

Agreeing to agree

EHUD LEHRER

School of Mathematical Sciences, Tel Aviv University

DOV SAMET

Faculty of Management, Tel Aviv University

Aumann (1976) shows that agents who have a common prior cannot have common knowledge of their posteriors for event E if these posteriors do not coincide. But given an event E , can the agents have posteriors with a common prior such that it is common knowledge that the posteriors for E do coincide? We show that a necessary and sufficient condition for this is the existence of a nonempty *finite* event F with the following two properties. First, it is common knowledge at F that the agents cannot tell whether E occurred. Second, this still holds true at F , when F itself becomes common knowledge.

KEYWORDS. Agreement theorem, common knowledge, common prior, no trade theorem.

JEL CLASSIFICATION. C70, D82.

1. INTRODUCTION

1.1 *Agreeing to disagree and agreeing to agree*

Can agents have common knowledge of their beliefs? In a seminal paper Aumann (1976) demonstrates the impossibility of *agreeing to disagree*: For *any* posteriors with a common prior, if the agents' posteriors for an event E are different (= *they disagree*), then the agents cannot have common knowledge (= *agreeing*) of these posteriors. Thus, the short answer to our opening question is that agents cannot have common knowledge of their beliefs when they are different.

But can there be other reasons why agents may fail to have common knowledge of their beliefs, even when these beliefs are the same? Obviously, agents may fail incidentally to know each others posteriors, and thus, a fortiori they cannot have common knowledge of them, even if they coincide. Our question here is not concerned with such an incidental lack of common knowledge of the posteriors, but rather with the impossibility of having such common knowledge in *principle*. That is, we examine the impossibility of having such common knowledge regardless of the beliefs held by the agents.

We investigate the conditions under which there *exist* posteriors with a common prior, such that it is common knowledge that the agents' posteriors for an event E

Ehud Lehrer: ehudlehrer@gmail.com

Dov Samet: dovsamet@gmail.com

Copyright © 2011 Ehud Lehrer and Dov Samet. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](http://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: 10.3982/TE578

coincide. Using the terminology of Aumann's theorem, we ask, "When is *agreeing to agree* about the posteriors of E possible?" We clarify this question by examining several examples.

EXAMPLE 1. Each of two firms can be either profitable or losing. Thus, there are four possible states of the world (profitable, profitable), (profitable, losing), (losing, profitable), and (losing, losing). Each firm knows only how well it does.

Consider first the event E that both firms have the same financial situation. That is, $E = \{(\text{profitable, profitable}), (\text{losing, losing})\}$. There are many posteriors that have a common prior such that the firms have common knowledge that the posterior probabilities of E coincide. For instance, consider a common prior, μ , with

$$\mu(\text{profitable, profitable}) = \mu(\text{losing, losing}) = p$$

and

$$\mu(\text{profitable, losing}) = \mu(\text{losing, profitable}) = q,$$

where $2p + 2q = 1$. Then, *in each of the four states*, both firms have the same posterior for E : $p/(p + q)$. Thus, it is common knowledge in every state that the posteriors for E are the same. \diamond

EXAMPLE 2. In the state space of the previous example, consider the event $E = \{(\text{losing, profitable})\}$. Note that the only event about which there is common knowledge is the whole state space. Therefore, if for some posteriors it is common knowledge that both posteriors for E are p , then the posteriors for E must be p in all four states. However, for any posterior on this space, when firm 1 is profitable, its posterior for E is 0. Thus, the only value p can have is 0.

Consider the complement of E , denoted $\neg E$. When firm 2 is losing, its posterior for $\neg E$ is 1. Thus, by the same argument as above, the only p for which there can be common knowledge that both posteriors for $\neg E$ are p , is 1.

As opposed to the event E in Example 1, in this example the possibility of agreeing to agree about the posteriors of E , or $\neg E$, is limited to trivial ones, i.e., 0 or 1. \diamond

1.2 A necessary condition for the possibility of agreeing to agree

The difference between the events in Examples 1 and 2 suggests a simple necessary condition for agreeing to agree. In any state at which the firm's posterior for E is nontrivial, the firm cannot tell whether E is the case or not E is the case. We say in this case that the firm is *ignorant* of E . Therefore, we can state the following condition.

If there exist posteriors with a common prior for which it is common knowledge that both posteriors for E are p for some $0 < p < 1$, then it is necessarily common knowledge that both firms are ignorant of E .

Now, in the first example, the firms are ignorant of E in each state. Thus, it is common knowledge that they are ignorant of E and the necessary condition for agreeing to

agree is satisfied. In the second example, in contrast, in all states there is no common knowledge that the firms are ignorant of E . Therefore, no matter what nontrivial posteriors for E the agents may have, the event where their posteriors coincide cannot be common knowledge.

Surprisingly, this necessary condition is not sufficient, as demonstrated by the next example.

EXAMPLE 3. The profit of each of two firms can be any integer, positive or negative. Each firm is informed only of its profit. Let E be the event that firm 2's profit is higher. That is, E is the event $(i < j)$. In each state, no firm can tell whether 2's profit is higher. Therefore, it is common knowledge in each state that both firms are ignorant of E . Nevertheless, no matter what posteriors with a common prior the firms may have, there can be no common knowledge that both posteriors for E are p for some $0 < p < 1$.

To see this, assume to the contrary that μ is a common prior and it is common knowledge that both posteriors of E are p with $0 < p < 1$. As the only event that is common knowledge is the whole space, it follows that the posteriors of E are p at each state of the world. Fix k and consider the right half-space $R_k = (i \geq k)$, the upper half-space $U_k = (j > k)$, the quadrants $Q_k^1 = Q^1 = U_k \cap \neg R_k$ and $Q_k^2 = Q^2 = R_k \cap \neg U_k$, and the cones $C_k^1 = C^1 = E \cap R_k$ and $C_k^2 = C^2 = \neg E \cap U_k$. It is easy to see that $E \cap U_k = C^1 \cup Q^1$ and $\neg E \cap R_k = C^2 \cup Q^2$, and the unions are disjoint. By our assumption, $\mu(C^1 \cup Q^1)/\mu(C^2) = \mu(E \cap U_k)/\mu(\neg E \cap U_k) = p/(1 - p)$ and $\mu(C^2 \cup Q^2)/\mu(C^1) = \mu(\neg E \cap R_k)/\mu(E \cap R_k) = (1 - p)/p$. It follows from these equations that $\mu(Q^1) = [p/(1 - p)]\mu(C^2) - \mu(C^1)$ and $\mu(Q^2) = [(1 - p)/p]\mu(C^1) - \mu(C^2)$. Thus, $\mu(Q^1) = -[p/(1 - p)]\mu(Q^2)$. This implies that $\mu(Q^1) = 0$. As this holds for each k and $E = \bigcup_k Q_k^1$, it follows that $\mu(E) = 0$, which contradicts our assumption. \diamond

1.3 Necessary and sufficient conditions

Our first result provides a necessary and sufficient condition for the possibility of agreeing to agree in terms of the information structure only. Although the notion of agreeing to agree involves posteriors and common priors, our characterization does not employ any probabilistic notions.

