Events Concerning Knowledge

Edward J. Green*
Department of Economics
The Pennsylvania State University
University Park, PA 16802, USA
eug2@psu.edu

DRAFT: 2010.04.05

Abstract

Aumann's knowledge operator can map a measurable event to one that is not measurable. Thus his theorem regarding agreement of subjective probabilities does not apply to uncountable information partitions. The event that an agent knows a measurable event is measurable, if the agent's information partition is induced by a measurable function from the state of the world to the agent's type. Nevertheless, "usual" set theory (Zermelo-Fraenkel axioms and axiom of choice) can be extended in two different ways, one of which imples the generalization to uncountable partitions of Aumann's "agreement theorem" and the other of which is conjectured to imply its negation.

JEL classification: D82, D83, D84

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

1 Introduction

The research to be reported here generalizes Robert Aumann's theory of *common knowledge* (1976; 1999a; 1999b). In his 1976 paper, Aumann restricted his attention to environments in which each agent has a countable information partition.[1] That assumption is restrictive in

*I am grateful to Marcin Peski for a discussion that led toward the positive results in the paper. Support for this research has been provided by a gift from the Human Capital Foundation (`http://www.hcfoundation.ru/en/`) to the Center for the Study of Auctions, Procurements, and Competition Policy at Penn State University.

[1]This is a partition of the possible states of the world. Aumann defines an agent to know an event in a state of the world, if the event is a superset of the block in the agent's information partition that contains

some applications, such as in applications to financial economics where it is desired to represent an agent's information by the value of a nonatomically distributed random variable. The range of such a random variable must be uncountable, so Aumann's "agreement theorem"—the corollary of which is the important theorem in financial economics that traders who share identical prior beliefs cannot have common knowledge that there are mutual gains to purely speculative trade among them—cannot be applied. The 1999 papers do not successfully fix the problem.[2] It will be shown here that, for the class of environments in which information partitions are uncountable, the agreement theorem can be proved by assuming an additional set-theoretic axiom to the axioms typically invoked in economics (that is, to the Zermelo-Fraenkel axioms along with the axiom of choice). On the other hand, it is conjectured that the generalization of Aumann's result is inconsistent with an alternate extension of set theory.

For Aumann, 'knowledge' is synonymous with 'information', and both words mean *direct observation.* Beginning with research by Brandenburger and Dekel (1987), a largely parallel body of game theory has arisen, in which *certainty*—belief with subjective probability 1—replaces knowledge as an object of investigation and as a concept invoked in the statements of substantive results. Especially in view of the issue regarding the set-theoretic foundation of common knowledge that I have just described, a legitimate question is: why not dispense with knowledge and with common knowledge, and instead focus exclusively on certainty and on common certainty? The theorem regarding impossibility of mutually beneficial, purely speculative trade can be reformulated in that way. However, except for strategic-form games in which players share common prior beliefs, the analogy between common knowledge and common certainty is not exact.[3] Moreover, regardless of how close a parallel might exist between the *results* regarding common knowledge and common belief in a game or economic model, in the *formulation* of a game, it is important to maintain the conceptual distinction between what what agents observe and what they believe. There are situations in which the distinction between observing an event to have occurred and merely being subjectively certain that it has occurred is crucial to the specification of agents' payoffs. For example, in Anglo-American judicial systems, a criminal defendent can be convicted on the evidence of someone's testimony that she has observed him committing the crime with which he is charged, but her testimony that she is subjectively certain that the defendant committed the crime would be inadmissable if it were not based on observation. That is, what the person who is testifying believes is irrelevant to, but what she has observed is decisive for, whether the defendant will be acquitted or convicted. This example suggests that the issues considered here will arise in the formalization of such situations as decision problems or games, even in cases where optimal actions and equilibrium strategy profiles can be defined in terms of belief rather than of knowledge.

that state of the world. An event is common knowledge among a group of agent if each agent knows the event, each knows that everyone else knows it, each knows that everyone else knows that everyone else knows it, and so forth.

