Journal of Logic and Computation Advance Access published January 28, 2016

Limit-agreeing to disagree*

CHRISTIAN W. BACH, *University of Liverpool Management School, University of Liverpool, Chatham Street, Liverpool L69 7ZH, UK.*
EpiCenter, School of Business and Economics, Maastricht University, 6200 MD Maastricht, The Netherlands.
E-mail: c.w.bach@liverpool.ac.uk; bach@epicenter.name

JÉRÉMIE CABESSA, *Laboratory of Mathematical Economics (LEMMA), University Paris 2 – Panthéon-Assas, 4 Rue Blaise Desgoffe, 75006 Paris, France.*
E-mail: jeremie.cabessa@u-paris2.fr

Abstract

We reconsider Aumann' s seminal impossibility theorem that agents cannot agree to disagree in a topologically extended epistemic model. In such a framework, a possibility result on agreeing to disagree actually ensues. More precisely, agents with a common prior belief satisfying limit knowledge instead of common knowledge of their posterior beliefs may have distinct posterior beliefs. Since limit knowledge is defined as the limit of iterated mutual knowledge, agents can thus be said to limit-agree to disagree. Besides, an example is provided in which limit knowledge coincides with Rubinstein's (1989) notion of almost common knowledge, and the agents have almost common knowledge of posteriors yet distinct posterior beliefs. More generally, an epistemic-topological foundation for almost common knowledge is thus provided.

Keywords: Agreeing to disagree, agreement theorems, almost common knowledge, epistemic-topological framework, interactive epistemology, limit knowledge.

1 Introduction

The impossibility for two agents to agree to disagree is established by Aumann's [1] so-called agreement theorem. More precisely, it is shown that if two Bayesian agents equipped with a common prior belief receive private information and have common knowledge of their posterior beliefs, then these posteriors must be equal. In other words, distinct posteriors cannot be common knowledge among Bayesian agents with a common prior. In this sense, agents cannot agree to disagree.

Along these lines, Milgrom and Stokey [20] establish an impossibility theorem of speculative trade. Intuitively, their result states that if two traders agree on a prior efficient allocation of goods, then upon receiving private information, it cannot be common knowledge that both traders have an incentive to trade. From an empirical or quasi-empirical point of view, the agreement theorem seems quite startling since real world agents do frequently disagree on a large variety of issues. It is then natural to scrutinize whether Aumann's basic result still holds with weakened or slightly modified assumptions.

In this spirit, Lewis [18] show that without assuming common knowledge of the posteriors, agents following a specific communication procedure can nevertheless not agree to disagree. Furthermore, Monderer and Samet [21] replace common knowledge by the weaker concept of common p-belief and establish an agreement theorem with such an approximation of common knowledge. Indeed, they show that if the posteriors of Bayesian agents equipped with a common prior are common p-belief

*An extended abstract of a preliminary version of this paper appears under the title 'Agreeing to Disagree with Limit Knowledge' in Hans van Ditmarsch *et al*. (eds.), *Proceedings of the 3rd International Workshop on Logic, Rationality and Interaction (LORI 2011)*, 51–90, LNCS vol. 6953, Springer.

doi:10.1093/logcom/exv082

for large enough p, then these posteriors cannot differ significantly. Besides, a bounded rationality approach is taken by Samet [23] who drops the implicit negative introspection assumption — which states that for every proposition that an agent does not know, the agent actually knows that he does not know it — and establishes that Aumann's agreement theorem remains valid with agents ignorant of their own ignorance. Yet a different generalization is provided by Bacharach's [10] non-probabilistic agreement theorem which states that if two agents follow a common decision procedure in line with the sure thing principle — which states that for every event and every partition of it, whenever each cell of the partition induces a same decision, the event itself generates precisely this decision — and if their particular decisions are common knowledge, then these decisions must coincide. In addition, Bonanno and Nehring [11] as well as Ménager [19] provide comprehensive surveys on works about agreeing to disagree. More recently, Bach and Perea [9] show that Aumann's impossibility result is not robust with respect to the common prior assumption in the sense that two Bayesian agents with arbitrarily close prior beliefs can have common knowledge of completely opposed posteriors, and also provide a lexicographic agreement theorem. Furthermore, the agreement theorem has been analysed from the perspective of dynamic epistemic logic. Notably, Dégremont and Roy [13] obtain a non-probabilistic impossibility result with common belief — instead of common knowledge — of posteriors within the framework of epistemic plausibility models, when the common priors satisfy a specific well-foundedness assumption. Also, several probabilistic agreement theorems are established by Demey [14] using enriched probabilistic Kripke models.

