and (T, \mathcal{T}) . Events in $\Sigma_1^1(S)$ and in $\Pi_1^1(S)$ are universally measurable. Each of $\Sigma_1^1(S)$ and $\Pi_1^1(S)$ is closed under countable unions and intersections. If $X \in \Sigma_1^1(S)$ and $Y \in \Sigma_1^1(T)$, then $X \times Y \in \Sigma_1^1(S \times T)$. If $f: S \rightarrow T$ is measurable, then $Y \in \Sigma_1^1(T)$ if and only if $f^{-1}(Y) \in \Sigma_1^1(S)$. If $Y \in \Sigma_1^1(S \times T)$ and $X = \{\sigma \mid \exists \tau \in T (\sigma, \tau) \in Y\}$, then $X \in \Sigma_1^1(S)$. $\Delta_1^1(S) = \mathcal{S}$, but $\mathcal{S} \subsetneq \Sigma_1^1(S)$ and $\mathcal{S} \subsetneq \Pi_1^1(S)$.⁹*

4.2. Equivalence relations generalize Harsanyi's agent-type framework. It is an elementary fact that the identity

$$(7) \quad \exists \pi \in \mathfrak{P} [\{\psi, \omega\} \subseteq \pi] \iff (\psi, \omega) \in E$$

associates an equivalence relation $E \subseteq \Omega \times \Omega$ to each partition of Ω , and vice versa. If an equivalence relation E is specified, then it *determines* partition \mathfrak{P} according to this relation. Both E and \mathfrak{P} are called *closed* (resp. *Borel*, *analytic*) if E is a closed (resp. Borel, analytic) subset of $\Omega \times \Omega$.

In much of the modeling of information in game theory, and in almost all of its modeling in applied economics, an agent's knowledge is represented implicitly in terms of the agent's state-contingent *type*. The space of types is posited to be a standard Borel space (T, \mathcal{T}) . A T -valued random variable $\tau: \Omega \rightarrow T$ (that is, a measurable function from (Ω, \mathcal{B}) to (T, \mathcal{T})) determines \mathfrak{P} according to $\mathfrak{P} = \{\tau^{-1}(t) \mid t \in T\} \setminus \{\emptyset\}$.¹⁰

⁸This is one of several, equivalent, conditions by which $\Sigma_1^1(S)$ can be defined. Cf. Bertsekas and Shreve [1978, Proposition 7.41]

⁹These assertions are proved in Bertsekas and Shreve [1978, corollary 7.42.1, corollary 7.35.2, corollary 7.36.1, proposition 7.39], Moschovakis [2009, theorem 1E3, theorem 2E.2]. Moschovakis' assumption that S is perfect in its metric topology (embedded in his specification of the product space \mathcal{X} in his theorem) can easily be removed, both from this result and from other results to be cited below. Cf. Moschovakis [2009, p. 60, fn. 4].

¹⁰If (Ω, \mathcal{B}) and (T, \mathcal{T}) are standard Borel spaces, then $f: \Omega \rightarrow T$ is measurable if, for every $A \in \mathcal{T}$, $f^{-1}(A) \in \mathcal{B}$. In the case that $f: \Omega \times \Omega \rightarrow T$ that will occur subsequently at some points, note that $\Delta_1^1(\Omega) \times \Delta_1^1(\Omega) = \Delta_1^1(\Omega \times \Omega)$ (where the left side denotes the smallest σ -algebra that contains all product "rectangles" $A \times B \subseteq \Omega \times \Omega$ and the right side is the σ -algebra generated by the open sets in the product topology). Cf. Aliprantis and Border [2006, theorem 4.44].

If τ is a such a measurable assignment of a type, then *the agent's type is the same in state ψ as in state ω* is an equivalence relation. Specifically, if $f: \Omega \times \Omega \rightarrow \mathcal{T} \times \mathcal{T}$ is defined by $f(\psi, \omega) = (\tau(\psi), \tau(\omega))$, then $E = f^{-1}(D)$ is the equivalence relation that determines the partition of states according to the agent's type. By lemma 1, E is analytic. In fact, E so defined belongs to a particularly simple sub-class of the Borel equivalence relations.¹¹ However, placing a restriction of being Borel, or even closed, on partitions the information partitions mentioned in the hypotheses of results to be proved below would not enable the conclusions of those results to be strengthened in any significant way.¹² For that reason, the remainder of this article will be predominantly concerned with analytic partitions, equivalence relations, and events.

A partition consisting of blocks in some class of events (closed events, for instance) does not necessarily belong to the corresponding class of partitions, according to this definition. For example, the partition defined by (4) consists of closed blocks, but the equivalence class associated to it by (7) does not have a closed graph, or even a graph in $\Sigma_1^1(\Omega \times \Omega)$, if $X \notin \Sigma_1^1(\Omega)$. The way to see this fact, is to prove its contrapositive: if E is defined by (4) and (7) and $E \in \Sigma_1^1(\Omega \times \Omega)$, then $X \in \Sigma_1^1(\Omega)$, and hence is universally measurable. The proof proceeds as follows.

Define the diagonal (or identity) relation $D \subseteq \Omega \times \Omega$ by

$$(8) \quad (\psi, \omega) \in D \iff \psi = \omega.$$

D is a closed, and therefore Borel, subset of $\Omega \times \Omega$; so $(\Omega \times \Omega) \setminus D \in \Sigma_1^1(\Omega \times \Omega)$ by lemma 1. Then, also by lemma 1, $E \setminus D \in \Sigma_1^1(\Omega \times \Omega)$. Finally, X can be defined from E by $\omega \in X \iff \exists \psi (\psi, \omega) \in E \setminus D$. Finally, lemma 1 shows both that $X \in \Sigma_1^1(\Omega)$ and that X is universally measurable.

Note that, if $\omega \in \pi \in \mathfrak{P}$ and E determines \mathfrak{P} , then $\pi = \{\psi \mid (\psi, \omega) \in E\}$. That is, π is defined by

$$(9) \quad \psi \in \pi \iff (\psi, \omega) \in E.$$

Defining $f: \Omega \rightarrow \Omega \times \Omega$ by $f(\psi) = (\psi, \omega)$, $\pi = f^{-1}(\Omega \times \{\omega\})$, which proves the following lemma.

Lemma 2. *If π is defined from E by (9), and if $E \in \Sigma_1^1(\Omega \times \Omega)$, then π is in $\Sigma_1^1(\Omega)$, and hence is universally measurable.*

To summarize, it has been shown in this section that the pathologies of non-measurability that were displayed in sections 2.2 and 3 are avoided by requiring agents' information partitions to be determined by analytic equivalence relations. This observation generalizes to the following conjectures.

- (1) If each \mathfrak{P}_i is determined by an analytic equivalence relation, then \mathfrak{P}_I is also determined by an analytic equivalence relation and every $\pi \in \mathfrak{P}_I$ is analytic.
- (2) If \mathfrak{P} is determined by an analytic equivalence relation and B is Borel, then $\kappa(B)$ is co-analytic.
- (3) If each \mathfrak{P}_i is determined by an analytic equivalence relation and B is Borel, then $\kappa_I(B)$ is co-analytic.