[2]The formal theory specified here is compared to Aumann (1999a,b) in section 9.

[3]For example, Ben-Porath (1997) shows that common certainty of players' rationality cannot substitute for common knowledge of their rationality in justifying the backward-induction solution of a game.

2 Events regarding knowledge and certainty

Consider an agent whose *prior beliefs* are specified by a probability space $(\Omega, \mathcal{B}, P)$, where $\mathcal{B}$ is the Borel σ-algebra on Ω and $P : \mathcal{B} \to [0,1]$ is a countably additive probability measure. It is assumed that[4]

$$\begin{aligned} &\Omega \text{ is a complete, separable metric space.} \\ &\mathcal{B} \text{ is the } \sigma\text{-algebra of Borel subsets of } \Omega. \end{aligned} \tag{1}$$

The agent's *information partition* Π satisfies

$$\Pi \subseteq \mathcal{B} \setminus \{\emptyset\} \text{ and } \forall \pi \in \Pi \; \forall \pi' \in \Pi \; [\pi = \pi' \text{ or } \pi \cap \pi' = \emptyset] \text{ and } \bigcup \Pi = \Omega. \tag{2}$$

That is, Π partitions Ω into Borel-measurable subsets.[5] For each $\omega \in \Omega$, the agent's *information set at* ω is the unique element $\pi \in \Pi$ such that $\omega \in \pi$.

Define the σ-algebra $B_\Pi \subseteq \mathcal{B}$ by

$$\mathcal{B}_\Pi = \{X \mid X \in \mathcal{B} \text{ and } \forall \omega \in X \;\; \exists \pi \in \Pi \;\; [\omega \in \pi \subseteq X]\}. \tag{3}$$

$\mathcal{B}_\Pi$ is the σ-algebra of Borel-measurable sets, regarding the occurrence or non-occurrence of which the agent is informed. Note that $\Pi \subseteq \mathcal{B}_\Pi$.

The agent's *posterior beliefs* are specified by a regular conditional probability function $p : \mathcal{B} \times \Omega \to [0,1]$, that is, by a function that satisfies the following conditions.

$$\text{For every } \omega,\; p_\omega(X) = p(X, \omega) \text{ defines a Borel probability measure.} \tag{4}$$

$$\text{For every } X \in \mathcal{B},\; p_X(\omega) = p(X, \omega) \text{ is } B_\Pi\text{-measurable.} \tag{5}$$

$$p_\omega(\pi(\omega)) = 1 \text{ a.s.} \tag{6}$$

$$\text{For every } X \in \mathcal{B} \text{ and } Y \in \mathcal{B}_\Pi,\; \int_Y p_X(\omega)\, dP = P(X \cap Y). \tag{7}$$

Consider an event $X \in \mathcal{B}$. That the agent knows X is also an event, as is the event that the agent is certain of X. Let $K(X)$ and $C(X)$ denote the knowledge and certainty events respectively. Here, and henceforth in this paper, *event* means simply a subset of Ω. It does not necessarily mean (as in general mathematical usage) that the subset is a Borel set or an element of some other σ-algebra. Indeed, to find a sufficient condition for measurability of $K(X)$ is the major theoretical problem to be addressed here.

Following Aumann (1976), the agent is specified to know an event X when he is in state ω if his information set in ω is a subset of X. That is, all of the states of the world that his information does not exclude from being the true state are in X. Formally,

$$K(X) = \{\omega \mid \exists \pi \in \Pi \; [\omega \in \pi \subseteq X]\}. \tag{8}$$

[4]More generally, it can be assumed that $(\Omega, \mathcal{B})$ isomorphic to a Borel subset of a *Polish space*, that is, of a topological space that could be characterized in terms of a metric satisfying this assumption.