Agreeing to agree on nontrivial posteriors for E is possible if and only if there is a nonempty finite event F at which it is common knowledge that the agents are ignorant of E , and this holds true at F also after F becomes common knowledge.

Since agreeing to agree is impossible in [Example 3](#), such a nonempty finite event F cannot exist in this example. Indeed, suppose that there exists such an event F . Let i_m be the maximal profit of 1 in all the states of F . By the property of F , 1 is still ignorant of E after being informed of F . In particular, 1 cannot tell that $\neg E$ is the case. Therefore, there must exist a point (i_m, j_0) in F that is also in E . By the definition of E , $i_m < j_0$. By the definition of i_m , for each point $(i, j_0) \in F$, $i \leq i_m$ and hence $i < j_0$. Thus, $(i, j_0) \in E$. But this means that when firm 2's profit is j_0 and it is informed of F , it knows E . This contradicts the assumption that it is ignorant of E .

Another characterization of the possibility of agreeing to agree is given in probabilistic terms.

Agreeing to agree about nontrivial posteriors for E is possible if and only if it is possible with a common prior that has a finite support.

While firms cannot agree on the probability that one is more profitable than the other, as shown in the previous example, firms can still agree on the probability that they are equally profitable. To see this, consider the following example.

EXAMPLE 4. In the same state space as in [Example 3](#), define E to be the event where both firms are equally profitable, that is, $E = (i = j)$. The conditions for the possibility of agreeing to agree hold in this case. Indeed, consider the finite set $F = \{(1, 1), (1, 0), (0, 1), (0, 0)\}$. It is common knowledge at every state, and in particular at every state of F , that the agents are ignorant of E . Second, when this finite set becomes common knowledge, the firms are still ignorant of E ; moreover, their ignorance is also common knowledge.

Using the set F , it is easy to construct posteriors with a common prior for which agreeing to agree holds for E . Consider the common prior μ , where $\mu(1, 1) = \mu(1, 0) = \mu(0, 1) = \mu(0, 0) = \frac{1}{4}$. When a firm's profit is either 0 or 1, the posterior for E is $\frac{1}{2}$. Since any information set of a firm whose profit is neither 0 nor 1 has probability zero according to μ , one may define the posteriors in these elements in an arbitrary fashion. In particular, one may define these posteriors for E to be $\frac{1}{2}$. Defining the posteriors in this way makes the fact that the posteriors for E are $\frac{1}{2}$ common knowledge at each state. \diamond

1.4 Positive common priors

In the previous example, agreeing to agree is made possible by a degenerate common prior. It vanishes on all but two elements in each partition. Our next result gives a necessary and sufficient condition for the possibility of agreeing to agree with a positive common prior. That is, a common prior that assigns a positive probability to every information set.

For any given event E , we define simple-structured, nonempty, and finite events called *loops on E* (see [Definition 4](#) and [Figure 1](#)). One of the conditions that define a loop implies that when a loop on E becomes common knowledge, it is common knowledge at the loop that the agents are ignorant of E . This property links the notion of loops to the condition in our main result.

When the information set of an agent intersects a loop, we say that this agent does not exclude loops. Our third result is the following statement.

Agreeing to agree is possible with a common prior that is positive on all information sets if and only if it is common knowledge at some state that the agents do not exclude loops.

1.5 The role of finiteness

The less obvious part of our characterizations is the finiteness. It has to do with the way E and $\neg E$ are entangled in the information structure. When there is *no* finite F with the required property, it is possible to separate E and $\neg E$ by arranging the partitions in a way that resembles the structure of the events in [Example 3](#). When there *is* such F ,

the sets E and $\neg E$ are “mixed” together like those in Examples 1 and 4. The separation is described in Proposition 3. In Proposition 4 it is shown why this separation implies the nonexistence of posteriors that makes agreeing to agree possible. This proposition generalizes an argument made in Samet et al. (2004) for the special case of Example 3, which is also related to an intriguing puzzle due to Cover (1987).

1.6 No trade

Suppose the posterior probabilities of E at some state for player 1 and 2 are p and q , respectively, with $p > q$. Consider a bet on event E in which if E obtains, player 1 pays $x > 0$ to player 2, and if $\neg E$ obtains, player 2 pays $y > 0$ to player 1. If x/y is in the interval $((1 - q)/q, (1 - p)/p)$, then the expected payoff of the bet at this state is positive for both players and both are willing to bet.

However, by expressing their will to accept the bet, the players convey information to each other, which results in a refinement of their partitions. Given this new information, the two players reassess the probability of E and their willingness to bet. When the process is repeated, it converges to limit partitions for which the two posteriors of E are common knowledge, and hence, by Aumann’s disagreement theorem, they coincide. Obviously, at this stage there is no trade, namely, one of the parties refuses to bet (Sebenius and Geanakoplos 1983).

Thus, the dynamics that leads to no trade ends with an agreement to agree on the posterior of E . The result of this paper shows that the limit partitions of this process are not arbitrary, and provides the precise relationship between the limit partitions and the event E .

1.7 Agreement theorems

The literature abounds with generalizations of Aumann’s agreement theorem. In this theorem the disagreement concerns the posterior of an event. Milgrom and Stokey (1982), Sebenius and Geanakoplos (1983), and Rubinstein and Wolinsky (1990) study agreement theorems where the disagreement concerns the value or expected value of random variables.

Aumann’s theorem provides a necessary condition for the existence of a common prior. Feinberg (1995), Morris (1994), Bonanno and Nehring (1966), and Samet (1998) show that when disagreements are extended to random variables, then the impossibility of agreeing to disagree is a necessary condition for the existence of a common prior.

McKelvey and Page (1986) and Nielsen et al. (1990) show that when the agents have different posteriors, then not only can these posteriors not be common knowledge, but also certain aggregates of the posteriors cannot be common knowledge.

Geanakoplos and Polemarchakis (1982), Sebenius and Geanakoplos (1983), Parikh and Krasucki (1990), and Heifetz (1996) study the dynamics of information exchange that leads from disagreement to agreement.

In Bacharach (1985) and Samet (1990), Aumann’s theorem is generalized to a non-partition information structure. In Monderer and Samet (1989), common p -belief,

rather than common knowledge, is shown to imply a bound on the probabilistic disagreement.

All the above-mentioned agreement theorems study the impossibility of commonly sharing probabilistic *disagreements*. This paper, in contrast, studies the impossibility of common knowledge of probabilistic *agreements*. As it turns out, the impossibility of having common knowledge of agents' posterior for an event may be the result not of the probabilistic disagreement, but rather the structure of the event and the partitions. While the starting point of the literature on agreement theorems is the probabilistic structure juxtaposed with the knowledge structure, here it is the knowledge structure alone. We look for conditions that characterize the events for which there exists a probabilistic structure that makes it possible to have common knowledge of agreements regarding the posterior of these events.