Here, we analyse agreeing to disagree in a topologically extended epistemic framework. The epistemic operator common knowledge is replaced by the epistemic-topological operator limit knowledge introduced by Bach and Cabessa [7, 8]. Assuming common priors, Bayesian agents and limit knowledge of posteriors, we derive a possibility result allowing the agents' posteriors to differ. Since limit knowledge is defined as the topological limit of higher-order mutual knowledge, our result can be interpreted as establishing — in contrast to Aumann's impossibility theorem — that agents can agree to disagree, or more precisely, can limit-agree to disagree. These results show that Aumann's agreement theorem is not robust when considered from a more general epistemic-topological perspective with limit knowledge instead of common knowledge, with both concepts being based on the same sequence of iterated mutual knowledge. We further show that the notion of limit knowledge is able to capture relevant reasoning patterns allowing agents to agree to disagree on their posterior beliefs. In particular, an example is constructed in which limit knowledge is identical with Rubinstein's [22] notion of almost common knowledge, i.e. it coincides with m iterations of mutual knowledge for some finite number m. Thereby, we give an epistemic-topological foundation for Rubinstein's [22] notion of almost common knowledge. However, in general limit knowledge does not coincide with iterated mutual knowledge up to some finite level.

We proceed as follows. In Section 2, the basic framework of set-based interactive epistemology is presented. Aumann's agreement theorem is then given in two slightly modified formulations and graphically illustrated. Section 3 sketches the topological approach to interactive epistemology initiated by Bach and Cabessa [7, 8]. Moreover, the epistemic-topological operator limit knowledge is defined and used for a possibility result on agreeing to disagree. Section 4 presents a plausible epistemic-topological context in which limit knowledge coincides with Rubinstein's [22] almost common knowledge and agents can limit-agree to disagree. Finally, Section 5 offers some concluding remarks and possible directions for future research.

2 Aumann's agreeing to disagree

Set-based interactive epistemology provides the formal framework in which the agreement theorem is established. Having been introduced and notably developed by Aumann [1–5] the discipline furnishes tools to formalize epistemic notions in interactive situations. Before the agreement theorem is restated, the basic ingredients of interactive epistemology are now briefly presented.

A so-called Aumann structure $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i\in I},p)$ consists of a countable set Ω of possible worlds, also called states, which are complete descriptions of the way the world might be, a finite set of agents I, a possibility partition $\mathcal{I}_i$ of Ω for each agent $i\in I$ representing his information, and a common prior belief function $p:\Omega\to[0,1]$ such that $\sum_{\omega\in\Omega}p(\omega)=1$. The cell of $\mathcal{I}_i$ containing the world ω is denoted by $\mathcal{I}_i(\omega)$ and consists of all worlds considered possible by i at world ω. In other words, agent i cannot distinguish between any two worlds ω and ω' that are in the same cell of his partition $\mathcal{I}_i$. Moreover, an event $E\subseteq\Omega$ is defined as a set of possible worlds. For example, the event that it is raining in London contains all worlds in which it does in fact rain in London. Note that the common prior belief function p can naturally be extended to a common prior belief measure on the event space $p:\mathcal{P}(\Omega)\to[0,1]$ by setting $p(E)=\sum_{\omega\in E}p(\omega)$. In this context, it is supposed that each information set of each agent has non-zero prior probability, i.e. $p(\mathcal{I}_i(\omega))>0$ for all $i\in I$ and $\omega\in\Omega$. Such a hypothesis seems plausible since it ensures that no information is excluded a priori. Moreover, all agents are assumed to be Bayesian and to hence update the common prior belief given their private information according to Bayes's rule. More precisely, given some event E and some world ω, the posterior belief of agent i in E at ω is given by $p(E\,|\,\mathcal{I}_i(\omega))=\frac{p(E\cap\mathcal{I}_i(\omega))}{p(\mathcal{I}_i(\omega))}$. Furthermore, an Aumann structure $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i\in I},p)$ is called finite if Ω is finite and infinite otherwise.

In Aumann's epistemic framework, knowledge is formalized in terms of events. More precisely, the event of agent i knowing E, denoted by $K_i(E)$, is defined as $K_i(E):=\{\omega\in\Omega:\mathcal{I}_i(\omega)\subseteq E\}$. If $\omega\in K_i(E)$, then i is said to know E at world ω. Intuitively, i knows some event E if in all worlds that he considers possible the event E holds. Naturally, the event $K(E)=\bigcap_{i\in I}K_i(E)$ then denotes mutual knowledge of E among the set I of agents. Letting $K^0(E):=E$, m-order mutual knowledge of the event E among the set I of agents is inductively defined by $K^{m+1}(E):=K(K^m(E))$ for all $m\geq 0$. Accordingly, mutual knowledge can also be denoted as 1-order mutual knowledge. Furthermore, an event is said to be common knowledge among a set I of agents whenever all m-order mutual knowledge of it simultaneously holds. Formally, it is standard to define the event that E is common knowledge among the set I of agents as the intersection of all higher-order mutual knowledge, i.e. $CK^{inter}(E):=\bigcap_{m>0}K^m(E)$.

An alternative formalization of common knowledge is proposed by Aumann [1] in terms of the meet of the agents' possibility partitions.[1] Accordingly, an event E is called common knowledge at world ω among the set I of agents, if E includes the cell of the meet $\bigwedge_{i\in I}\mathcal{I}_i$ that contains ω. Formally, the meet definition of common knowledge of some event E can be stated as $CK^{meet}(E):=\{\omega\in\Omega:(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\subseteq E\}$. Aumann [1] states and intuitively shows that the standard and the meet definitions of common knowledge are in fact equivalent for finite state spaces. We now formally show the equivalence of the two definitions for the more general case also admitting the set Ω of possible worlds to be infinite, and thereafter, for any event E, the two identical events $CK^{inter}(E)$ and $CK^{meet}(E)$ are thus simply referred to as $CK(E)$.