[5]A *Borel-measurable set* is simply an element of $\mathcal{B}$. A *P-measurable set* with respect to a probability measure is a broader concept that will be defined (with respect to a probability measure P) below.

As discussed above, the agent is specified to be certain of X in ω if $p_\omega(X) = 1$. That is,

$$C(X) = \{\omega \mid p_\omega(X) = 1\}. \tag{9}$$

Definitions (8) and (9), along with properties (5)–(7) of the definition of regular conditional probability, imply the following lemma.

Lemma 1 *For every $X \in \mathcal{B}$, $K(X) \subseteq X$ and $C(X) \in \mathcal{B}_\Pi$. For all $X \in \mathcal{B}_\Pi$, $P(X \triangle C(X)) = 0$.*[6]

3 A measurability problem with $K(X)$

Knowledge of a Borel-measurable event is not necessarily a measurable event, even in the weaker (than Borel) sense of being measurable with respect to the completion of some Borel measure. An example of this problem will be constructed momentarily. To prepare for this example, recall what is the *completion* of a measure.

Let 2^Ω denote the power set of Ω, and let $P : \Omega \to [0,1]$ be a countably additive measure. Then define $P_* : 2^\Omega \to [0,1]$ by $P_*(X) = \sup\{P(Y) \mid Y \in \mathcal{B} \text{ and } Y \subseteq X\}$ and define $P^* : 2^\Omega \to [0,1]$ by $P^*(X) = \inf\{Z \mid Z \in \mathcal{B} \text{ and } X \subseteq Z\}$. The set $\mathcal{B}_P = \{X \mid P_*(X) = P^*(X)\}$ is a σ-algebra, the events in which will be called *P measurable*. The restriction of P^* to $\mathcal{B}_P$ is a countably additive measure that coincides with P on $\mathcal{B}$. To simplify notation, P will be considered to be defined on $\mathcal{B}_P$. The measure on this enlarged σ-algebra is called the *completion* of P.[7]

An *atom* of $\mathcal{B}$ is a smallest non-empty element of $\mathcal{B}$ (that is, a non-empty element of $\mathcal{B}$, no nonempty, proper subset of which is also an element of $\mathcal{B}$). A probability measure P *is purely atomic* if there is a countable set X of atoms of $\mathcal{B}$, such $P(\Omega \setminus \bigcup X) = 0$.

Lemma 2 *There is an event V such that*

$$\forall P\ [P \textit{ is not purely atomic } \Longrightarrow V \textit{ is not } P \textit{ measurable }].^8 \tag{10}$$

Example 1 *An information partition Π and a Borel-measurable event X, such that*

$$\forall P\ [P \textit{ is not purely atomic } \Longrightarrow K(X) \textit{ is not } P \textit{ measurable }].$$

---

[6] The symmetric-difference operation on sets is denoted by $\triangle$.

[7] Halmos (1970, Chapters II–III) presents this material in detail.

[8] Oxtoby (1980, Theorems 5.3, 5.4). The proof utilizes the Axiom of Choice.

Let $\Omega = [0,1]$, and let V be an event satisfying (10). Let $X = \{\omega/2 \mid \omega \in V \setminus \{0,1\}\}$. Clearly $X \subseteq (0, 1/2)$ and X satisfies (10), since that property is not affected by deletion of a finite set or by a linear transformation. Define Π by

$$\omega \in \pi \Longleftrightarrow \begin{cases} \pi = \{\omega, \omega + 1/2\} & \text{if } \omega \in X, \\ \pi = \{\omega, \omega - 1/2\} & \text{if } \omega - 1/2 \in X, \\ \pi = \{\omega\} & \text{otherwise.} \end{cases} \tag{11}$$

Now it will be shown that, although $[0, 1/2]$ is Borel measurable, $K([0, 1/2])$ is not P measurable for any P that is not purely atomic. $K([0, 1/2]) = [0, 1/2] \setminus X$, and therefore $[0, 1/2] \setminus K([0, 1/2]) = X$. Since the difference of two P-measurable sets is also P measurable, and since X is not P measurable if P is not purely atomic, $K([0, 1/2])$ cannot be P measurable if P is not purely atomic. ∎