¹¹Cf. Harrington et al. [1990].

¹²Cf. Green [2012].

These conjectures will be proved below. In fact, the hypothesis that B is Borel can be weakened to a hypothesis that B is co-analytic.

4.3. Transitive closure. For a relation $R \subseteq \Omega \times \Omega$, set¹³

$$(10) \quad R^{(1)} = R \text{ and, for each } n \in \mathbb{N}_+, \\ R^{(n+1)} = \{(\chi, \omega) \mid \exists \psi \in \Omega [(\chi, \psi) \in R \text{ and } (\psi, \omega) \in R^{(n)}]\}.$$

Then

$$(11) \quad R^+ = \bigcup_{n \in \mathbb{N}_+} R^{(n)}$$

is the *transitive closure* of R .

Let \mathcal{E}_R denote the set of equivalence relations $E \subseteq \Omega \times \Omega$ such that $R \subseteq E$.

Lemma 3. *If R is reflexive and symmetric, then R^+ is an equivalence relation. Specifically, $R^+ = \bigcap \{E \mid E \in \mathcal{E}_R\}$.*

Proof. By induction, for all $n \in \mathbb{N}_+$,

- $R^{(n)}$ is reflexive and symmetric;
- $\forall E \in \mathcal{E}_R \quad R^{(n)} \subseteq E$;
- $(\psi, \omega) \in R^{(n)} \iff \exists f [f: \mathbb{N} \rightarrow \Omega \text{ and } f(0) = \psi \text{ and } f(n) = \omega \text{ and } \forall k < n \quad (f(k), f(k+1)) \in R]$.

By (11), R^+ is reflexive and symmetric and $\forall E \in \mathcal{E}_R \quad R^+ \subseteq E$. Therefore $R^+ \subseteq \bigcap \{E \mid E \in \mathcal{E}_R\}$. If $(\chi, \psi) \in R^+$ and $(\psi, \omega) \in R^+$, then, for some m and n , $(\chi, \psi) \in R^{(m)}$ and $(\psi, \omega) \in R^{(n)}$. Therefore there are functions $f: \mathbb{N} \rightarrow \Omega$ and $g: \mathbb{N} \rightarrow \Omega$ such that $f(0) = \chi$, $f(m) = g(0) = \psi$, $g(n) = \omega$, $\forall k < m \quad [(f(k), f(k+1)) \in R]$, and $\forall k < n \quad [(f(k), f(k+1)) \in R]$. Defining $h(k) = f(k)$ if $k < m$ and $h(k) = g(k)$ if $k \geq m$ shows that $(\chi, \omega) \in R^{(m+n)} \subseteq R^+$, so R^+ is transitive (and an equivalence relation). Thus $R^+ \in \mathcal{E}_R$, so $R^+ = \bigcap \{E \mid E \in \mathcal{E}_R\}$. \square

Lemma 4. *If $R \in \Sigma_1^1(\Omega \times \Omega)$, then $R^+ \in \Sigma_1^1(\Omega \times \Omega)$.*

Proof. It will be proved by induction that $R^{(n)} \in \Sigma_1^1(\Omega \times \Omega)$. $R^{(1)} \in \Sigma_1^1(\Omega \times \Omega)$ by assumption. Define $f: \Omega \times \Omega \times \Omega \rightarrow \Omega \times \Omega \times \Omega \times \Omega$ by $f(\chi, \omega, \psi) = (\chi, \psi, \psi, \omega)$. This function is continuous for Polish topologies that generate the Borel algebras on the three-fold and four-fold products of Ω , so it is a Borel function. Let $G = \Omega \times D \times \Omega$. Consider the induction hypothesis that $R^{(n)} \in \Sigma_1^1(\Omega \times \Omega)$. Then, by lemma 1, $(R \times R^{(n)}) \cap G \in \Sigma_1^1(\Omega \times \Omega \times \Omega \times \Omega)$. Define $H^{(n)} = \{(\chi, \omega, \psi) \mid (\chi, \psi) \in R \text{ and } (\psi, \omega) \in R^{(n)}\}$. By lemma 1, $H^{(n)} = f^{-1}((R \times R^{(n)}) \cap G) \in \Sigma_1^1(\Omega \times \Omega \times \Omega)$ and $R^{(n+1)} = \{(\chi, \omega) \mid \exists \psi \in \Omega \quad (\chi, \psi, \omega) \in H^{(n)}\} \in \Sigma_1^1(\Omega \times \Omega)$. So, by induction, $R^{(n)} \in \Sigma_1^1(\Omega \times \Omega)$ for all n . Again by lemma 1, $R^+ = \bigcup_{n \in \mathbb{N}} R^{(n)} \in \Sigma_1^1(\Omega \times \Omega)$. \square

5. KNOWLEDGE AND COMMON KNOWLEDGE AS EQUIVALENCE RELATIONS

5.1. Characterizing \mathfrak{K}_1 . Conjecture (1) of section 4.2 follows immediately from lemma 2 and the following proposition.

¹³ \mathbb{N} denotes $\{0, 1, 2, \dots\}$ and \mathbb{N}_+ denotes $\{1, 2, \dots\}$.

Proposition 5. *If the information partition \mathfrak{P}_i of every agent i is determined by a $\Sigma_1^1(\Omega \times \Omega)$ equivalence relation E_i , then the common knowledge partition \mathfrak{P}_I is determined by the $\Sigma_1^1(\Omega \times \Omega)$ equivalence relation*

$$(12) \quad E_I = \left(\bigcup_{i \in I} E_i \right)^+.$$

Every block of \mathfrak{P}_I is an analytic (thus universally measurable) event.

The first step to proving this proposition is to show that E_I is the transitive closure of $\bigcup_{i \in I} E_i$ and that it determines \mathfrak{P}_I . This fact, stated as lemma 6, follows straightforwardly from lemma 3. Since $\bigcup_{i \in I} E_i$ is analytic by lemma 1, proposition 5 is an immediate consequence of lemmas 2, 4, and 6.

Lemma 6. *E_I is an equivalence relation. If π_ω^i is defined by (9) with respect to ω and E_i and π_ω^I is defined by (9) with respect to ω and E_I , then $\pi_\omega^i \subseteq \pi_\omega^I$. If E^* is any equivalence relation such that, for all $i \in I$, $\pi_\omega^i \subseteq \pi_\omega^*$ (where π_ω^* is defined by (9) with respect to ω and E^*), then $\pi_\omega^I \subseteq \pi_\omega^*$. $\mathfrak{P}_I = \{\pi_\omega^I | \omega \in \Omega\}$.*

From this point forward, in light of proposition 5, individual agents' information will be represented as $\Sigma_1^1(\Omega \times \Omega)$ equivalence relations, rather than as partitions. If E determines \mathfrak{P} , then $\kappa_E(A)$ is synonymous with $\kappa_{\mathfrak{P}}(A)$.