1.8 Outline

In the next section, we introduce the basics of the model of knowledge and belief. The main elements of our agreeing theorems, *ignorance* and the *possibility of agreeing to agree* are defined in Section 3, and the main results are stated. A counterexample for uncountable information structures is presented in Section 4. Two open problems are discussed in Section 5 and the proofs are given in the last section.

2. PRELIMINARIES

2.1 Information structures

We fix a *state space* (Ω, \mathcal{B}) , with a set of *states* Ω and a σ -field of *events* \mathcal{B} . An *information structure* on the state space, for agents 1 and 2, is a pair (Π_1, Π_2) of countable measurable partitions of Ω . For every state ω , we denote by $\Pi_i(\omega)$ the element in Π_i that contains ω .

For an event F and agent i , we denote by Π_i^F the partition $\Pi_i \vee \{F, \neg F\}$, which is the join of the partitions Π_i and $\{F, \neg F\}$.¹ It describes the information of the agent when, in addition to the information given by Π_i , she is also informed whether F occurred.

2.2 Knowledge and common knowledge

We say that agent i *knows* event E at ω when $\Pi_i(\omega) \subseteq E$. Thus, the event that i knows E is $K_i(E) = \{\omega \mid \Pi_i(\omega) \subseteq E\}$. The event that both agents know E is $K(E) = K_1(E) \cap K_2(E)$. The event that E is *common knowledge* is $K^\infty(E) = \bigcap_{n=1}^\infty K^n(E)$. It is the union of all the elements of the join $\Pi_1 \vee \Pi_2$ contained in E . For further discussion of knowledge and common knowledge, see Aumann (1976), Monderer and Samet (1989), Geanakoplos (1994), Fagin et al. (1995), and Aumann (1999).

Note that the epistemic operators K_i , K , and K^∞ are defined with respect to the information structure (Π_1, Π_2) . In the sequel we consider also the information structure (Π_1^F, Π_2^F) on the same state space. We make it clear when we use the epistemic operators with respect to this latter information structure.

¹The *join* of two partitions is their finest common coarsening.

The event F belongs to the join $\Pi_1^F \vee \Pi_2^F$. Thus, in the information structure (Π_1^F, Π_2^F) , $F = K^\infty(F)$ and $\neg F = K^\infty(\neg F)$. Indeed, this information structure can be defined as the one obtained from making F and $\neg F$ common knowledge.

2.3 Posteriors and priors

A *posterior* of agent i , for the information structure (Π_1, Π_2) , is a family of probability measures $\{\mu_i^\omega\}$ on (Ω, \mathcal{B}) for $i = 1, 2$ and each ω , such that $\mu_i^\omega(\Pi_i(\omega)) = 1$ and for each $\omega' \in \Pi_i(\omega)$, $\mu_i^{\omega'} = \mu_i^\omega$. A *common prior* for the posteriors of the agents is a probability measure μ on (Ω, \mathcal{B}) such that $\mu_i^\omega(\cdot) = \mu(\cdot \mid \Pi_i(\omega))$ whenever $\mu(\Pi_i(\omega)) > 0$.

3. AGREEING TO AGREE: DEFINITIONS AND MAIN RESULTS

3.1 Agreements

An agreement is an event that the posteriors of the agents coincide at some given probability. The formal definition follows.

DEFINITION 1. Let E be an event in the state space (Ω, \mathcal{B}) with information structure (Π_1, Π_2) and posteriors $\{\mu_i^\omega\}$. An *agreement on E* is an event

$$A = \{\omega \mid \mu_1^\omega(E) = \mu_2^\omega(E) = p\} \tag{1}$$

for some $0 < p < 1$.

Note that this definition requires that the agreed upon posteriors be strictly between 0 and 1.

Next we formally define the possibility that there is common knowledge of an agreement.

DEFINITION 2. Let E be an event in the state space (Ω, \mathcal{B}) with information structure (Π_1, Π_2) . *Agreeing to agree is possible for E* if there are posteriors with a common prior μ and an agreement A on E for these posteriors such that $\mu(K^\infty(A)) > 0$.

DEFINITION 3. The event that agent i is *ignorant* of E is

$$I_i(E) = \neg K_i(E) \cap \neg K_i(\neg E).$$

The event that both agents are ignorant of E is $I(E) = I_1(E) \cap I_2(E)$.

Thus, $I(E)$ is the event that none of the agents can tell whether E is the case or $\neg E$ is the case.²

²The negation of the operator I_i is the *knowing whether* operator, $J_i(E) = K_i(E) \cup K_i(\neg E)$, the properties of which were studied in Hart et al. (1996).

3.2 The main results

Our first theorem states that agreeing to agree is possible for E if and only if there exists a nonempty finite event F , at which it is common knowledge that both agents are ignorant of E . Moreover, this holds true also after F becomes common knowledge. Formally and succinctly, we can state the following theorem.

THEOREM 1. *Agreeing to agree is possible for event E if and only if there exists a nonempty finite event F such that*

$$F \subseteq K^\infty(I(E)) \quad (2)$$

with respect to both (Π_1, Π_2) and (Π_1^F, Π_2^F) .

The fact that at F it is common knowledge that both agents are also ignorant of E is expressed by $F \subseteq K^\infty(I(E))$, with respect to (Π_1, Π_2) . The fact that it remains common knowledge that both agents are also ignorant of E after F becomes common knowledge is conveyed by $F \subseteq K^\infty(I(E))$, with respect to (Π_1^F, Π_2^F) . When the information structure is finite, the condition for agreeing to agree becomes simpler.

COROLLARY. *If the information structure is finite, then agreeing to agree is possible for E , if and only if*

$$K^\infty(I(E)) \neq \emptyset. \quad (3)$$

To see this, note first that since $F \neq \emptyset$, (2) implies $K^\infty(I(E)) \neq \emptyset$. Conversely, if (3) holds, define $F = K^\infty(I(E))$, where F is obviously nonempty and finite. Moreover, since F is an element of the join partition, $\Pi_1 \vee \Pi_2$, it follows that $\Pi_i^F = \Pi_i$ for $i = 1, 2$. Therefore, (2) holds with respect to both (Π_1, Π_2) and (Π_1^F, Π_2^F) .

Notice that the condition for agreeing to agree in [Theorem 1](#) is formulated solely in terms of the information structure. The next theorem provides a condition in probabilistic terms.

THEOREM 2. *Agreeing to agree is possible for E if and only if it is possible with a common prior that has a finite support.*

3.3 Positive priors

A common prior determines the posteriors only on elements of the partition to which it assigns positive probability; on elements of probability zero, the posteriors can be defined arbitrarily. Of special interest are common priors that determine uniquely the posteriors, that is, common priors that assign a positive probability to all the elements in both partitions. We next characterize events for which agreeing to agree is possible with such a common prior. This characterization is done in terms of events called *loops*, which play a central role also in the proof of [Theorem 1](#).