Example 1 shows that, if Π satisfies (2) and $X \in \mathcal{B}$, nevertheless it is possible that $K(X)$ defined by (8) may not be measurable with respect to any prior belief that is not purely atomic. To embed the example in the full model of an agent specified in section 2, satisfying all of the conditions (1)–(7), it is sufficient to specify a non-purely-atomic prior belief P. Parthasarathy (1967, Theorem 8.1) shows that there is a regular conditional probability measure satisfying (4)–(7) that represents the agent's posterior belief.[9]

4 Measurable assignment of knowledge states induces measurable knowledge events

The measurability problem identified in example 1 does not occur (for any prior probability measure), if the agent's information partition is induced by a Borel-measurable function $\tau : \Omega \to T$, where each T is a separable, complete metric space of knowledge states (or *types*).[10]

Assume that

$$T \text{ is a separable, complete metric space, and } \mathcal{T} \text{ is the Borel } \sigma\text{-algebra on } T. \tag{12}$$

Let $\mathcal{A}$ denote the product σ-algebra $\mathcal{B} \times \mathcal{T}$ on $\Omega \times T$. Define a subset X of either Ω or T to be *analytic* if it is the projection of a set $Y \in \mathcal{A}$. Define a subset X of either Ω or T to be *universally measurable* if, for every measure μ, $\mu_*(X) = \mu^*(X)$.

Lemma 3 *A countable union of analytic sets is analytic. Every analytic set is universally measurable. The complement of a universally measurable set is universally measurable.*[11]

[9]To apply Parthasarathy's theorem, let X and Y both be Ω, let $\mathcal{C} = \mathcal{B}_\Pi$ and let π be the identity mapping.

[10]Π is induced by τ if $\Pi = \{\tau^{-1}(t) \mid t \in T\} \setminus \{\emptyset\}$.

[11]The first assertion is proved in Bertsekas and Shreve (1978, Corollary 7.35.2). The second is proved in Bertsekas and Shreve (1978, Corollary 7.42.1). The third is immediate from the definition of measurability with respect to μ.

Lemma 4 *If $\tau : \Omega \to T$ is Borel measurable, and $X \subseteq \Omega$ and $Y \subseteq T$ are analytic, then $\tau(X)$ and $\tau^{-1}(Y)$ are analytic.*[12]

Proposition 1 *If Π is the information partition induced by a Borel-measurable function $\tau : \Omega \to T$, where T is a separable, complete metric space, then for every $X \in \mathcal{B}$, $K(X)$ is universally measurable.*

Proof $K(X) = \{\omega \mid \tau(\omega) \notin \tau(\Omega \backslash X)\} = \Omega \backslash \tau^{-1}(\tau(\Omega \backslash X))$. Event $\tau^{-1}(\tau(\Omega \backslash X))$ is analytic by lemma 4, so it is univerally measurable by lemma 3. Therefore $K(X) = \Omega \backslash \tau^{-1}(\tau(\Omega \backslash X))$ is universally measurable, again by lemma 3. ∎

5 Common knowledge

Aumann (1976) formalizes the theory of common knowledge and proves an "agreement theorem" that, if agents' posterior probabilities of an event in some some state of the world are commonly known to them, then the posterior probabilities are identical if the agents share a common prior. Aumann makes an assumption that implies that the information partition of each agent is a countable set of measurable events, each of which has positive prior probility. In this section, it will be shown that a measurability problem related to example 1 can occur otherwise.

Consider a finite set I of agents, and let each agent i have information partition Π_i. Let Π_I be the *common-knowledge partition*, which is the finest partition that is at least as coarse as each Π_i.[13] Define $K_I(X) = \bigcap_{i \in I} K_i(X)$. The event that X is common knowledge is $\bigcap_{n \in \mathbb{N}} K_I^n(X)$. Denote this event by $K_I^\infty(X)$.