5.2. Characterizing $\kappa_E(\mathbf{A})$. The saturation with respect to equivalence relation E of set $A \subseteq \Omega$ is

$$(13) \quad [A]_E = \{\alpha | \exists \beta [(\alpha, \beta) \in E \text{ and } \beta \in A]\}.$$

The following lemma, trivially verified, is stated for subsequent reference.

Lemma 7. *Knowledge (or common knowledge) of event A is related to saturation by*

$$(14) \quad \kappa_E(A) = \Omega \setminus [\Omega \setminus A]_E,$$

where E is E_i or E_I .

Let's investigate where in the projective hierarchy $[A]_E$ and $\kappa_E(A)$ lie, if E is in $\Sigma_1^1(\Omega \times \Omega)$ and A is in either $\Sigma_1^1(\Omega)$ or $\Pi_1^1(\Omega)$. If $A \in \Sigma_1^1(\Omega)$, then there is a Borel $I \subseteq \Omega \times \Omega \times \Omega$ such that $(\alpha, \beta) \in E \iff \exists \gamma (\alpha, \beta, \gamma) \in I$.

Lemma 8. *There is a measurable function $f: \Omega \rightarrow \Omega \times \Omega \times \Omega$ that is 1-1 and onto.*

Proof. Both Ω and $\Omega \times \Omega \times \Omega$ are standard Borel spaces. By the Kuratowski's isomorphism theorem (Moschovakis [2009, theorem 1G.4]), there is a Borel isomorphism between them. \square

In the following proposition, adopt the notation that $f(\omega) = (\omega_{(0)}, \omega_{(1)}, \omega_{(2)})$ is the homeomorphism that lemma 8 asserts to exist. Also, let $t_{23}: \Omega \times \Omega \times \Omega \times \Omega \rightarrow \Omega \times \Omega \times \Omega \times \Omega$ denote the homeomorphism that transposes the second and third factors of the product space.

Proposition 9. *Suppose that $E \in \Sigma_1^1(\Omega \times \Omega)$. If $A \in \Sigma_1^1(\Omega)$, then $[A]_E \in \Sigma_1^1(\Omega)$ and $\kappa_E(A) \in \Pi_2^1(\Omega)$. If $A \in \Pi_1^1(\Omega)$, then $[A]_E \in \Sigma_2^1(\Omega)$ and $\kappa_E(A) \in \Pi_1^1(\Omega)$.*

Proof. Suppose that A is in $\Sigma_1^1(\Omega)$. Specifically suppose that, for some $B \in \mathcal{B}$, $\omega \in A \iff \exists \chi (\omega, \chi) \in B$. Define $C = (I \times \Omega) \cap t_{23}(\Omega \times \Omega \times B)$. $C \in \mathcal{B} \times \mathcal{B} \times \mathcal{B} \times \mathcal{B}$. Define $j: \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B} \times \mathcal{B} \times \mathcal{B} \times \mathcal{B}$ by $j(\omega, \chi) = (\omega, \chi_{(0)}, \chi_{(1)}, \chi_{(2)})$, and define $J = j^{-1}(C)$. $J \in \mathcal{B} \times \mathcal{B}$. Then

$$\begin{aligned}
 \omega \in [A]_E &\iff \exists \chi [(\omega, \chi) \in E \text{ and } \chi \in A] \\
 &\iff \exists \chi [(\omega, \chi_{(0)}, \chi_{(1)}) \in I \text{ and } (\chi_{(0)}, \chi_{(2)}) \in B] \\
 (15) \quad &\iff \exists \chi [(\omega, \chi_{(0)}, \chi_{(1)}, \chi_{(2)}) \in (I \times \Omega) \cap t_{23}(\Omega \times \Omega \times B)] \\
 &\iff \exists \chi [(\omega, \chi) \in J].
 \end{aligned}$$

Hence $[A]_E \in \Sigma_1^1(\Omega)$. Consequently, if instead $A \in \Pi_1^1(\Omega)$, then $\kappa_E(A) = \Omega \setminus [\Omega \setminus A]_E \in \Pi_1^1(\Omega)$ by lemma 7.

Suppose instead that A is in $\Pi_1^1(\Omega)$. Specifically suppose that, for some $B \in \mathcal{B}$, $\omega \in A \iff \forall \psi (\omega, \psi) \in B$. Define $k: B \times B \times B \rightarrow B \times B \times B \times B$ by $k(\omega, \chi, \psi) = (\omega, \chi_{(0)}, \chi_{(1)}, \psi)$, and define $K = k^{-1}(C)$. $K \in \mathcal{B} \times \mathcal{B} \times \mathcal{B}$. Then

$$\begin{aligned}
 \omega \in [A]_E &\iff \exists \chi [(\omega, \chi) \in E \text{ and } \chi \in A] \\
 &\iff \exists \chi \forall \psi [(\omega, \chi_{(0)}, \chi_{(1)}) \in I \text{ and } (\chi_{(0)}, \psi) \in B] \\
 (16) \quad &\iff \exists \chi \forall \psi [(\omega, \chi_{(0)}, \chi_{(1)}, \psi) \in (I \times \Omega) \cap t_{23}(\Omega \times \Omega \times B)] \\
 &\iff \exists \chi \forall \psi [(\omega, \chi, \psi) \in K].
 \end{aligned}$$

Hence $[A]_E \in \Sigma_2^1(\Omega)$. Consequently, if instead $A \in \Sigma_1^1(\Omega)$, then $\kappa_E(A) = \Omega \setminus [\Omega \setminus A]_E \in \Pi_2^1(\Omega)$ by lemma 7. \square

6. AUMANN'S AGREEMENT THEOREM

6.1. Framing the issue. Roughly speaking, Aumann has proved for a countable, standard Borel, state space that, if the posterior probabilities that two agents assign to some event are common knowledge between them at some state, then both agents must assign the same posterior probability to the event at that state.

If a standard Borel space is countable, then every singleton is a Borel event, every arbitrary subset of the space is a countable union of singletons, and so every subset of the space is a Borel event. In contrast, as has already been mentioned, every uncountable standard Borel space has non-Borel subsets. Some of those events are not measurable with respect to any nonatomic probability measure.

Aumann's proof makes use of this special feature of a countable space. Specifically, the proof incorporates the assumption that the event of agents having common knowledge is a Borel event. The question is, is this assumption a mathematical convenience to provide a simple proof of a special case of a general result, or is the agreement theorem a fragile proposition that cannot be extended to apply to the uncountable state spaces in which models in decision theory, game theory, and economics are framed?

Proposition 9 suggests the answer to this question may depend on what is the class of events, of which agents' common knowledge is to be considered. If that class is the class of all analytic events, then the only information that proposition 9 provides in general, is that common knowledge events are in $\Pi_2^1(\Omega)$. It turns out that, apparently, whether or not every such event is measurable is independent of the axioms of set theory that are typically invoked in decision theory and related

fields.¹⁴ If an event of agents having common knowledge is not measurable, then there is little hope of generalizing Aumann's proof to cover the situation. That is, Aumann's analysis involving a Borel event $\kappa_I(A)$ could perhaps be generalized to an analysis involving a measurable event, but it clearly could not be generalized to apply to a non-measurable event.