DEFINITION 4. *A loop on E is an event L that consists of $2n$ distinct points $\omega_1, \omega'_1, \dots, \omega_n, \omega'_n$ for some integer $n \geq 1$, such that*

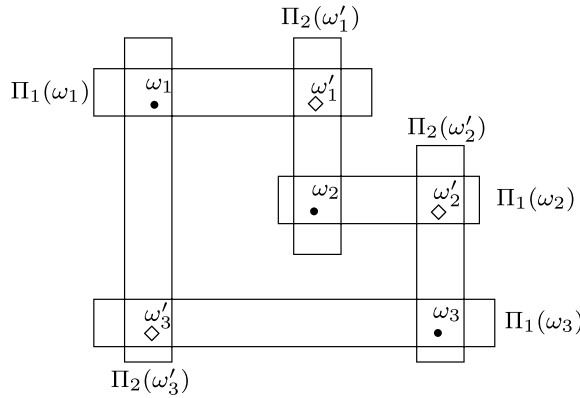


FIGURE 1. A loop on E with 6 points. The dot shaped states, $\omega_1, \omega_2, \omega_3$, are in E , and the diamond shaped states, $\omega'_1, \omega'_2, \omega'_3$, are in $\neg E$.

- (i) $\omega_k \in E$ and $\omega'_k \in \neg E$ for any $k = 1, \dots, n$, and
- (ii) $\omega'_k \in \Pi_1(\omega_k)$ and $\omega_k \in \Pi_2(\omega'_{k-1})$ for any $k = 1, \dots, n$, where $\omega'_0 = \omega'_n$.

Figure 1 illustrates a loop. The following two propositions relate loops to events F that satisfy (2).

PROPOSITION 1. *If F is a nonempty finite event such that $F \subseteq K^\infty(I(E))$ with respect to (Π_1^F, Π_2^F) , then F contains a loop on E .*

PROPOSITION 2. *If L is a loop on E , then $L \subseteq K^\infty(I(E))$ with respect to (Π_1^L, Π_2^L) .*

In the next theorem, we state formally that agreeing to agree is possible for E with a common prior that assigns positive probability to each element of each partition, if and only if it is common knowledge at some state that neither of the players excludes the possibility of loops.

THEOREM 3. *Agreeing to agree is possible for E with a common prior that assigns a positive probability to each element of Π_1 or Π_2 if and only if $K^\infty(\neg K_1(\neg \hat{L}) \cap \neg K_2(\neg \hat{L})) \neq \emptyset$, where \hat{L} is the union of all loops on E .*

4. UNCOUNTABLE INFORMATION STRUCTURES

So far we have assumed that the partitions are countable. To study the question of agreeing to agree in models with uncountable partitions, we need first to define in such models posteriors and common priors.

4.1 Posteriors and common prior for the uncountable case

Posteriors $\{\mu_i^\omega\}$ are defined as in the countable case with an additional measurability requirement: $\mu_i^\omega(E)$ is measurable as a function of ω for any fixed event E and $i = 1, 2$. Equivalently, for $i = 1, 2$, any fixed event E , and p , $\{\omega \mid \mu_i^\omega(E) \geq p\}$ is a measurable set.

A probability measure μ is a common prior for these posteriors if for $i = 1, 2$ and each event E ,

$$\mu(E) = \int \mu_i^\omega(E) d\mu(\omega). \tag{4}$$

In particular, using the properties of posteriors,

$$\mu(\Pi_i(\omega_0) \cap E) = \int_{\Pi_i(\omega_0)} \mu_i^\omega(E) d\mu(\omega) = \mu_i^{\omega_0}(E)\mu(\Pi_i(\omega_0)).$$

This implies that if $\mu(\Pi_i(\omega_0)) > 0$, then $\mu_i^{\omega_0}(E) = \mu(E | \Pi_i(\omega_0))$, which shows that (4) indeed generalizes the definition of a common prior in the countable partitions case.³

4.2 Theorem 1 and the uncountable case

The proof that the condition in Theorem 1 is sufficient does not make use of the countability of the information structure and, therefore, it holds also for uncountable information structures. The proof that the condition is necessary relies on the countability assumption. As shown by the following example, depicted in Figure 2, the assumption is essential: the condition is not necessary in the uncountable case.

EXAMPLE 5. Consider a state space Ω that consists of the union of the following four disjoint sets A_i , $i = 1, \dots, 4$, in \mathbb{R}^2 : $A_1 = \{(x, x + 1) \mid -1 < x < 0\}$, $A_2 = \{(x, x) \mid$




FIGURE 2. Theorem 1 fails in the uncountable case. The state space consists of the three diagonals A_1 , A_2 , and A_3 , and of A_4 . The latter is obtained by a rightward shift of the top-right diagonal by an irrational number c . Agent 1 is informed of the first coordinate of the state and agent 2 is informed of the second. The event E is the union of A_1 and A_3 . There are no loops on E ; nevertheless, agreeing to agree is possible with the common prior that is uniform on the sets A_i .

³See Samet (1999, 2000) for studies of belief systems with uncountable information structures.

$-1 < x < 0$), $A_3 = \{(x, x - 1) \mid 0 < x < 1\}$, and $A_4 = \{(x, \varphi(x)) \mid 0 < x < 1\}$, where $\varphi(x) = x - c \pmod{1}$ for a fixed irrational c in $(0, 1)$. The σ -field on Ω is the Borel σ -field on the sets A_i .

Agent 1 is informed of the first coordinate of the state and agent 2 is informed of the second. Thus, each element of Π_1 contains the two points on the vertical line that contains the state. Similarly, Π_2 contains the two points on the horizontal line that includes the state.

The posterior μ_i^ω for each of the two points in $\Pi_i(\omega)$ is $\frac{1}{2}$. To see that the measurability condition is satisfied, consider $P: \Omega \rightarrow \mathbb{R}$, the projection of the state space on the first coordinate. Note that both P and P^{-1} map measurable sets to measurable sets. Now, if $p \in (\frac{1}{2}, 1]$ and E is a measurable set, then $\{\omega \mid \mu_1^\omega(E) \geq p\} = P^{-1}(P(E \cap A_1) \cap P(E \cap A_2)) \cup P^{-1}(P(E \cap A_3) \cap P(E \cap A_4))$, which is measurable. Alternatively, if $p \in (0, \frac{1}{2}]$, then $\{\omega \mid \mu_1^\omega(E) \geq p\} = P^{-1}(P(E \cap A_1) \cup P(E \cap A_2)) \cup P^{-1}(P(E \cap A_3) \cup P(E \cap A_4))$, which is also measurable. A similar argument shows that the measurability condition holds for the second agent's posterior as well.

Let μ be the probability measure $\frac{1}{4} \sum_{i=1}^4 \mu_i$, where μ_i is the Lebesgue measure on A_i . It is easy to see that μ is a common prior for $\{\mu_i^\omega\}$. Denote $E = A_1 \cup A_3$. For $i = 1, 2$ and every ω , $\mu_i^\omega(E) = \frac{1}{2}$. Thus agreeing to agree is possible for E . We show that there is no loop for E and, therefore, by Proposition 1, Theorem 1 does not hold in this case.