Analogously to (8) (which defines $K(X)$ for an individual agent), the event that event X is commonly known is defined by

$$\omega \in K_I^\infty(X) \iff \exists \pi \in \Pi_I\ [\omega \in \pi \subseteq X]. \tag{13}$$

The following example, which builds on example 1, shows that an element of Π_I may fail to be measurable with respect to a prior belief that is not purely atomic.

Example 2 *A non-measurable event in a common-knowledge partition*

Consider two agents. Let Π_1 be the information partition constructed in example 1. (Recall that the construction is based on an event $X \subseteq (0, 1/2)$ that is not measurable with respect to any measure that is not purely atomic.) Let $\Pi_2 = \{[1/2, 1]\} \cup \{\{\omega\} \mid \omega < 1/2\}$. Both agents' information partitions consist solely of events in $\mathcal{B}$. However, $\Pi_I = \{[1/2, 1] \cup X\} \cup \{\{\omega\} \mid \omega < 1/2 \text{ and } \omega \notin X\}$. Since $[1/2, 1] \cup X$ is the disjoint union of a Borel set and X, it is not measurable with respect to any measure that is not purely atomic. (Otherwise X would also be measurable, contrary to assumption.) ∎

[12] Bertsekas and Shreve (1978, Proposition 7.40).

[13] Π is as fine as Π' if $\forall \pi \in \Pi\ \exists \pi' \in \Pi'\ \pi \subseteq \pi'$. It is strictly finer if, furthermore, the two partitions are not identical. Π' is as coarse as Π if Π is not strictly finer than Π'.

6 The common-knowledge partition under measurable assignment of knowledge states

As in the case of events defined by an individual agent's knowledge, the non-measurability problem illustrated by example 2 can be ruled out by assuming that there is a measurable assignment of knowledge states. Specifically, if each agent's information partition is induced by a Borel-measurable type function $\tau_i : \Omega \to T_i$ (where each T_i is a separable complete metric space), then the common-knowledge partition consists of analytic events. The key to proving this fact is stated in terms of a function $\pi_I : 2^\Omega \to 2^\Omega$ defined by

$$\pi_I(X) = \bigcup_{i \in I} [\Omega \setminus K_i(\Omega \setminus X)] = \bigcup_{i \in I} \tau_i^{-1}(\tau_i(X)). \tag{14}$$

Lemma 5 *The element π of Π_I such that $\omega \in \pi$ satisfies*

$$\pi = \bigcup_{n \in \mathbb{N}} \pi_I^n(\{\omega\}).^{14} \tag{15}$$

Proposition 2 *Every element of Π_I is analytic, and is thus universally measurable.*

Proof For an arbitrary event X, $\pi_I(X) = \bigcup_{i \in I} \tau_i^{-1}(\tau_i(X))$. Therefore, by lemma 3, $\pi_I(X)$ is analytic if X is analytic. Suppose that $\omega \in \pi \in \Pi_I$. By induction, every set $\pi_I^n(\{\omega\})$ is analytic. Therefore, by lemma 3 and lemma 5, π is analytic, and thus universally measurable. ∎

7 Generalizing Aumann's agreement theorem

Aumann's (1976) theorem concerns the pairwise comparison of agents' posterior probabilities of an event X, in states of the world in which their posterior probabilities of that event are common knowledge. Given this topic, there is no loss of generality from making the simplifying assumption that there are only two agents. Following Aumann, assume in this section that $I = \{1, 2\}$ and that P is the agents' common prior probability measure, and let $p_i(\cdot, \omega)$ be the regular conditional probability of agent i (conditioning on $\mathcal{B}_{\Pi_i}$). Fix an event X, and define

$$L = \{\omega \mid p_1(X,\omega) < p_2(X,\omega)\} \text{ and } G = \{\omega \mid p_1(X,\omega) > p_2(X,\omega)\} \text{ and } E = \{\omega \mid p_1(X,\omega) = p_2(X,\omega)\} \tag{16}$$