LEMMA 1
Let $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i\in I},p)$ be an Aumann structure and $E\subseteq\Omega$ be an event. Then, $CK^{inter}(E)=CK^{meet}(E)$.

PROOF. See Appendix. ■

[1] Given two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ of a set S, partition $\mathcal{P}_1$ is called *finer* than partition $\mathcal{P}_2$ or $\mathcal{P}_2$ *coarser* than $\mathcal{P}_1$, if each cell of $\mathcal{P}_1$ is a subset of some cell of $\mathcal{P}_2$. Given n partitions $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_n$ of S, the finest partition that is coarser than $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_n$ is called the *meet* of $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_n$ and is denoted by $\bigwedge_{j=1}^n\mathcal{P}_j$. Moreover, given $x\in S$, the cell of the meet $\bigwedge_{j=1}^n\mathcal{P}_j$ containing x is denoted by $(\bigwedge_{j=1}^n\mathcal{P}_j)(x)$.

Aumann's agreement theorem states that if two agents have a common prior and their posterior beliefs in some event are common knowledge, then these posterior beliefs must coincide. In other words, if two agents with common prior beliefs hold distinct posterior beliefs, then these posterior beliefs cannot be common knowledge among them. Intuitively, it is impossible for agents to consent to distinct beliefs. Thus, agents cannot agree to disagree.

Aumann's result is now formalized in a slightly modified way. More precisely, the introduction of arbitrary values for the agents' posterior beliefs as in Aumann's original statement is dispensed with. Instead, these arbitrary values are substituted by conceivable posterior beliefs, namely updated prior beliefs induced by some auxiliary world. The values of the posteriors are thus made endogenous, as nothing external to the formal structure is needed for their determination. In this sense, the following statement of the agreement theorem is inherent to the formal structure it is embedded in.

AUMANN'S AGREEMENT THEOREM (VERSION 1).
Let $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i=1}^n,p)$ be an Aumann structure, $E\subseteq\Omega$ be an event, and $\hat{\omega}$ be a world. If $CK(\bigcap_{i=1}^n\{\omega'\in\Omega:p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_i(\hat{\omega}))\})\neq\emptyset$, then $p(E\,|\,\mathcal{I}_1(\hat{\omega}))=p(E\,|\,\mathcal{I}_2(\hat{\omega}))=\cdots=p(E\,|\,\mathcal{I}_n(\hat{\omega}))$.

PROOF. See Appendix. ■

The previous theorem states that, if common knowledge of the agents' posterior beliefs being equal to the values induced by some auxiliary world is non-empty, then the agents' posterior beliefs at this given world coincide. Yet as a consequence, the agents' posterior beliefs do not only coincide at the auxiliary world, but also at every possible world inducing the same values as the auxiliary world. In particular, every world contained in common knowledge of the agents' posterior beliefs satisfies equality of posterior beliefs. Thus, a second version of Aumann's agreement theorem ensues as follows.

AUMANN'S AGREEMENT THEOREM (VERSION 2).
Let $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i=1}^n,p)$ be an Aumann structure, $E\subseteq\Omega$ be an event, and $\hat{\omega},\omega\in\Omega$ be worlds such that $CK(\bigcap_{i=1}^n\{\omega'\in\Omega:p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_i(\hat{\omega}))\})\neq\emptyset$ and $\omega\in CK(\bigcap_{i=1}^n\{\omega'\in\Omega:p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_i(\hat{\omega}))\})$. Then, $p(E\,|\,\mathcal{I}_1(\omega))=p(E\,|\,\mathcal{I}_2(\omega))=\cdots=p(E\,|\,\mathcal{I}_n(\omega))$.

PROOF. See Appendix. ■

The two preceding versions of Aumann's Agreement Theorem provide slightly different formulations of the impossibility for agents to agree to disagree on their posterior beliefs. In the first version, emphasis is put on the posterior beliefs induced by some auxiliary world, whereas the second version focuses on the posterior beliefs that the agents actually hold. Basically, the first version intuitively states that, if the agents' posterior beliefs are common knowledge *somewhere*, then they must coincide, while the second version states that, if the agents' *actual* posterior beliefs are common knowledge, then they must coincide.