Suppose that $\omega_1 = (a - 1, a) \in A_1 \subseteq E$ is the first point in a loop. Then ω'_1 is necessarily $(a - 1, a - 1) \in A_2 \subseteq \neg E$. The next two points must be $\omega_2 = (a, a - 1) \in A_3$ and $\omega'_2 = (a, \varphi(a)) \in A_4$. Thus $\omega_3 = (\varphi(a) - 1, \varphi(a)) \in A_1$. Continuing this way, we obtain that the second coordinates of the points $\{\omega_{2k+1}\}_{k \geq 0}$ are $\{\varphi\}^k(a)_{k \geq 0}$. Since c is irrational, all the numbers in the latter sequence are distinct and, therefore, so are all the points in the first sequence. A similar argument shows that each of the sequences $\{\omega'_{2k+1}\}_{k \geq 0}$, $\{\omega_{2k}\}_{k \geq 1}$, and $\{\omega'_{2k}\}_{k \geq 1}$ also consists of distinct points. These four sequences are disjoint as they belong to different sets A_i . Thus, the sequence starting with $\omega_1 \in A_1$ is infinite. This also shows that any sequence starting with $\omega_1 \in A_3$ is also infinite and, therefore, there cannot be a loop L on E . \diamond

4.3 The pathology in this example

Each element of the join $\Pi_1 \vee \Pi_2$ can be written as a doubly infinite countable sequence

$$\dots, \omega_{-k}, \omega'_{-k}, \dots, \omega_1, \omega'_1, \dots, \omega_k, \omega'_k, \dots,$$

such that $(\{\omega_k, \omega'_k\})_{k=-\infty}^\infty \subseteq \Pi_1$ and $(\{\omega'_k, \omega_{k+1}\})_{k=-\infty}^\infty \subseteq \Pi_2$. This sequence is constructed in both directions in the same way as the sequence above. If such a sequence is considered a subspace, there is no common prior on this subspace. Indeed, a common prior ν should satisfy, for each k , $\nu(\omega_k) = \nu(\omega'_k)$ and also $\nu(\omega'_k) = \nu(\omega_{k+1})$. Thus, all the states must have the same probability, which is impossible.

The pathology in this example concerns the question of the existence of a common prior for given posteriors. On one hand, there exists a common prior since the measure μ satisfies (4). On the other hand, for any element of the join $\Pi_1 \vee \Pi_2$, there is no common prior. It seems that the subspace relevant to the agents at a given state is

the event of the join containing this state. This is the smallest event that is commonly known by the agents at this state. In our space, at any state there is no common prior on the relevant subspace containing it.

This example raises conceptual difficulties concerning the notion of a common prior in the case of uncountable information structures. The problem was first observed (in a different state space) and discussed in Simon (1998).

5. OPEN PROBLEMS

5.1 *Uncountable information structures*

The necessity of the condition in Theorem 1 is proved by negation. We show that if agreeing to agree holds for E and the condition fails, then the prior probability of each element of Π_1 is zero. This constitutes a contradiction because the partition is countable. Here we use the countability of only one of the partitions, but as we explain next, the proof requires that both be countable.

Propositions 3 and 4 are the two main parts of the proof that the condition in Theorem 1 is necessary. In the proof of the first one, we define a binary relationship on Π_1 . Using this relation, we coarsen both partitions by grouping together the elements of the partitions. The second proposition uses this grouping to show that elements of Π_1 that do not intersect a loop must be assigned zero probability by the common prior.

The grouping of the elements of the partitions in the proof of Proposition 3 does not require countable partitions. However, the countability of the partitions is used to establish that the grouped elements are *measurable* sets. Although the statement of Proposition 4 refers solely to elements in Π_1 , the proof makes use also of the measurability of events resulting from the grouping procedure over Π_2 . These events are measurable due to countability. Therefore, our proofs hinge on the countability of *both* partitions.

The countability assumption in Theorem 1 is essential, because, as has been shown in the previous section, the condition in the theorem is not necessary in the uncountable case. It has been noted, though, that the common prior, which enables the agreeing to agree in the previous section, is pathological in the sense that it fails to induce common priors on the elements of the join of the partitions. This suggests the following problem.

OPEN PROBLEM 1. Consider a state space with an *uncountable* information structure. Assume that there is a common prior *that induces a common prior on each nonnull element of the join of the partitions*.⁴ Does this imply that agreeing to disagree is possible if and only if the condition in Theorem 1 is satisfied?

5.2 *More than two agents*

All our results were formulated and proved for models with two agents. Generalizing the notion of a loop to many agents may enable a generalization of the sufficiency part of Theorem 1 to more than two agents. However, the proof of the necessity part strongly depends on having two agents.

⁴This is the case, for instance, when the whole space is the only element of the join.

The binary relationship employed in the proof of Proposition 3, as well as the grouping of the elements of the partitions, is defined only for two agents. We know of no way to generalize it to more than two agents.

OPEN PROBLEM 2. Find a necessary and sufficient condition in terms of the information structure for the possibility of agreeing to agree in models with more than two agents.

6. PROOFS

PROOF OF PROPOSITION 1. Let F be a nonempty event that satisfies (2) with respect to the information structure (Π_1^F, Π_2^F) . For each $\omega \in F$ and $i = 1, 2$, we have $\Pi_i^F(\omega) \subseteq F \subseteq K^\infty(I(E)) \subseteq I(E)$ (where K^∞ and I are considered as operators with respect to (Π_1^F, Π_2^F)). Hence, $\Pi_i^F(\omega) \cap E \neq \emptyset$ and $\Pi_i^F(\omega) \cap \neg E \neq \emptyset$, $i = 1, 2$. Therefore, there is a sequence⁵ $\omega_1, \omega'_1, \dots, \omega_k, \omega'_k, \dots$ in F such that $\omega_k \in E$, $\omega'_k \in \neg E$, $\omega'_k \in \Pi_1^F(\omega_k)$, and $\omega_{k+1} \in \Pi_2^F(\omega'_k)$, for every $k \geq 1$. Since for each $\omega \in F$ and $i = 1, 2$, $\Pi_i^F(\omega) = \Pi_i(\omega) \cap F$, it follows that $\omega'_k \in \Pi_1(\omega_k)$ and $\omega_{k+1} \in \Pi_2(\omega'_k)$ for every $k \geq 1$.