It will be convenient to reformulate Aumann's result by studying the event that the two agents' posterior probabilities of X are commonly known *not* to be equal. Aumann proves

[14] This equation formalizes the conclusion of the paragraph analyzing *reachable from* ω, in Aumann (1976, p. 1237).

that this event is empty. That strong result is a consequence of his assumption that every non-empty, measurable event has positive measure. Without that assumption, the conclusion would be that the commonly known non-agreement event is null (that is, that it has prior probability zero). Thus, Aumann's theorem is tantamount to the following proposition and its corollary.

Proposition 3 *If each agent's information partition is induced by a Borel-measurable type function $\tau_i : \Omega \to T_i$ (where each T_i is a separable complete metric space),*

$$\Pi_I \text{ is countable,} \tag{17}$$

and M is an event that satisfies

$$M = \bigcup \{\pi \mid \pi \in \Pi_I \text{ and } \pi \subseteq M\}, \tag{18}$$

then $P(M \cap L) = P(M \cap G) = 0$ and $P(M \cap E) = P(M)$.

Proof By proposition 2, every element of Π_I is P-measurable. By (17), every union of elements of Π_I is P-measurable. In particular, by (18), M, $M \cap L$, $M \cap G$, and $M \cap E$ are all P-measurable. The last three sets partition M.

For each of these sets $M \cap K$ ($K \in \{L, G, E\}$), the condition (7) of the definition of regular conditional probability states that

$$\int_{M \cap K} p_1(X, \omega)\, dP(\omega) = P(M \cap K \cap X) = \int_{M \cap K} p_2(X, \omega)\, dP(\omega). \tag{19}$$

Since $p_1(X, \omega) < p_2(X, \omega)$, for all $\omega \in L$, $\int_{M \cap L} p_1(X, \omega)\, dP(\omega) < \int_{M \cap L} p_2(X, \omega)\, dP(\omega)$ if $P(M \cap L) > 0$. Thus, by (19), $P(M \cap L) = 0$. An analogous argument shows that $P(M \cap G) = 0$, so $P(M \cap E) = P(M)$. ■

The following proposition is virtually a corollary of proposition 3, but it is not quite a special case of that proposition because hypothesis (17) is not invoked.

Proposition 4 *If each agent's information partition is induced by a Borel-measurable type function $\tau_i : \Omega \to T_i$ (where each T_i is a separable complete metric space), $Z \in \mathcal{B}$, $\omega \in \pi \in \Pi_I$, $P(\pi) > 0$, and $\pi \subseteq p_{1Z}^{-1}(p_1^*) \cap p_{2Z}^{-1}(p_2^*)$ (where p_{iZ} denotes i's posterior probability of Z as represented in clause (5) of the definition of regular conditional probability), then $p_1^* = p_2^*$.*

Proof In proposition 3, take $X = p_{1Z}^{-1}(p_1^*) \cap p_{2Z}^{-1}(p_2^*)$ and $M = \pi$. Note that, because M is the union of a countable set (specifically, of a singleton) of elements of Π_I, the countability of Π_I does not have to be assumed. ■

Remarks:

1. Proposition 4 corresponds most closely to Aumann's statement of his theorem. It is included—with its extremely short proof—to substantiate that proposition 3 expresses the main content of the result and is the appropriate version of it to be generalized. Proposition 4 itself is a poor candidate for generalization because typically, in a setting where P is nonatomic, no element of Π_I will have positive measure.

2. Aumann (1976) invokes an even stronger countability assumption than condition (17). His condition implies that the functions τ_i satisfying the first hypothesis of the proposition can be constructed, although that hypothesis is not explicit in his theorem.