Agreeing to disagree can be graphically illustrated for the case of two agents. In Figure 1, the set of all possible worlds Ω is partitioned horizontally in equivalence classes of worlds that yield a same posterior belief for agent *Alice* in some fixed event E. Similarly, the vertical slices represent equivalence classes with respect to worlds that induce a same posterior belief for *Bob* in E. Observe that the partition formed by the horizontal slices is coarser than *Alice*'s possibility partition, since Bayesian updating ensures that *Alice*'s posteriors remain constant throughout any cell of her possibility partition. Similarly, the partition formed by the vertical slices is coarser than *Bob*'s possibility partition. Moreover, the intersection of the horizontal and vertical slices forms a refined partition whose cells represent equivalence classes of worlds that induce a same posterior

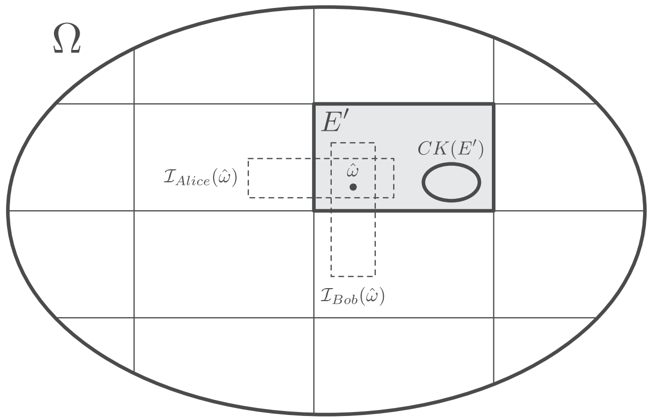

FIGURE 1. Illustration of agreeing to disagree for the case of two agents.

belief profile, i.e. the posterior of each agent remains constant throughout the cell. Given an auxiliary world $\hat{\omega}$ — a world which merely serves the supply of posterior beliefs that can potentially be generated given the constraints of the formal structure — the cell of this refined partition containing $\hat{\omega}$ represents the event $E' = \bigcap_{i \in I} \{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}$ and includes the event $CK(E')$. In particular, note that $CK(E')$ can be empty.

The claim of Aumann's agreement theorem can now be understood graphically: given a world $\hat{\omega}$, if the corresponding cell E' of the refined partition includes a non-empty $CK(E')$, then identical values for the agents' posterior beliefs obtain at $\hat{\omega}$, and hence also at all worlds throughout cell E'. Conversely, common knowledge of the agents' posteriors equals the empty set in all cells of the refined partition in which the posteriors of *Alice* and *Bob* differ. In particular, the equality of the two agents' posteriors throughout $CK(E')$ is established. Finally, note that in the specific Aumann structure represented in Figure 1, at world $\hat{\omega}$, none of the two agents know the true fact that their posteriors coincide, since each agent, respectively, considers possible a world at which the other agent's posterior belief is different. In contrast, for any world $\omega \in CK(E')$, the agents' posterior beliefs are not merely equal at such a world, but the agents also know that they coincide, know that they know that they coincide, etc.

3 An epistemic-topological approach to agreeing to disagree

The standard set-based approach to interactive epistemology lacks a general framework providing some formal notion of closeness between events.[2] By adding a topological dimension to an epistemic structure, it is actually possible to introduce a perception of closeness of events into the reasoning of agents.[3] In such an enriched epistemic-topological framework, the reasoning of agents may thus also depend on topological instead of mere epistemic features of the underlying interactive situation.

[2]For the specific purpose of dealing with counterfactuals, Lewis [18] and Stalnaker [24] consider closeness between possible worlds as a primitive in the semantics of their conditional logics. The basic idea is that for every possible world and every statement, a selection function picks the closest world such that the statement holds true. In contrast to the classical models of Stalnaker and Lewis, we consider closeness between events; have closeness determined by an underlying topology; and do not restrict attention concerning closeness to counterfactuals.

[3]Note that topological spaces can be seen as generalizations of metric spaces. While closeness between elements is explicitly measured via the respective distance function in metric spaces, it is only implicitly determined by open neighbourhoods in topological spaces.

For instance, suppose an agent is reasoning about the weather in London. Intuitively, the event *It is cloudy in London* seems to be closer to the event *It is raining in London* than the event *It is sunny in London*. Now, the agent may make identical decisions being informed only of the truth of some event within a class of *close* events. In fact, the agent might decide to stay at home not only in the case of it raining outside, but also in the case of events perceived by him to be similar, i.e. *close*, such as it being cloudy outside. In a topologically enriched epistemic framework, the notion of closeness is induced by the topology.

Note that such an epistemic-topological approach is of descriptive not normative character. Due to the heterogeneity of real-world agents, different people may have different perceptions of closeness and ways of reasoning. It therefore seems implausible to claim that there is some kind of universal topology that represents the correct perception of the event space that agents should hold. In contrast, different intuitive cognitive-topological patterns of closeness can be formalized and serve as the agents' perceptions of the event space. Such a descriptive use of topologies is in line with Rubinstein's [22] view that topology can be used as a substantial tool to formalize natural intuitions about closeness.[4]

In this context, Bach and Cabessa [7, 8] consider Aumann structures equipped with topologies[5] on the event space and introduce the operator limit knowledge, which is linked to epistemic features as well as topological aspects of the event space. More precisely, limit knowledge is defined as the topological limit of higher-order mutual knowledge.

Definition 1
Let $(\Omega,(\mathcal{I}_i)_{i\in I},p)$ be an Aumann structure, $\mathcal{T}$ a topology on $\mathcal{P}(\Omega)$, and E an event. If the limit point of the sequence $(K^m(E))_{m>0}$ is unique, then $LK(E) := \lim_{m\to\infty} K^m(E)$ is the event that E is limit knowledge among the set I of agents.