Note that since $\omega_k \in E$ and $\omega'_k \in \neg E$, then $\omega_k \neq \omega'_k$. Thus, as F is finite, the sequence must contain a subsequence $\omega_{k+1}, \omega'_{k+1}, \dots, \omega_{k+n}, \omega'_{k+n}$ of $2n$ distinct points, for $n \geq 1$ such that $\omega_{k+n+1} = \omega_{k+1}$. This sequence is a loop on E . □

PROOF OF PROPOSITION 2. Consider a loop on E ,

$$L = \{\omega_1, \omega'_1, \dots, \omega_n, \omega'_n\}.$$

By the first condition of Definition 4, $L \subseteq I(E)$. By the second condition, some coarsening of $\{\{\omega_1, \omega'_1\}, \dots, \{\omega_n, \omega'_n\}\}$ is included in Π_1^F and some coarsening of $\{\{\omega'_1, \omega_2\}, \dots, \{\omega'_n, \omega_1\}\}$ is included in Π_2^F . Therefore, $L \subseteq \Pi_1 \vee \Pi_2$. Thus, $L \subseteq K^\infty(I(E))$ with respect to (Π_1^F, Π_2^F) . □

The proof of Theorem 1 requires the following two propositions.

PROPOSITION 3. Let $P_0 \in \Pi_1$. Suppose that for every loop L on E , $L \cap P_0 = \emptyset$. Then there is a coarsening of Π_1 , $\{P_E, P_{\neg E}, P_0\}$, and a coarsening of Π_2 , $\{Q_E, Q_{\neg E}\}$, such that⁶

$$(P_E \cup P_0) \cap Q_E \subseteq E \tag{5}$$

and

$$(P_{\neg E} \cup P_0) \cap Q_{\neg E} \subseteq \neg E. \tag{6}$$

PROOF. In this proof, P and Q denote generic elements of Π_1 and Π_2 , respectively. We define a binary relation ρ on Π_1 : $P \rho P'$ if for some Q , $P \cap Q \cap E \neq \emptyset$ and $P' \cap Q \cap \neg E \neq \emptyset$. We denote by \succ the transitive closure⁷ of ρ .

⁵The sequence may be constructed inductively.

⁶We allow each of the sets $P_E, P_{\neg E}, Q_E$, and $Q_{\neg E}$ to be empty.

⁷That is, $P \succ P'$ if there are P_1, \dots, P_n in Π_1 , with $n \geq 2$, such that $P_1 = P, P_n = P'$, and, for $j = 1, \dots, n - 1$, $P_j \rho P_{j+1}$.

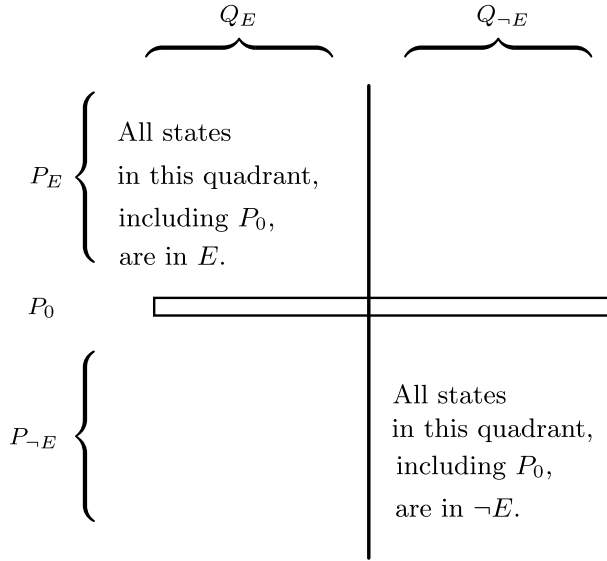


FIGURE 3. An illustration for Proposition 3. The elements of Π_1 are horizontal; those of Π_2 are vertical.

Obviously, \succ is transitive. We show that $P_0 \not\succeq P_0$. Suppose to the contrary that $P_0 \succ P_0$. Then there is a sequence P_0, \dots, P_n , such that

$$P_0 = P_n \rho \cdots \rho P_1 \rho P_0. \tag{7}$$

Moreover, by taking the shortest sequence of this kind, we can assume that P_1, \dots, P_n are distinct. By (7), there are Q_0, \dots, Q_{n-1} in Π_2 such that for $k = 1, \dots, n$, $P_k \cap Q_{k-1} \cap E \neq \emptyset$ and $P_{k-1} \cap Q_{k-1} \cap \neg E \neq \emptyset$.

For each $k = 1, \dots, n$ choose $\omega_k \in P_k \cap Q_{k-1} \cap E$ and $\omega'_k \in Q_k \cap P_k \cap \neg E$, where $Q_n = Q_0$. All $2n$ points are distinct. Indeed, $\omega_k \neq \omega'_k$ because $\omega_k \in E$ and $\omega'_k \in \neg E$. Moreover, since $\{\omega_k, \omega'_k\} \subseteq P_k$ and P_1, \dots, P_n are all distinct, each pair $\{\omega_k, \omega'_k\}$ belongs to a different element of Π_1 . Thus, the points $\omega_1, \omega'_1, \dots, \omega_n, \omega'_n$ form a loop on E that intersects P_0 at ω_n . This contradicts the assumption and we conclude that $P_0 \not\succeq P_0$.

Define P_E to be the union of all elements P that satisfy $P \succ P_0$. Since $P_0 \not\succeq P_0$, it follows that $P_0 \cap P_E = \emptyset$. The set P_{-E} is the complementary set of $P_E \cup P_0$. An element Q is a subset of Q_{-E} if and only if there exists P , such that $P \subseteq P_E \cup P_0$ and $P \cap Q \cap \neg E \neq \emptyset$. The set Q_E is the complement of Q_{-E} .

By the definition of Q_E , for each $Q \subseteq Q_E$ and $P \subseteq P_E \cup P_0$, $P \cap Q \cap \neg E = \emptyset$, which proves (5). To prove (6), suppose to the contrary that for some $P \subseteq P_{-E} \cup P_0$ and $Q \subseteq Q_{-E}$, $P \cap Q \cap E \neq \emptyset$. By the definition of Q_{-E} , there is $P' \subseteq P_E \cup P_0$ such that $P' \cap Q \cap \neg E \neq \emptyset$. Thus, $P \succ P'$. (See Figure 3.)

Now, if $P' = P_0$, then $P \succ P_0$. Else $P' \subseteq P_E$, in which case $P \succ P' \succ P_0$. Thus, in either case $P \succ P_0$, which implies that $P \subseteq P_E$. This contradicts the assumption that $P \subseteq P_{-E} \cup P_0$ and the proof of (6) is also complete. \square

PROPOSITION 4. *Let E be an event, let $\{\mu_i^\omega\}$ be posteriors with a common prior μ , and let A be the agreement $\{\omega \mid \mu_1^\omega(E) = \mu_2^\omega(E) = p\}$ for some $0 < p < 1$. If $P_0 \in \Pi_1$, $P_0 \subseteq K^\infty(A)$, and $P_0 \cap L = \emptyset$ for all loops L on E , then $\mu(P_0) = 0$.*

PROOF. In case $\mu(K^\infty(A)) = 0$, the proof is complete since $P_0 \subseteq K^\infty(A)$. Otherwise, define $\nu(\cdot) = \mu(\cdot \mid K^\infty(A))$. Note that ν is also a common prior for the posteriors⁸ $\{\mu_i^\omega\}$. Moreover, if $\nu(\Pi_i(\omega)) > 0$, then $\mu_i^\omega(E) = p$.