Therefore, to answer the question of whether or not Aumann's result is a fragile one, two more basic questions may have to be answered. The first question is whether or not *every* event in $\mathbf{\Pi}_2^1(\Omega)$ can be obtained as $\kappa_I(A)$ for some analytic event A ? (Conceivably, all such events $\kappa_I(A)$ lie in a sub-class of $\mathbf{\Pi}_2^1(\Omega)$ that consists entirely of measurable events.) This is an open question, and its answer may also be independent of ZFC set theory. If the answer to the first question is affirmative or depends on "strong" set-theoretic axioms beyond ZFC, then the second question is a philosophical question: which set-theoretic axioms beyond ZFC are the appropriate ones to adopt as a framework for studying knowledge and belief in a community of rational agents?

But, in practice, the questions about knowledge and common knowledge that arise in the modeling of rational agents, and of communities of rational agents, predominantly have to do with common knowledge of Borel events.¹⁵ By lemma 1, those events are in $\mathbf{\Delta}_1^1(\Omega)$, hence in $\mathbf{\Sigma}_1^1(\Omega)$. By proposition 9, then, the event of an agent having knowledge of such an event, or of the event being common knowledge among a community of agents, is in $\mathbf{\Sigma}_1^1(\Omega)$. By lemma 1, then, that event is universally measurable. So, if common knowledge of Borel events is all that needs to be considered, then the most obvious reason to suspect that the agreement theorem cannot be generalized is sidestepped. But a difficult problem remains, having to do with the conditional-probability representation of agents' posterior beliefs. A review of Aumann's proof will be helpful for stating this problem.

6.2. Formal statement of Aumann's theorem. To facilitate subsequent comparison with generalizations, it will be helpful to re-state and prove Aumann's [1976] version of the agreement theorem. Consider a countable probability space $(\mathbb{N}, 2^{\mathbb{N}}, \mu)$ and let $I = \{1, 2\}$. For $i \in I$, let $\mathfrak{P}_i = \{\pi_k^i\}_{k \in N_i}$, where $N_i \subseteq \mathbb{N}$, with $j \neq k \implies \pi_j^i \cap \pi_k^i = \emptyset$ for all j and k in N_i .¹⁶

Consider an event $A \subseteq \mathbb{N}$.

If μ represents the prior beliefs common to the two agents, then the posterior probability that agent i assigns to A in event $\pi \in \mathfrak{P}_i$ is

$$(17) \quad \mu_i(A|\pi) = \mu(A \cap \pi) / \mu(\pi)$$

if $\mu(\pi) > 0$.

Theorem 10 (Aumann's agreement theorem). *If*

- (1) $q_1 \in [0, 1]$ and $q_2 \in [0, 1]$, and
- (2) $\mu(\{\pi\}) > 0$ for every $\pi \in \mathfrak{P}_1 \cup \mathfrak{P}_2$, and

¹⁴These are the ZFC axioms, including the axiom of choice. Independence is a consequence of Jech [2002, corollary 25.28, theorem 26.14]. The question is *apparently* independent because Theorem 26.14 has a hypothesis that cannot be proved to be consistent with ZFC.

¹⁵In some cases, questions regarding events that are not Borel do arise naturally. Cf. Stinchcombe and White [1992].

¹⁶ N_i is introduced, rather than simply using \mathbb{N} as the index set, because the partition might possibly be finite.

- (3) $Q \subseteq \mathbb{N}$ is defined by $Q = [\bigcup\{\pi \in \mathfrak{P}_1 | \mu_1(A|\pi) = q_1\} \cap [\bigcup\{\pi \in \mathfrak{P}_2 | \mu_2(A|\pi) = q_2\}]$,
and
(4) $\kappa_I(Q) \neq \emptyset$,

then $q_1 = q_2$.

Proof. First observe that, for any sets $B \subseteq S$ and partition \mathfrak{P} of S , $[S \setminus [B]_{\mathfrak{P}}]_{\mathfrak{P}} = S \setminus [B]_{\mathfrak{P}}$. Therefore, taking $B = \Omega \setminus Q$ and $S = \Omega$ and $\mathfrak{P} = \mathfrak{P}_I$,

$$(18) \quad [\kappa_I(Q)]_{\mathfrak{P}_I} = \kappa_I(Q).$$

Since each \mathfrak{P}_i refines \mathfrak{P}_I , $[\kappa_I(Q)]_{\mathfrak{P}_i} = \kappa_I(Q)$ for each i . Therefore, for $i \in I$, there is a subset $M_i \subseteq N_i$ such that $\bigcup_{n \in M_i} \pi_n^i = \kappa_I(Q)$. Note that $q_i = \mu(A|\pi_n^i)$ for all $n \in M_i$.

$$(19a) \quad q_1 \mu(\kappa_I(Q)) = q_1 \sum_{n \in M_1} \mu(\pi_n^1)$$

$$(19b) \quad = \sum_{n \in M_1} \mu(A|\pi_n^1) \mu(\pi_n^1)$$

$$(19c) \quad = \mu(A \cap \kappa_I(Q))$$

$$(19d) \quad = \sum_{n \in M_2} \mu(A|\pi_n^2) \mu(\pi_n^2)$$

$$(19e) \quad = q_2 \sum_{n \in M_2} \mu(\pi_n^2)$$

$$(19f) \quad = q_2 \mu(\kappa_I(Q))$$

Because $\mu(B) > 0$ for every non-empty subset B of \mathbb{N} and $\kappa_I(Q) \neq \emptyset$, the equation between the left side of (19a) and (19f) can be divided by $\mu(\kappa_I(Q))$, yielding that $q_1 = q_2$. \square

7. GENERALIZING THE AGREEMENT THEOREM TO UNCOUNTABLE SPACES

7.1. A futile attempt. In general, when \mathfrak{P} is an analytic partition of an uncountable standard Borel space, then the conditional probability (conditioning from prior probability μ) of an event A with respect to \mathfrak{P} is defined in two steps. First, a σ -algebra $\mathcal{B}_{\mathfrak{P}}$ of μ -measurable subsets of Ω that are observable with respect to the partition is defined. Specifically, $\mathcal{B}_{\mathfrak{P}}$ can be specified as the sub- σ -algebra of \mathcal{B} that comprises the events in \mathcal{B} that can be tiled by blocks of \mathfrak{P} .¹⁷

$$(20) \quad \mathcal{B}_{\mathfrak{P}} = \left\{ B | B \in \mathcal{B} \text{ and } \exists \mathcal{T} \subseteq \mathfrak{P} \ B = \bigcup \mathcal{T} \right\} = \{ B | B \in \mathcal{B} \text{ and } B = [B]_{\mathfrak{P}} \}.$$

Second, a $\mathcal{B}_{\mathfrak{P}}$ -measurable random variable $a: \Omega \rightarrow [0, 1]$ such that

$$(21) \quad \forall B \in \mathcal{B}_{\mathfrak{P}} \quad \int_B a(\omega) d\mu(\omega) = \mu(A \cap B).$$

is asserted by the Radon-Nikodym theorem to exist.¹⁸

Typically, that random variable is nonatomically distributed. That is, for any q , $\mu(\{\omega | a(\omega) = q\}) = 0$. In that case, if Q is defined as above, then $\mu(\kappa_I(Q)) = 0$. Therefore it should be anticipated that the hypothesis of theorem 10 will virtually

¹⁷Blackwell [1956] and Blackwell and Dubins [1975] study conditional probability in relation to σ -algebras of this form.