Accordingly, limit knowledge of an event E is constituted by — whenever unique — the limit point of the sequence of iterated mutual knowledge, and thus linked to both epistemic as well as topological aspects of the event space. Note that $\mathcal{T}$ being a topology on the event space $\mathcal{P}(\Omega)$ means that $\mathcal{T}\subseteq\mathcal{P}\big(\mathcal{P}(\Omega)\big)$. In contrast, a topology $\mathcal{U}$ on the state space Ω means that $\mathcal{U}\subseteq\mathcal{P}(\Omega)$.

Limit knowledge can be understood as the event which is approached by the sequence of iterated mutual knowledge, according to the notion of closeness between events provided by a given topology on the event space.[6] Thus, the higher the iterated mutual knowledge, the closer this latter epistemic event is to limit knowledge.

Note that limit knowledge should not be amalgamated with common knowledge. Indeed, both operators can be perceived as sharing distinct implicative properties with regards to iterated mutual knowledge. Common knowledge bears a standard implicative relation in terms of set inclusion to all

[4](Cf. Rubinstein [22], p. 390).

[5]For sake of self-containedness, some basic notions from topology are recalled. Given some set X, a *topology* $\mathcal{T}$ on X consists of a family of subsets of X, i.e. $X\subseteq\mathcal{P}(X)$, such that the empty set and X belong to $\mathcal{T}$, and the family $\mathcal{T}$ is closed under finite intersection as well as under arbitrary union. If $\mathcal{T}$ is a topology on X, the pair $(X,\mathcal{T})$ is called *topological space*, where the elements of $\mathcal{T}$ are called *open sets*. For any $p\in X$, an *open neighbourhood* of p is an open set containing p. Moreover, given some subset $S\subseteq X$, an element $p\in X$ is a *limit point* of S if every open neighbourhood of p contains some element of S different from p. Given some sequence $s=(x_i)_{i\geq 0}$ of elements of X, if the limit point of s is unique, it is denoted by $\lim_{i\to\infty} x_i$.

[6]In our epistemic-topological framework, there exist various ways in which a notion of closeness between events can be defined based on their open neighbourhoods. For instance, two events could be said to be close, if they violate the so-called Hausdorff-condition, i.e., if there exist no disjoint neighbourhoods of these two events. Also, degrees of closeness could possibly be defined in the sense that the more neighbourhoods contain two events, the closer the respective events are to each other. In case of the considered topology on the event space to be metrizable, there exists a distance function which could explicitly measure the closeness between events.

iterated mutual knowledge. In contrast, limit knowledge entertains an implicative relation in terms of set proximity with iterated mutual knowledge: the higher the iteration, the closer the respective higher-order mutual knowledge to limit knowledge.

In general, limit knowledge also differs from approximations of common knowledge such as Monderer and Samet's [21] common p-belief as well as Rubinstein's [22] almost common knowledge. Indeed, common p-belief — as infinitely iterated mutual p belief with p-belief being weaker than knowledge — and almost common knowledge — as iterated mutual knowledge up to some finite level only — are both implied by common knowledge, whereas limit knowledge is not.

It is possible to link limit knowledge to reasoning patterns of agents based on closeness of events. In fact, agents satisfying limit knowledge of some event are in a situation infinitesimally close to having arbitrarily-high iterated mutual knowledge of this event, and the agents' reasoning may be influenced accordingly. Note that a reasoning pattern associated with limit knowledge depends on the particular topology on the event space, which fixes the closeness relation between events. For instance, in Section 4 a specific topological structure on the event space is provided, such that limit knowledge corresponds to Rubinstein's [22] almost common knowledge. An epistemic-topological foundation for almost common knowledge is thus provided, which can also be considered a concern of its own sake independent from agreeing to disagree.

The operator limit knowledge is shown by Bach and Cabessa [7, 8] to be able to provide relevant epistemic-topological characterizations of solution concepts in games. Despite being based on the same sequence of higher-order mutual knowledge claims, the distinguished interest of limit knowledge resides in its capacity to potentially differ from the purely epistemic operator common knowledge. Notably, it can be proven that such differing situations necessarily require an infinite event space as well as sequences of higher-order mutual knowledge that are strictly shrinking.[7]

In contrast to the purely epistemic operator common knowledge, for which factiveness holds, i.e. $CK(E) \subseteq E$, the epistemic-topological operator limit knowledge does in general not bear this property.[8] However, if the sequence of higher-order mutual knowledge of some event E is not strictly shrinking, then limit knowledge is equal to common knowledge, and therefore, limit knowledge is factive, i.e. $LK(E) \subseteq E$.[9] In particular, limit knowledge is factive in all finite Aumann structures.

Now, the question whether agents with a common prior belief can agree to disagree on their posterior beliefs is addressed from a topological point of view. The original hypotheses of Aumann's result are modified in that the epistemic operator common knowledge is replaced by the epistemic-topological operator limit knowledge. It is now shown that agents can indeed limit-agree to disagree.