By **Proposition 3**, there are partitions $\{P_E, P_{-E}, P_0\}$ and $\{Q_E, Q_{-E}\}$ that satisfy (5) and (6). The events P_E and Q_{-E} are unions of elements of Π_1 and Π_2 , respectively, and since the partitions are countable, these events are measurable. From the definitions of A and ν , it follows that $\nu(E \cap P_E) = p\nu(P_E)$ and $\nu(\neg E \cap P_E) = (1 - p)\nu(P_E)$. Thus (recall, $1 - p > 0$),

$$\nu(E \cap P_E) = (p/(1 - p))\nu(\neg E \cap P_E). \tag{8}$$

By a similar argument,

$$\nu(E \cap Q_{-E}) = (p/(1 - p))\nu(\neg E \cap Q_{-E}). \tag{9}$$

By (6), the relation between the events on the left sides of (8) and (9) is $E \cap Q_{-E} \subseteq E \cap P_E$. Hence, by comparing the right sides of (8) and (9), we obtain (after dividing by $p/(1 - p) \neq 0$)

$$\nu(\neg E \cap P_E) \geq \nu(\neg E \cap Q_{-E}).$$

By (5), the event on the left side is a subset of the event on the right side. Thus their difference must have probability 0. This difference is

$$\begin{aligned} (\neg E \cap Q_{-E}) \setminus (\neg E \cap P_E) &= \neg E \cap (Q_{-E} \setminus P_E) \\ &= \neg E \cap (Q_{-E} \cap (P_{-E} \cup P_0)) \\ &= Q_{-E} \cap (P_{-E} \cup P_0), \end{aligned}$$

where the last equality holds by (6).

We conclude that $\nu(Q_{-E} \cap (P_{-E} \cup P_0)) = 0$ and, in particular, $\nu(Q_{-E} \cap P_0) = 0$. By (5), $Q_{-E} \cap P_0 \subseteq \neg E$. Hence, $Q_{-E} \cap P_0 = \neg E \cap P_0$. Thus, $\nu(\neg E \cap P_0) = 0$ and, therefore, $\mu(\neg E \cap P_0) = 0$. However, by assumption, $P_0 \subseteq K^\infty(A)$ and, therefore, for every $\omega \in P_0$, $\mu_1^\omega(\neg E) = 1 - p > 0$. This is possible only when $\mu(P_0) = 0$. \square

PROOF OF THEOREM 1. Suppose that a nonempty finite F is a subset of $K^\infty(I(E))$ with respect to both (Π_1, Π_2) and (Π_1^F, Π_2^F) . Thus, $K^\infty(I(E))$ with respect to (Π_1, Π_2) is nonempty.

From now on (in this proof), the operators K^∞ and I are used with respect to (Π_1, Π_2) . We construct posteriors $\{\mu_i^\omega\}$ with a common prior μ such that $\mu(K^\infty(I(E))) >$

⁸This is so since $K^\infty(A)$ is in the join of the partitions. Therefore, $\nu(\cdot)$ induces over $K^\infty(A)$ the same posteriors as μ , while on the complement of $K^\infty(A)$, the posteriors can be defined so as to coincide with $\{\mu_i^\omega\}$.

0 and $K^\infty(I(E)) \subseteq A$, where $A = \{\omega \mid \mu_1^\omega(E) = \mu_2^\omega(E) = \frac{1}{2}\}$. This inclusion implies $K^\infty(I(E)) \subseteq K^\infty(A)$, which shows that $\mu(K^\infty(A)) > 0$ and thereby establishes the fact that agreeing to agree is possible for E .

By Propositions 1 and 2, there exists a loop $L = \{\omega_1, \omega'_1, \dots, \omega_n, \omega'_n\}$ on E such that $L \subseteq K^\infty(I(E))$. Define a common prior μ such that $\mu(\omega_k) = \mu(\omega'_k) = 1/(2n)$ for $k = 1, \dots, n$. Consider $P \in \Pi_1$ such that $P \subseteq K^\infty(I(E))$.

Suppose that $P \cap L \neq \emptyset$. Then $\omega_k \in P$ if $\omega'_k \in P$. Since all the points in L are distinct, it follows that P contains the same number of points in $E \cap F$ and in $\neg E \cap F$. Thus for $\omega \in P$, $\mu_1^\omega(E) = \frac{1}{2}$.

If $P \cap L = \emptyset$, then $\mu(P) = 0$ and, therefore, the posteriors on P can be defined arbitrarily. Since $P \subseteq K^\infty(I(E))$, it follows that $P \subseteq I(E)$ and, therefore, $P \cap E \neq \emptyset$ and $P \cap \neg E \neq \emptyset$. We define the posterior μ_1^ω for $\omega \in P$ such that $\mu_1^\omega(E) = \mu_1^\omega(\neg E) = \frac{1}{2}$.

A similar argument holds for elements $Q \in \Pi_2$. This shows that $K^\infty(I(E)) \subseteq A$, as required.

The proof of the converse direction makes use of Propositions 3 and 4. Suppose that agreeing to agree is possible for E for posteriors with a common prior μ and the agreement A (recall (1)). This means that

$$\mu(K^\infty(A)) > 0. \tag{10}$$

Obviously, $A \subseteq I(E)$ and, therefore,

$$K^\infty(A) \subseteq K^\infty(I(E)). \tag{11}$$

By (10) and (11), $K^\infty(I(E))$ is not empty.

Since $K^\infty(I(E))$ is a union of elements of the join $\Pi_1 \vee \Pi_2$, it follows that for a loop L ,

$$\text{if } L \cap K^\infty(I(E)) \neq \emptyset, \text{ then } L \subseteq K^\infty(I(E)). \tag{12}$$

Assume that there exists a loop L on E such that $L \subseteq K^\infty(I(E))$. By Proposition 2, $F = L$ is a nonempty finite set that satisfies (2) and the proof is complete. Otherwise, for every loop L on E , $L \not\subseteq K^\infty(I(E))$, which implies by (12) that $L \cap K^\infty(I(E)) = \emptyset$. Thus, by (11), $L \cap K^\infty(A) = \emptyset$. Therefore, by Proposition 4, for each $P_0 \in \Pi_1$ such that $P_0 \subseteq K^\infty(A)$, $\mu(P_0) = 0$. Since the partitions are countable, this implies that $\mu(K^\infty(A)) = 0$, which contradicts (10). \square

PROOF OF THEOREM 2. If agreeing to agree is possible for E , then the condition of Theorem 1 holds. As was shown in the first part of its proof, this condition implies that agreeing to agree is made possible by a common prior supported on a loop, which is a finite event. \square

PROOF OF THEOREM 3. Assume first that the event

$$\hat{\Omega} = K^\infty(\neg K_1(\neg \hat{L}) \cap \neg K_2(\neg \hat{L}))$$

is not empty. Then, for each $P \in \Pi_i$, $i = 1, 2$, if $P \subseteq \hat{\Omega}$, then $P \subseteq \neg K_1(\neg \hat{L}) \cap \neg K_2(\neg \hat{L})$ and thus $P \cap \hat{L} \neq \emptyset$.