¹⁸This construction is due to Kolmogorov [1956, chapter IV]. Any two random variables satisfying (21) agree almost surely.

never be satisfied in the setting of an uncountable Borel space. A more general statement of the theorem is called for. But merely to generalize the statement will not resolve conceptual issues regarding the meaning of the theorem or the logic of its proof. Therefore, a first step to extend the theorem is to provide a proof of a result that differs from theorem 10 only by the substitution of an uncountable standard Borel space for \mathbb{N} .

Many results can be generalized from the context of a countable probability space to that of an uncountable standard Borel space with a nonatomic measure by a routine procedure of transforming probability-weighted summations into integrals with respect to a density. Consider how (19) in the proof of theorem 10 would be transformed in that way.

Conjecture. *If*

- (1) $q_1 \in [0, 1]$ and $q_2 \in [0, 1]$, and
- (2) (Ω, \mathcal{B}) is an uncountable standard Borel space and $\mu: \mathcal{B} \rightarrow [0, 1]$ is a nonatomic probability measure, and
- (3) the random variables a_1 and a_2 satisfy (21) with respect to $\mathcal{B}_{\mathfrak{P}_1}$ and $\mathcal{B}_{\mathfrak{P}_2}$ respectively, and
- (4) $Q = a_1^{-1}(\{q_1\}) \cap a_2^{-1}(\{q_2\})$, and
- (5) $\mu^*(\kappa_I(Q)) > 0$,¹⁹

then $q_1 = q_2$.

Attempted proof. $Q \in \mathcal{B}$, so, by lemma 1 and proposition 9, $\kappa_I(Q) \in \mathbf{\Pi}_1^1(\Omega)$. Therefore $\kappa_I(Q)$ is universally measurable, and is thus in the domain of μ^* as hypothesis 5 implies.²⁰

Because a co-analytic set is universally measurable, the Lebesgue integral of a Borel measurable function over a co-analytic domain, $K \in \mathbf{\Pi}_1^1(\Omega)$, is uniquely defined. Specifically, because K is universally measurable by lemma 1, there are Borel events J and L such that $J \subseteq K \subseteq L$ and $\mu(J) = \mu(L)$. For every Borel-measurable $f: \Omega \rightarrow \mathbb{R}_+$, $\int_J f(\omega) d\mu(\omega) = \int_L f(\omega) d\mu(\omega)$. Thus $\int_K f(\omega) d\mu(\omega)$ should be defined by

$$(22) \quad \int_K f(\omega) d\mu(\omega) = \int_J f(\omega) d\mu(\omega) = \int_L f(\omega) d\mu(\omega).$$

This definition of the integral is extended in the usual way to functions of form $f - g$, where both f and g are nonnegative-valued and have finite integrals. Since $\kappa_I(Q) \in \mathbf{\Pi}_1^1(\Omega)$, integration over $\kappa_I(Q)$ is well defined.

¹⁹Recall that μ^* denotes the extension by outer measure of μ to the σ -algebra that completes \mathcal{B} with respect to μ . To streamline notation in the following discussion, μ will be used to refer to its extension, μ^* . Wherever ' $\mu(B)$ ' is written but the possibility that $B \notin \mathcal{B}$ is countenanced, a reference to $\mu^*(B)$ is implicit.

²⁰Lemma 1 asserts that $\mathbf{\Delta}_1^1(\Omega) \subsetneq \mathbf{\Pi}_1^1(\Omega)$, so μ must be interpreted here as a measure completion that extends the domain, \mathcal{B} , of μ to a larger σ -algebra.

On that account, and because $\kappa_I(Q) \subseteq Q \subseteq a_i^{-1}(\{q_i\})$, the steps at (23a), (23b), (23e), and (23f) of the following argument are sound.

$$\begin{aligned}
(23a) \quad q_1 \mu(\kappa_I(Q)) &= q_1 \int_{\kappa_I(Q)} 1 \, d\mu(\omega) \\
(23b) \quad &= \int_{\kappa_I(Q)} a_1(\omega) \, d\mu(\omega) \\
(23c) \quad &= \mu(A \cap \kappa_I(Q)) \\
(23d) \quad &= \int_{\kappa_I(Q)} a_2(\omega) \, d\mu(\omega) \\
(23e) \quad &= q_2 \int_{\kappa_I(Q)} 1 \, d\mu(\omega) \\
(23f) \quad &= q_2 \mu(\kappa_I(Q))
\end{aligned}$$

Hypothesis 5 justifies dividing the left side of line (23a) and line (23f) by $\mu(\kappa_I(Q))$, which yields that $q_1 = q_2$. ?

7.2. A gap and a counterexample. There is a gap in this attempted proof where it is asserted, at stages (23c) and (23d), that

$$\int_{\kappa_I(Q)} a_1(\omega) \, d\mu(\omega) = \mu(K \cap \kappa_I(Q)) = \int_{\kappa_I(Q)} a_2(\omega) \, d\mu(\omega).$$

These stages may seem to be justified by the defining condition (21) of conditional probability, but that condition is not necessarily applicable. If $\kappa_I(Q) \in \mathcal{B}$, then $\kappa_I(Q) \in \mathcal{B}_{\mathfrak{F}}$ by (18) and (20). However, all that has been proved (in proposition 9) is that $\kappa_I(Q) \in \mathbf{\Pi}_1^1(\Omega)$. Since $\mathcal{B}_{\mathfrak{F}_i} \subseteq \mathcal{B} \subsetneq \mathbf{\Pi}_1^1(\Omega)$ (by lemma 1), $\kappa_I(Q)$ (as the value of B) might not be in the domain of the quantification over $\mathcal{B}_{\mathfrak{F}_i}$ (as the value of $\mathcal{B}_{\mathfrak{F}}$) in (21).²¹ In that case, (20) cannot be used to argue in a single step that $\int_{\kappa_I(Q)} a_1(\omega) \, d\mu(\omega) = \mu(A \cap \kappa_I(Q))$ (that is, step (23c)) or that the corresponding assertion in (23d) holds.

It might be hoped that a three-step argument for step (23c) (and analogously for step (23d)) would work. Find an event $J \in \mathcal{B}_{\mathfrak{F}}$ and an event $L \in \mathcal{B}$ such that $J \subseteq \kappa_I(Q) \subseteq L$ and $\mu(J) = \mu(L)$. From (22), argue that $\int_J a_1(\omega) \, d\mu(\omega) = \mu(A \cap J)$. Finally, infer from $\mu(J) = \mu(L)$ and (22), and from (6) applied to $A \cap \kappa_I(Q)$, that $\int_{\kappa_I(Q)} a_1(\omega) \, d\mu(\omega) = \mu(A \cap \kappa_I(Q))$. But this hope is in vain.