THEOREM 1

There exist an Aumann structure $\mathcal{A} = (\Omega, (\mathcal{I}_i)_{i \in I}, p)$ equipped with a topology $\mathcal{T}$ on the event space $\mathcal{P}(\Omega)$, an event $E \subseteq \Omega$, and worlds $\omega, \hat{\omega} \in \Omega$ such that $\omega \in LK(\bigcap_{i \in I} \{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\})$, as well as both $p(E \mid \mathcal{I}_i(\hat{\omega})) \neq p(E \mid \mathcal{I}_j(\hat{\omega}))$ and $p(E \mid \mathcal{I}_i(\omega)) \neq p(E \mid \mathcal{I}_j(\omega))$ for some agents $i, j \in I$.

[7] Given some event E, the sequence of higher-order mutual knowledge $(K^m(E))_{m>0}$ is called *strictly shrinking* if $K^{m+1}(E) \subsetneq K^m(E)$ for all $m \geq 0$.

[8] For example, in the proof of Theorem 1 an Aumann structure is constructed, the event space is furnished with a topology, and an event E' is considered for which $LK(E') \not\subseteq E'$ holds.

[9] In fact, Bach and Cabessa [8] show that, if limit knowledge and common knowledge are distinct events, then the sequence of higher-order mutual knowledge is strictly shrinking.

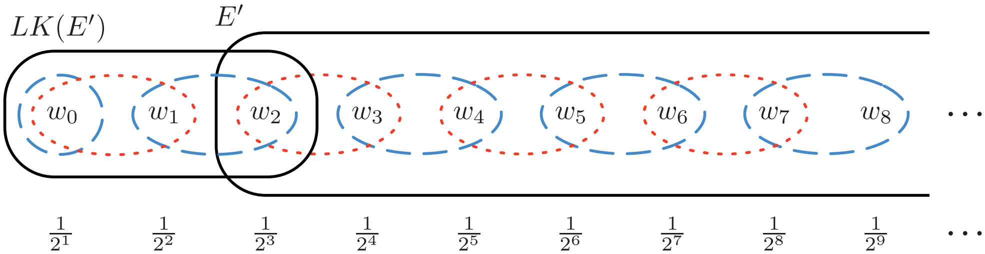

FIGURE 2. Illustration of the Aumann structure described in the proof of Theorem 1. The dotted and dashed sets represent the possibility partitions of *Alice* and *Bob*, respectively. The fractions correspond to the prior probabilities associated to the possible worlds.

PROOF. Consider the Aumann structure $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i\in I},p)$, which is illustrated in Figure 2 and where $\Omega=\{\omega_n : n\geq 0\}$, $I=\{Alice, Bob\}$, $\mathcal{I}_{Alice}=\{\{\omega_{2n},\omega_{2n+1}\}: n\geq 0\}$, $\mathcal{I}_{Bob}=\{\{\omega_0\}\}\cup\{\{\omega_{2n+1},\omega_{2n+2}\}: n\geq 0\}$, as well as $p:\Omega\to\mathbb{R}$ is given by $p(\omega_n)=\frac{1}{2^{n+1}}$ for all $n\geq 0$. Note that the common prior belief function p is well defined since $\sum_{n\geq 0}\frac{1}{2^{n+1}}=1$. Now, consider the event $E=\{\omega_{2n}: n\geq 1\}$, and the world $\omega_2\in\Omega$. Besides, for sake of notational convenience, let the event $\bigcap_{i\in I}\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_i(\omega_2))\}$ be denoted by E'. First of all, observe that $p(E\,|\,\mathcal{I}_{Alice}(\omega_2))=\frac{2}{3}$ and $p(E\,|\,\mathcal{I}_{Bob}(\omega_2))=\frac{1}{3}$. Moreover, $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Alice}(\omega'))=p(E\,|\,\mathcal{I}_{Alice}(\omega_2))=\frac{2}{3}\}=\Omega\setminus\{\omega_0,\omega_1\}$ and $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Bob}(\omega'))=p(E\,|\,\mathcal{I}_{Bob}(\omega_2))=\frac{1}{3}\}=\Omega\setminus\{\omega_0\}$, thus $E'=(\Omega\setminus\{\omega_0,\omega_1\})\cap(\Omega\setminus\{\omega_0\})=\Omega\setminus\{\omega_0,\omega_1\}$. Furthermore, the definitions of the possibility partitions of *Alice* and *Bob* ensure that $K^m(E')=K^m(\Omega\setminus\{\omega_0,\omega_1\})=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_{m+1}\}$, for all $m>0$. Consequently, the sequence $(K^m(E'))_{m>0}$ is strictly shrinking and $CK(E')=\{\omega\in\Omega : (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\subseteq E'\}=\emptyset$. Now, consider the topology $\mathcal{T}$ on $\mathcal{P}(\Omega)$ defined by $\mathcal{T}=\{O\subseteq\mathcal{P}(\Omega) : \{\omega_0,\omega_1,\omega_2\}\notin O\}\cup\{\mathcal{P}(\Omega)\}$. Then, the only open neighbourhood of the event $\{\omega_0,\omega_1,\omega_2\}$ is $\mathcal{P}(\Omega)$, and all terms of the sequence $(K^m(E'))_{m>0}$ are contained in $\mathcal{P}(\Omega)$. Thus $(K^m(E'))_{m>0}$ converges to $\{\omega_0,\omega_1,\omega_2\}$. Moreover, for every event $F\in\mathcal{P}(\Omega)$ such that $F\neq\{\omega_0,\omega_1,\omega_2\}$, the singleton $\{F\}$ is open, and since $K^{m+1}(E')\subsetneq K^m(E')$ for all $m>0$, the strictly shrinking sequence $(K^m(E'))_{m>0}$ will never remain in the open neighbourhood $\{F\}$ of F from some index onwards. Hence $(K^m(E'))_{m>0}$ does not converge to any such event F. Therefore, the limit point $\{\omega_0,\omega_1,\omega_2\}$ of the strictly shrinking sequence $(K^m(E'))_{m>0}$ is unique, and $LK(E')=\lim_{m\to\infty}K^m(E')=\{\omega_0,\omega_1,\omega_2\}$. The event E' and its limit point $LK(E')$ are also illustrated in Figure 2. Next, consider the world ω_1. Note that $\omega_1\in LK(E')$. Also, observe that $p(E\,|\,\mathcal{I}_{Alice}(\omega_2))=\frac{2}{3}\neq\frac{1}{3}=p(E\,|\,\mathcal{I}_{Bob}(\omega_2))$ as well as $p(E\,|\,\mathcal{I}_{Alice}(\omega_1))=0\neq\frac{1}{3}=p(E\,|\,\mathcal{I}_{Bob}(\omega_1))$. Finally, taking $\omega=\omega_1$ and $\hat{\omega}=\omega_2$ concludes the proof. ■