3. For the no-mutual-gains-to-speculative-trade theorem, the result about L and M rather than the result about E is directly relevant. That is, for it to be commonly known that there is a mutual gain to speculative trade, the direction of the inequality between agents' posterior beliefs—not just the fact that they are not equal—must be commonly known. In order to be commonly known, such an event must be null.

4. In proposition 4 and in proposition 5 below, condition (17) is replaced by alternate conditions that imply that M is P-measurable. Because the propositions are identical to proposition 3 in other respects, the proof of proposition 3 establishes those propositions as well, once this measurability condition has been proved.

8 A set-theoretic condition for $K_I^\infty(X)$ to be measurable

A full-strength generalization of Aumann's agreement theorem would be simply to remove the countability hypothesis of the statement that $K_I^\infty(X)$ is P-measurable. The usual axioms of set theory (ZFC)—the Zermelo-Fraenkel axioms (ZF) together with the Axiom of Choice—are apparently not strong enough to prove or disprove such a generalization. It can be proved, however, from an additional axiom, the *axiom of projective determinacy* (PD). On the other hand, it seems likely to be inconsistent with another axiom, the *axiom of constructability* (V=L) It has been proved in ZF that the axiom of constructability implies the axiom of choice. The prevailing conjecture among set theorists is that PD is also logically consistent with ZFC.[15]

8.1 Projective sets

The results to be derived below are consequences of the fact—to be proved momentarily—that $K_I^\infty(X)$ is a countable intersection of *projective sets.*

A subset Z of a Polish space Ω is projective if either

1. Z is analytic; or

2. $\Omega \setminus Z$ is projective; or

3. there is a projective subset $V \subseteq \Omega \times \Omega'$ (where Ω' is some Polish space) such that $Z = \{v \in \Omega \mid \exists w \in W \; (v, w) \in V\}$.

That is, a set is projective if it is an analytic subset of some Polish space, or if it can be obtained from analytic sets via finitely many operations of complementation and projection.

[15]Jech (2002) is among the textbooks that state these axioms and examine them in detail.

Lemma 6 *If each agent's information partition is induced by a Borel-measurable type function $\tau_i : \Omega \to T_i$ (where each T_i is a separable complete metric space), then for every Borel-measurable set X and every number n, $K_I^n(X)$ is a projective set. $K_I^\infty(X)$ is a countable intersection of projective sets.*

Proof That every $K_I^n(X)$ is projective is proved by induction. The basis step is that $K_I^0(X) = X$, and X is a Borel set and thus an analytic set. The induction step closely resembles the proof of proposition 1. (Note that the graph of a Borel-measurable function is an analytic set, so the images and inverse images of projective sets under τ_i are projective sets.) $K_I^\infty(X)$ is a countable intersection of projective sets because $K_I^\infty(x) = \bigcap_{n\in\mathbb{N}} K_I^n(X)$. ∎

As has been mentioned above, two axioms have opposite implications regarding whether or not all projective sets are universally measurable. ZF+V=L implies that a non-measurable projective set exists. I conjecture that, assuming V=L, an example can be constructed of a set $X \in \mathcal{B}$ such that $K_I^\infty(X)$ is not universally measurable. ZFC+PD implies that every projective set is universally measurable. In that case, since every $K_I^\infty(X)$ (for $X \in \mathcal{B}$) is a countable intersection of projective sets,ZFC+PD entails that every $K_I^\infty(X)$ is universally measurable.

8.2 Agreement of agents' posterior probabilities

The following generalization of Aumann's theorem is provable in ZFC. Its hypotheses (20) is implied by the axiom of projective determinacy. The proof is parallel to the proof of proposition 3, with lemma 6 and the fact that a countable intersection of P-measurable sets is P-measurable playing the role that the countability hypothesis of proposition 3 plays in the proof of that proposition.