Proposition 11. *The universal measurability of $\kappa_I(Q)$ entails that there are events J and L in \mathcal{B} such that $J \subseteq \kappa_I(Q) \subseteq L$ and $\mu(J) = \mu(L)$, but it does not entail that J can be chosen from $\mathcal{B}_{\mathfrak{F}}$, even though $\kappa_I(Q) = [\kappa_I(Q)]_{\mathfrak{F}}$.*

This proposition is an immediate consequence of the following lemma, which shows that, if one tries to approximate an event in $\mathbf{\Pi}_1^1(\Omega) \setminus \mathcal{B}$ from inside and from outside in this way, there may be no approximation except by countable and co-countable (and hence of probabilities 0 and 1 with respect to the nonatomic measure μ) Borel events.

²¹An example in which $\kappa_I(Q) \notin \mathcal{B}$ can be obtained by adapting the construction provided in Green [2012].

Lemma 12. *For every $K \in (\Sigma_1^1(\Omega) \cup \Pi_1^1(\Omega)) \setminus \mathcal{B}$, if both K and $\Omega \setminus K$ have cardinality 2^{\aleph_0} (the same cardinality as Ω has), then there is a partition \mathfrak{P} of Ω such that*

- (1) $K = [K]_{\mathfrak{P}}$;
- (2) \mathfrak{P} consists of singletons and pairs;
- (3) if $B \in \mathcal{B}_{\mathfrak{P}}$ and $B \subseteq K$, then B is countable;
- (4) if $C \in \mathcal{B}_{\mathfrak{P}}$ and $K \subseteq C$, then $\Omega \setminus C$ is countable.

Proof. The first step is to construct a partition Ω that satisfies conclusions 1–3. Define $\mathcal{U} = \{B \mid B \in \mathcal{B} \text{ and } B \subseteq K \text{ is uncountable}\}$. Every $B \in \mathcal{U} \cup \{\emptyset\}$ has cardinality 2^{\aleph_0} , and \mathcal{U} itself has cardinality 2^{\aleph_0} .²² Let κ denote the first ordinal of this cardinality.

Let $f: \kappa \rightarrow \mathcal{U}$ be 1–1 and onto.

Let $g: 2^\Omega \setminus \{\emptyset\} \rightarrow \Omega$ be a choice function.

Define Ω recursively.

- $\Omega_0 = \emptyset$.
- If $\alpha < 2^{\aleph_0}$, then

$$\Omega_{\alpha+1} = \Omega_\alpha \cup \{g(f(\alpha) \setminus \bigcup \Omega_\alpha), g(K \setminus (\bigcup \Omega_\alpha \cup f(\alpha)))\}.$$
- If $\lambda \leq \kappa$ is a limit ordinal, then $\Omega_\lambda = \bigcup_{\alpha < \lambda} \Omega_\alpha$.

For $\alpha < \kappa$, define the pair $\pi_\alpha = \Omega_{\alpha+1} \setminus \Omega_\alpha$. Note the general principle that, if $\iota < \kappa$ and, for every $\theta < \iota$, F_θ is finite, then the cardinality of $\bigcup_{\theta < \iota} F_\theta$ is less than 2^{\aleph_0} . Thus $\Omega_\alpha = \Omega_0 \cup \bigcup_{\beta < \alpha} \pi_\beta$ has cardinality less than 2^{\aleph_0} for every $\alpha < \kappa$. Consequently the sets $f(\alpha) \setminus \bigcup \Omega_\alpha$ and $K \setminus (\bigcup \Omega_\alpha \cup f(\alpha))$ are both non empty, so this recursion is well defined.

$\bigcup \Omega_\kappa \subseteq K$.

Define $\Omega = \Omega_\kappa \cup \{\{\omega\} \mid \omega \in \Omega \setminus \bigcup \Omega_\kappa\}$.

$K = [K]_\Omega$ because $K = (\bigcup \Omega_\kappa) \cup \{\{\omega\} \mid \omega \in K \setminus \bigcup \Omega_\kappa\}$.

Suppose that $B \in \mathcal{U} \cap \mathcal{B}_\Omega$. Then $B = f(\alpha)$ for some $\alpha < 2^{\aleph_0}$

Therefore $\{g(B \setminus \bigcup \Omega_\alpha), g(K \setminus (\bigcup \Omega_\alpha \cup B))\} \in \Omega_{\alpha+1} \subseteq \Omega$.

Hence $g(K \setminus (\bigcup \Omega_\alpha \cup B)) \in [B] = B$.

This is a contradiction, since g , as a choice function, must satisfy $g(K \setminus (\bigcup \Omega_\alpha \cup B)) \in K \setminus (\bigcup \Omega_\alpha \cup B)$.

Similarly, define a partition \mathfrak{R} such that $[\Omega \setminus K]_{\mathfrak{R}} = \Omega \setminus K$ and, for every uncountable subset B of $\Omega \setminus K$, $[B]_{\mathfrak{R}} \neq B$.

Finally, define $\mathfrak{P} = (\Omega \upharpoonright K) \cup (\mathfrak{R} \upharpoonright (\Omega \setminus K))$.

Since $\mathfrak{P} \upharpoonright K = \Omega \upharpoonright K$, Ω satisfies conclusions 1–3, and those conclusions concern subsets of K , conclusion 4 is all that remains to be proved. Suppose that $K \subseteq C \in \mathcal{B}$, and let $B = \Omega \setminus C$. Then $B \subseteq \Omega \setminus K$. By construction of \mathfrak{P} and \mathfrak{R} , $[B]_{\mathfrak{P}} = [B]_{\mathfrak{R}}$ and therefore $B = [B]_{\mathfrak{P}} \implies B$ is countable. Since $C = [C]_{\mathfrak{P}} \implies B = [B]_{\mathfrak{P}}$, $C = [C]_{\mathfrak{P}} \implies \Omega \setminus C$ is countable. \square

If $\kappa_I(Q) \in \Pi_1^1(\Omega) \setminus \mathcal{B}$, and if both $\kappa_I(Q)$ and $\Omega \setminus \kappa_I(Q)$ have cardinality 2^{\aleph_0} and $\mathcal{B}_{\mathfrak{P}_1}$ is related to $\kappa_I(Q)$ as in lemma 12, then the situation described in proposition 11 obtains. That is, the argument via (23) for the conjectured generalization of Aumann's agreement theorem is unsound.

²²The first assertion is Moschovakis [2009, corollary 2C.3]. The second assertion follows quickly from that result and Moschovakis [2009, 1E.3]. Cardinal and ordinal numbers and transfinite recursion are explained in chapters 1–2 of Jech [2002]

7.3. Proposition 11 reconsidered. The partition, constructed in the proof of lemma 12, which leads to a difficulty of measurability despite consisting of innocuous sets (singletons and pairs), strongly resembles the problematic partition constructed in section 2.2. While the earlier counterexample has to do with violation of measurability, lemma 12 illuminates a more subtle problem: that a universally measurable set that is saturated with respect to a partition may not be approximable in measure by Borel sets that are saturated with respect to it.