The preceding theorem counters Aumann's impossibility result in the sense of showing that agents can limit-agree to disagree. More precisely, agents may hold distinct actual posterior beliefs, while at the same time satisfying limit knowledge of their posteriors being equal to the specific values induced by a given auxiliary world. Hence, agents may agree in the sense of satisfying limit knowledge of their posteriors, while at the same time disagree in the sense of actually having different posterior beliefs.

The interactive situation depicted in the proof of Theorem 1 shows that agreeing to disagree becomes possible when common knowledge is substituted by limit knowledge. Thus, Aumann's impossibility result no longer holds when moving to a topologically enriched context. In such an epistemic-topological framework, agents can now be seen to have cognitive access to a further dimension in their reasoning that permits them to agree to disagree on their posterior beliefs. More precisely, the agents are in a limit situation of having higher-order mutual knowledge of their posteriors, which, in connection with the particular notion of closeness provided by the topology, leads them to actually possess different posterior beliefs.

Note that the epistemic model constructed in the proof of Theorem 1 could actually be strengthened in the sense of generating different posterior beliefs for the agents not only at some but at all worlds contained in limit knowledge. For instance, taking as topology on the event space $\mathcal{T}=\{O\subseteq \mathcal{P}(\Omega):\{\omega_n : n>0\}\notin O\}\cup\{\mathcal{P}(\Omega)\}$ yields $LK(\bigcap_{i\in I}\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_{i\in I}(\hat{\omega}))\})=\Omega\setminus\{\omega_0\}$, and hence different posterior beliefs for *Alice* and *Bob* at all worlds in $\Omega\setminus\{\omega_0\}$. More generally, it would actually be possible to make limit knowledge correspond to any event F by equipping the event space with the corresponding excluded point topology $\mathcal{T}=\{O\subseteq\mathcal{P}(\Omega):F\notin O\}\cup\{\mathcal{P}(\Omega)\}$. This topology does not have any intuitive interpretation, but has merely be chosen as to keep the proof of Theorem 1 as simple as possible. However, in Section 4, a topological structure on the event space is given, which is based on intuitive properties and under which limit-agreeing to disagree also obtains.

The proof of Theorem 1 illustrates an epistemic situation in which limit knowledge of the agents' posterior beliefs holds concurrently with distinct actual posterior beliefs of the agents, in the particular case of the agents' posterior beliefs being nowhere common knowledge in the structure. Hence, the impression might arise that the relevance of limit knowledge for a possible disagreement only occurs in epistemic models in which common knowledge of the agents' posterior beliefs does not hold anywhere at all. However, it can be shown that the possibility of such a disagreement may also emerge in situations in which common knowledge of the agents' posteriors actually holds somewhere in the structure. Hence, disagreement induced by limit knowledge does not depend on the existence or non-existence of common knowledge of the agents' posteriors beliefs in the epistemic model. From an interpretative point of view, limit knowledge of the agents' posteriors might be consistent with a disagreement between the agents irrespective of whether these posteriors would have been publicly disclosed somewhere, or not. Also, the proof of Theorem 1 describes an interactive situation in which limit knowledge of distinct agents' posterior beliefs obtains simultaneously with differing actual posterior beliefs of the agents. Thus, the impression might arise that a possible disagreement could only be elicited by limit knowledge of already distinct posteriors. However, it can be shown that agents may actually disagree, while having limit knowledge of identical posterior beliefs. Countering the preceding two possible impressions, the following example indeed depicts an epistemic situation in which a disagreement on the agents' actual posterior beliefs is induced by limit knowledge of identical posterior beliefs, while common knowledge of these posteriors also holds somewhere in the structure.