Since $\hat{\Omega}$ is a union of elements of the join of $\Pi_1 \vee \Pi_2$, it follows that for a loop L , if $L \cap \hat{\Omega} \neq \emptyset$, then $L \subseteq \hat{\Omega}$. Let $\{L_1, L_2, \dots\}$ be the set of all the loops contained in $\hat{\Omega}$. For each n , let μ_n be a probability distribution on L_n such that for each $\omega \in L_n$, $\mu_n(\omega) = 1/(2|L_n|)$.⁹ Let $\mu = \sum 2^{-n}\mu_n$. Then, for $P \in \Pi_i$, $i = 1, 2$, $\mu(P \cap E) = \sum 2^{-n}\mu_n(P \cap E) = \sum 2^{-n}(1/2)\mu_n(P) = (1/2)\mu(P)$.

Since for each $P \in \Pi_i$ that is a subset of $\hat{\Omega}$, $P \cap \hat{L} \neq \emptyset$, it follows that $\mu(P) > 0$. In addition, μ vanishes on the complement of $\hat{\Omega}$. Let μ' be any measure on Ω that vanishes on $\hat{\Omega}$ and is positive on the elements of Π_1 and Π_2 that are not subsets of $\hat{\Omega}$. Define $\nu = \mu/2 + \mu'/2$. Then ν is positive on Π_1 and Π_2 . Let $\{\nu_i^\omega\}$ be the posteriors defined by ν . Then $\hat{\Omega} \subseteq K^\infty(A)$ for $A = \{\omega \mid \nu_1^\omega(E) = \nu_1^\omega(E) = 1/2\}$ and $\nu(\hat{\Omega}) > 0$, which shows that agreeing to agree is possible for E with the prior ν that assigns a positive probability to every element of either partition.

As for the converse direction, assume that there are posteriors $\{\mu_i^\omega\}$ with a common prior μ that is positive on every element of Π_1 or Π_2 . Furthermore, assume that A is an event of the sort $\{\omega \mid \mu_1^\omega(E) = \mu_2^\omega(E) = p\}$ for some $0 < p < 1$ that satisfies $\mu(K^\infty(A)) > 0$. It is sufficient to show that $K^\infty(A) \subseteq K^\infty(\neg K_1(\neg \hat{L}) \cap \neg K_2(\neg \hat{L}))$. For this it suffices to show that for $i = 1, 2$ and any $P \in \Pi_i$ such that $P \subseteq K^\infty(A)$, $P \cap \hat{L} \neq \emptyset$. Indeed, if to the contrary, $P \cap \hat{L} = \emptyset$ for such P , then by Proposition 4, $\mu(P) = 0$. This contradicts the positivity of μ on every element of either partition. \square

REFERENCES

Aumann, Robert J. (1976), “Agreeing to disagree.” *Annals of Statistics*, 4, 1236–1239. [269, 274]

Aumann, Robert J. (1999), “Interactive epistemology I: Knowledge.” *International Journal of Game Theory*, 28, 263–300. [274]

Bacharach, Michael (1985), “Some extensions to a claim of Aumann in an axiomatic model of knowledge.” *Journal of Economic Theory*, 37, 167–190. [273]

Bonanno, Giacomo and Klaus Nehring (1996), “Fundamental agreement: A new foundation for the Harsanyi doctrine.” Working Paper 96-02, Department of Economics, University of California, Davis. [273]

Cover, Thomas M. (1987), “Pick the largest number.” In *Open Problems in Communication and Computation* (Thomas M. Cover and B. Gopinath, eds.), 152, Springer-Verlag, New York. [273]

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi (1995), *Reasoning About Knowledge*. MIT Press, Cambridge, Massachusetts. [274]

Feinberg, Yossi (1995), “A converse to the agreement theorem.” Discussion Paper 83, Center for Rationality and Interactive Decision Theory, The Hebrew University. [273]

⁹Note that this is the same probability used in Theorem 1’s proof.

Geanakoplos, John D. (1994), "Common knowledge." In *Handbook of Game Theory*, volume II (Robert J. Aumann and Sergiu Hart, eds.), 1437–1496, North-Holland, Amsterdam. [274]

Geanakoplos, John D. and Heraklis M. Polemarchakis (1982), "We can't disagree forever." *Journal of Economic Theory*, 28, 192–200. [273]

Hart, Sergiu, Aviad Heifetz, and Dov Samet (1996), "'Knowing whether,' 'knowing that,' and the cardinality of state spaces." *Journal of Economic Theory*, 70, 249–256. [275]

Heifetz, Aviad (1996), "Comment on consensus without common knowledge." *Journal of Economic Theory*, 70, 273–277. [273]

McKelvey, Richard D. and Talbot Page (1986), "Common knowledge, consensus, and aggregate information." *Econometrica*, 54, 109–127. [273]

Milgrom, Paul R. and Nancy L. Stokey (1982), "Information, trade and common knowledge." *Journal of Economic Theory*, 26, 17–27. [273]

Monderer, Dov and Dov Samet (1989), "Approximating common knowledge with common beliefs." *Games and Economic Behavior*, 1, 170–190. [273, 274]

Morris, Stephen (1994), "Trade with heterogeneous prior beliefs and asymmetric information." *Econometrica*, 62, 1327–1347. [273]

Nielsen, Lars Tyge, Adam Brandenburger, John D. Geanakoplos, Richard D. McKelvey, and Talbot Page (1990), "Common knowledge of an aggregate of expectations." *Econometrica*, 58, 1235–1239. [273]

Parikh, Rohit and Paul Krasucki (1990), "Communication, consensus, and knowledge." *Journal of Economic Theory*, 52, 178–189. [273]

Rubinstein, Ariel and Asher Wolinsky (1990), "On the logic of 'Agreeing to disagree' type results." *Journal of Economic Theory*, 51, 184–193. [273]

Samet, Dov (1990), "Ignoring ignorance and agreeing to disagree." *Journal of Economic Theory*, 52, 190–207. [273]

Samet, Dov (1998), "Common priors and the separation of convex sets." *Games and Economic Behavior*, 24, 172–174. [273]

Samet, Dov (1999), "Bayesianism without learning." *Research in Economics*, 53, 227–242. [278]

Samet, Dov (2000), "Quantified beliefs and believed quantities." *Journal of Economic Theory*, 95, 169–185. [278]

Samet, Dov, Iddo Samet, and David Schmeidler (2004), "One observation behind two-envelope puzzles." *American Mathematical Monthly*, 111, 347–351. [273]

Sebenius, Julius and John D. Geanakoplos (1983), "Don't bet on it: Contingent agreements with asymmetric information." *Journal of the American Statistical Association*, 78, 424–426. [273]

Simon, Robert (1998), "The common prior assumption in belief spaces." Unpublished paper. [280]

Submitted 2009-6-16. Final version accepted 2010-8-18. Available online 2010-8-18.