Again in this case, a sound analysis of an uncountable state space requires that agents' information partitions must be constrained by conditions stated in terms of the equivalence relations that determine them. Proposition 14, below, is identical to the conjecture framed in section 7.1, except that two new hypotheses have been added. Hypothesis 6 states that each agent's information partition is determined by an analytic equivalence relation, and hypothesis 7 states that each of those partitions consists of countable blocks. Recall that the counterexample constructed in lemma 12 consists of countable blocks.

The proof of proposition 14 will involve the notion of a relation being *Borel measurable*. $R \subseteq \Omega \times \Omega$ possesses this property if, for every $B \in \mathcal{B}$, $\{\psi | \exists \omega \in B (\psi, \omega) \in R\} \in \mathcal{B}$. Note that a countable union of Borel-measurable relations is Borel measurable.

Lemma 13. *If E_i is Borel measurable for each $i \in I$, and if $Q \in \mathcal{B}$, then $\kappa_I(Q) \in \mathcal{B}$.*

Proof. Define $R = \bigcup_{i \in I} E_i$. R , a union of finitely many Borel-measurable relations, is Borel measurable. By proposition 5, $E_I = R^+$. Define $B = \Omega \setminus Q \in \mathcal{B}$ and $K_0 = B$. For $m > 0$, $K_m = \{\psi | \exists \omega \in B (\psi, \omega) \in R^{(m)}\}$. Note that $K_{m+1} = \{\psi | \exists \omega \in K_m (\psi, \omega) \in R\}$, which is a Borel event by Borel-measurability of R . Thus, by induction on m , for any $B \in \mathcal{B}$, $K_m \in \mathcal{B}$. $[B]_{\mathfrak{P}_I} = \bigcup_{m \in \mathbb{N}} K_m \in \mathcal{B}$. Therefore $\kappa_I(Q) = \Omega \setminus [B]_{\mathfrak{P}_I} \in \mathcal{B}$. \square

Proposition 14. *If*

- (1) $q_1 \in [0, 1]$ and $q_2 \in [0, 1]$, and
- (2) (Ω, \mathcal{B}) is an uncountable standard Borel space and $\mu: \mathcal{B} \rightarrow [0, 1]$ is a nonatomic probability measure, and
- (3) the random variables a_1 and a_2 satisfy (21) with respect to $\mathcal{B}_{\mathfrak{P}_1}$ and $\mathcal{B}_{\mathfrak{P}_2}$ respectively, and
- (4) $Q = a_1^{-1}(\{q_1\}) \cap a_2^{-1}(\{q_2\})$,
- (5) $\mu^*(\kappa_I(Q)) > 0$,
- (6) each agent's information partition \mathfrak{P}_i is determined by an equivalence relation $E_i \in \Sigma_1^1(\Omega \times \Omega)$, and
- (7) for each $i \in I$ and for every $\omega \in \Omega$, the block $[\{\omega\}]_{\mathfrak{P}_i}$ is countable,

then $q_1 = q_2$.

Proof. By hypotheses 6 and 7, and by Moschovakis [2009, exercises 2E.4 and 4F.17], the graph of each E_i is the union of a countable set $\{f_{ik}: \Omega \rightarrow \Omega\}_{k \in \mathbb{N}}$ of Borel-measurable functions.²³ The graph of each of these functions is a Borel-measurable relation. E_i , the union of this countable set, is then Borel measurable, so lemma 13 implies that $\kappa_I(Q) \in \mathcal{B}$. Therefore the calculation (23) is sound in this case. \square

²³Moschovakis [2009, exercise 4F.17] is a recursion-theoretic version of the result needed here. Moschovakis [2009, theorem 3D.7, theorem 3E.4] explains how it justifies the needed result.

7.4. Generalization to a sigma-compact Polish space. Proposition 14 is an extension of the agreement theorem that places no restriction on the space Ω of states of the world, except that it should be a standard Borel space. It places a mild restriction, that is, to be analytic, on agents' information partitions. However, it also places a much tighter restriction on those partitions, that they should consist of countable blocks. For the purpose of shedding light on a counterexample with countable blocks (that is, the one in lemma 12), that restriction is innocuous. However, it would be untenable in settings, envisioned in the introduction of this article and in section 7.1, in which the type of one agent is nonatomically distributed conditional on the type of the other. For example, if each of two agents privately observes a normal random variable, and if those observations are not perfectly correlated, then there must be a continuum of states (corresponding to the possible observations by agent 2) in which agent 1 observes 0.

The following proposition places significant restrictions on Ω and the relations E_i , but ones that could be satisfied in many prospective applications, including to finite-strategy games of incomplete information (discussed below in section 7.5).

Proposition 15. *If*

- (1) $q_1 \in [0, 1]$ and $q_2 \in [0, 1]$, and
- (2) (Ω, \mathcal{B}) is an uncountable standard Borel space, \mathcal{B} is generated by a metric that induces a σ -compact topology on Ω ,²⁴ and $\mu: \mathcal{B} \rightarrow [0, 1]$ is a nonatomic probability measure, and
- (3) the random variables a_1 and a_2 satisfy (21) with respect to $\mathcal{B}_{\mathfrak{P}_1}$ and $\mathcal{B}_{\mathfrak{P}_2}$ respectively, and
- (4) $Q = a_1^{-1}(\{q_1\}) \cap a_2^{-1}(\{q_2\})$,
- (5) $\mu^*(\kappa_I(Q)) > 0$,
- (6) each agent's information partition \mathfrak{P}_i is determined by an equivalence relation E_i that is closed in the topology on Ω

then $q_1 = q_2$.

Proof. By Aliprantis and Border [2006, theorem 18.20], $R = \bigcup_{i \in I} E_i$ is Borel measurable. Therefore lemma 13 implies that $\kappa_I(Q) \in \mathcal{B}$, and the calculation (23) is sound. \square

7.5. Compact Polish universal belief-type spaces. The purpose of this section is to sketch a proof that proposition 15 is applicable to a class of environments that includes finite-strategy games of incomplete information.

Recall that, roughly speaking, a *state of nature* completely specifies the objective environment (such as the feasible strategy profiles of a game and the players' payoffs from each profile). A *state of the world* completely specifies both the state of nature and agent's beliefs (including higher-order beliefs about the beliefs of others). A *type space* is a set of states of the world. Mertens and Zamir [1985] construct, from compact space S of states of nature, an essentially unique *universal type space* into which any other space of states of the world (that are genuinely distinct from one another) incorporating these states of nature can be embedded. Brandenburger and Dekel [1993] provide an analogous construction based on a Polish space S . Both constructions represent the universal type space as an inverse limit of a sequence of finite-level belief hierarchies. Thus, if S is a compact Polish space (that is, a space

²⁴A topology on Ω is σ -compact if Ω is a countable union of compact sets.

to which both constructions apply), then, because the inverse limit is essentially unique, the two constructions yield the same type space.