Proposition 5 *If each agent's information partition is induced by a Borel-measurable type function $\tau_i : \Omega \to T_i$ (where each T_i is a separable complete metric space) and*

$$\textit{every projective event is } P\textit{-measurable} \tag{20}$$

then, for the events L, G, and E defined in equation (16), $K_I^\infty(L)$, $K_I^\infty(G)$, and $K_I^\infty(E)$ are all P-measurable, and $P(L) = P(G) = 0$.

9 Comparison between section 2 and Aumann's framework

The framework of section 2 closely follows a single-agent specialization of Aumann (1999a,b). Now several differences in exposition, and minor differences in specification, are discussed.

In (1999a), Aumann defines an information partition in terms of a mapping from the states of the world to an abstract set of *knowledge states*. The information partition is induced by the equivalence relation of mapping to the same knowledge state. The framework of this paper is equivalent in this regard, since information sets can be taken to be the knowledge states. When Π is viewed in a natural way as a correspondence from each state of the world to its equivalence class in the information partition, it functions as Aumann's mapping from states of the world to knowledge states.

In (1999b, section 12), Aumann defines a *knowledge-belief system.* While the mapping to knowledge states in (1999a) is not assumed to possess any measurability property, the information sets of a knowledge-belief system are assumed to be measurable events as in (2). However, $(\Omega, \mathcal{B})$ is assumed only to be an abstract measurable space, rather than to be derived from a separable, complete metric space.

A knowledge-belief system includes a function π that satisfies conditions (4) that p_ω is a measure for every ω and (5) that p_X is measurable for every X. Aumann assumes that $K(X) \subseteq C(X)$, which is slightly stronger than assumption (6) here.[16] Specifically, (6) and the definitions of K and C imply that $P(K(X) \setminus C(X)) = 0$, while Aumann's assumption can be reformulated as $K(X) \setminus C(X) = \emptyset$. Aumann also assumes that, for every $X \in \mathcal{B}$ and $\alpha \in [0, 1]$, $K(p_X^{-1}(\alpha)) = p_X^{-1}(\alpha)$. This assumption is implied by assumption (5) that p_X is $\mathcal{B}_\Pi$ measurable (which entails that $p_X^{-1}(\alpha) \in \mathcal{B}_\Pi$ because $\{\alpha\}$ is the intersection of countably many open sets) and the definition (8) of K (which entails that $\forall X \in \mathcal{B}_\Pi \; K(X) = X$).

Aumann's knowledge-belief system does not have an analogue of (7), which is the condition of Bayesian consistency between prior and posterior beliefs. Rather, the system specifies only posterior beliefs. If desired, prior beliefs of a single agent could be specified by taking (7) as a definition. However, that approach would have limitations in a multi-agent setting. For example, if it is desired to impose a common-prior assumption (cf. Harsanyi (1967–68)), then it is more straightforward to do so directly than to verify that the implications of the assumption regarding agents' posterior probabilities are satisfied.

[16] Note that example 1 and example 2 are counterexamples to measurability within Aumann's framework as well as within the framework of this paper, since the examples satisfy this assumption.

References

Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.

Robert J. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28:263–300, 1999a.

Robert J. Aumann. Interactive epistemology II: Probability. *International Journal of Game Theory*, 28:301–314, 1999b.

Elchanan Ben-Porath. Rationality, Nash equilibrium, and backwards induction in perfect-information games. *Review of Economic Studies*, 64:23–46, 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`.

Adam Brandenburger and Eddie Dekel. Common knowledge with probability 1. *Journal of Mathematical Economics*, 16:237–245, 1987.

Paul R. Halmos. *Measure Theory*. Springer-Verlag, 1970.

John C. Harsanyi. Games with incomplete information played by Bayesian players. *Management Science*, 14:159–182, 320–334, 486–502, 1967–68.

Thomas Jech. *Set Theory*. Springer-Verlag, third edition, 2002.

John C. Oxtoby. *Measure and Category*. Springer-Verlag, second edition, 1980.

K. R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, 1967.