The inverse limit of a sequence of Hausdorff spaces is a closed subset of a countable product of Hausdorff spaces. Therefore, for any class of Hausdorff spaces that is closed under taking both countable products and closed subspaces, the inverse limit of a sequence of spaces in that class belongs to the class. Both compact spaces and Polish spaces satisfy those conditions.²⁵ Thus, if S is a compact Polish space, then its universal type space is also a compact Polish space, so proposition 15 applies.

8. CONCLUSION

From the proof of proposition 15, it is clear that the proposition is mathematically not very general. Specifically, the proof shows that $\kappa_I(Q)$ is a Borel event under its hypotheses. In contrast, proposition 9 restricts $\kappa_I(Q)$ only to be co-analytic, and it follows from lemma 1 and Green [2012, corollary 4.3] that there is a non-Borel common-knowledge event, even under the closedness assumption 5 of proposition 15. Thus assumption 2 of the proposition, although it superficially has nothing to do with common knowledge, actually places a binding constraint on the class of common knowledge events. There is little to choose, between Aumann's device of assuming that information partitions are countable in order to ensure that common-knowledge events must be Borel and the present device of assuming that state spaces are σ -compact in order to coerce the same result.

Of course the goal is not mathematical generality for its own sake, but relevance to economics. Proposition 15 is also restrictive in that sense. There are practical situations to which it does not apply, such as continuous-time asset trading. The mathematical framework for studying that situation is the theory of continuous-time stochastic processes, of which diffusion processes built from a Wiener process are the most tractable sub-class. But the canonical Wiener process is a time-indexed family of random variables defined on the space of continuous, real-valued functions with the topology of uniform convergence on compact sets. That topological space is not σ -compact.

Possibly a more satisfactory framework for studying probabilities of events concerning knowledge and common knowledge would be the concept of a countably additive, normal conditional probability introduced by Blackwell and Dubins [1975]. The idea would be to take $\mathcal{C} \subset 2^\Omega$ to be a σ -algebra to which all analytic events belong, and to specify $\mathcal{B}_{\mathfrak{F}}$ to be a sub- σ -algebra of \mathcal{C} , rather than of \mathcal{B} .²⁶

$$(24) \quad \mathcal{B}_{\mathfrak{F}} = \left\{ B \mid B \in \mathcal{C} \text{ and } \exists \mathcal{T} \subseteq \mathfrak{F} \ B = \bigcup \mathcal{T} \right\} = \{ B \mid B \in \mathcal{C} \text{ and } B = [B]_{\mathfrak{F}} \}.$$

Then, the random variable a in (21) would be specified to be measurable with respect to $\mathcal{B}_{\mathfrak{F}}$ as defined in (24), rather than as defined in (20).

²⁵The condition regarding closed subspaces is immediate from the definition of the relative topology. The condition regarding countable products follows from Aliprantis and Border [2006, theorem 2.61 (Tychonoff), corollary 3.39].

²⁶Blackwell and Dubins specify \mathcal{C} to be Lebesgue's completion of \mathcal{B} by all subsets of Borel events of μ measure 0.

REFERENCES

- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, third edition, 2006.
- Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976. URL <http://projecteuclid.org/euclid.aos/1176343654>.
- Robert J. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3):263–300, 1999a. URL www.ma.huji.ac.il/raumann/pdf/Interactive%20epistemology1.pdf.
- Robert J. Aumann. Interactive epistemology II: Probability. *International Journal of Game Theory*, 28(3):301–314, 1999b. URL www.ma.huji.ac.il/raumann/pdf/Interactive%20epistemology2.pdf.
- Elchanan Ben-Porath. Rationality, Nash equilibrium, and backwards induction in perfect-information games. *Review of Economic Studies*, 64(1):23–46, 1997.
- Dimitri P. Bertsekas and Steven E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Academic Press, 1978. URL <http://web.mit.edu/dimitrib/www/soc.html>.
- David Blackwell. On a class of probability spaces. In *Proc. Third Berkeley Symp. on Math. Statist. and Prob.*, volume 2, pages 1–6. University of California Press, 1956. URL <http://projecteuclid.org/euclid.bsmmsp/1200502002>.
- David Blackwell and Lester E. Dubins. On existence and non-existence of proper, regular, conditional distributions. *The Annals of Probability*, 3(5):741–752, 1975. URL <http://projecteuclid.org/euclid.aop/1176996261>.
- Adam Brandenburger and Eddie Dekel. Common knowledge with probability 1. *Journal of Mathematical Economics*, 16(3):237–245, 1987.
- Adam Brandenburger and Eddie Dekel. Hierarchies of belief and common knowledge. *Journal of Economic Theory*, 59(1):189–198, 1993.
- Michael J. Brennan and H. Henry Cao. Information, trade, and derivative securities. *Review of Financial Studies*, 9(1):163–208, 1996.
- David Easley, Maureen O'Hara, and P. S. Srinivas. Open volume and stock prices: evidence on where informed traders trade. *Journal of Finance*, 53(2):431–465, 1998.
- Edward J. Green. Embedding an analytic equivalence relation in the transitive closure of a Borel relation. Draft, 2012. URL <http://arxiv.org/abs/1204.1451v1>.
- L. A. Harrington, A. S. Kechris, and A. Louveau. A Glimm-Effros dichotomy for Borel equivalence relations. *Journal of the American Mathematical Society*, 3(4):903–928, 1990. URL <http://www.ams.org/journals/jams/1990-03-04/S0894-0347-1990-1057041-5/home.html>.
- Thomas Jech. *Set Theory*. Springer-Verlag, third edition, 2002.
- A. N. Kolmogorov. *Foundations of Probability Theory*. Chelsea, second English edition, 1956.
- Saul A. Kripke. A completeness theorem in modal logic. *Journal of Symbolic Logic*, 24(1):1–14, 1959.
- David Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
- Jean-François Mertens and Shmuel Zamir. Formulation of Bayesian games with incomplete information. *International Journal of Game Theory*, 14(1):1–29, 1985. URL http://www.ma.huji.ac.il/~zamir/papers/22_IJGT85.pdf.

- Paul Milgrom and Nancy Stokey. Trade, information, and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.
- Dov Monderer and Dov Samet. Approximating common knowledge with common belief. *Games and Economic Behavior*, 1(2):170–190, 1989. URL <http://www.tau.ac.il/~samet/>.
- Yannis N. Moschovakis. *Descriptive Set Theory*. American Mathematical Association, second edition, 2009. URL <http://www.math.ucla.edu/~ynm/books.htm>.
- John C. Oxtoby. *Measure and Category*. Springer-Verlag, second edition, 1980.
- Maxwell B. Stinchcombe and Halbert White. Some measurability results for extrema of random functions over random sets. *Review of Economic Studies*, 59(3):495–514, 1992.
- Jean Tirole. On the possibility of trade under rational expectations. *Econometrica*, 50(5):11, 1982.

DEPARTMENT OF ECONOMICS, THE PENNSYLVANIA STATE UNIVERSITY, UNIVERSITY PARK, PA 16802, USA

E-mail address: eug2@psu.edu