Example 1

Consider the Aumann structure $\mathcal{A}=(\Omega,(\mathcal{I}_i)_{i\in I},p)$, where $\Omega=\{\omega_n : n\geq 0\}$, $I=\{Alice, Bob\}$, $I_{Alice}=\{\{\omega_0\},\ \{\omega_1\}\}\ \cup\ \{\{\omega_{2n},\ \omega_{2n+1}\}: n>0\}$, $I_{Bob}=\{\{\omega_0\}\}\ \cup\ \{\{\omega_{2n+1},\omega_{2n+2}\}: n\geq 0\}$, and $p:\Omega\rightarrow\mathbb{R}$ is given by $p(\omega_n)=\frac{1}{2^{n+1}}$ for all $n\geq 0$. This structure is illustrated in Figure 3. Note that the common prior belief function p is well defined since $\sum_{n\geq 0}\frac{1}{2^{n+1}}=1$. Now, consider the event $E=\{\omega_1\}$ and the world ω_0. Besides, for sake of notational convenience, let the event $\bigcap_{i\in I}\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_i(\omega'))=p(E\,|\,\mathcal{I}_i(\omega_0))\}$ be denoted by E'. First of all, observe that $p(E\,|\,\mathcal{I}_{Alice}(\omega_0))=0$ as well as $p(E\,|\,\mathcal{I}_{Bob}(\omega_0))=0$. Moreover, $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Alice}(\omega'))=p(E\,|\,\mathcal{I}_{Alice}(\omega_0))=0\}=\Omega\setminus\{\omega_1\}$ and $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Bob}(\omega'))=p(E\,|\,\mathcal{I}_{Bob}(\omega_0))=0\}=\Omega\setminus\{\omega_1,\omega_2\}$, whence $E'=(\Omega\setminus\{\omega_1\})\cap(\Omega\setminus\{\omega_1,\omega_2\})=\Omega\setminus\{\omega_1,\omega_2\}$. Furthermore, the definitions of the possibility partitions of *Alice* and *Bob* ensure that $K^m(E')=K^m(\Omega\setminus\{\omega_1,\omega_2\})=\Omega\setminus\{\omega_1,\omega_2,\ldots,\omega_{m+2}\}$, for all $m>0$. Consequently, the sequence $(K^m(E'))_{m>0}$ is strictly shrinking. Now, consider the topology $\mathcal{T}$ on $\mathcal{P}(\Omega)$ defined by $\mathcal{T}=\{O\subseteq\mathcal{P}(\Omega):\{\omega_1,\omega_2\}\notin O\}\cup\{\mathcal{P}(\Omega)\}$. Then, the only open neighbourhood of the event $\{\omega_1,\omega_2\}$

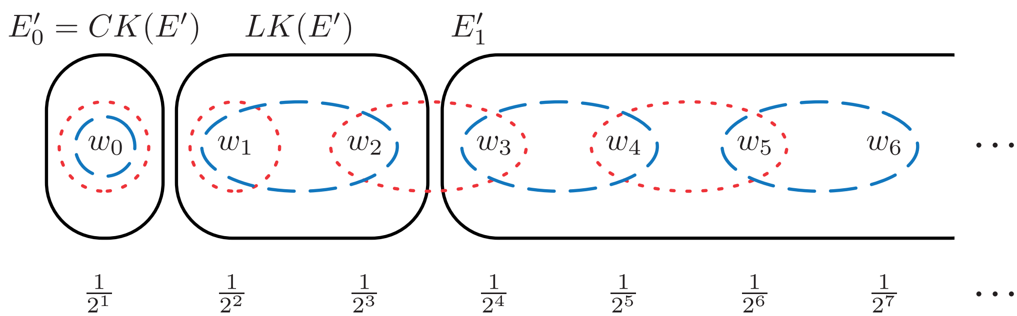

FIGURE 3. Illustration of the Aumann structure used in Example 1. The dotted and dashed sets represent the possibility partitions of *Alice* and *Bob*, respectively. The fractions correspond to the prior probabilities associated with the possible worlds. Note that $E' = E'_0 \cup E'_1$.

is $\mathcal{P}(\Omega)$, and all terms of the sequence $(K^m(E'))_{m>0}$ are contained in $\mathcal{P}(\Omega)$. Thus $(K^m(E'))_{m>0}$ converges to $\{\omega_1, \omega_2\}$. Moreover, for every event $F \in \mathcal{P}(\Omega)$ such that $F \neq \{\omega_1, \omega_2\}$, the singleton $\{F\}$ is open, and since $K^{m+1}(E') \subsetneq K^m(E')$ for all $m > 0$, the strictly shrinking sequence $(K^m(E'))_{m>0}$ will never remain in the open neighbourhood $\{F\}$ of F from some index onwards. Hence $(K^m(E'))_{m>0}$ does not converge to any such event F. Therefore, the limit point $\{\omega_1, \omega_2\}$ of the strictly shrinking sequence $(K^m(E'))_{m>0}$ is unique, and $LK(E') = \lim_{m\to\infty} K^m(E') = \{\omega_1, \omega_2\}$. Besides, note that $CK(E') = \{\omega \in \Omega : (\bigwedge_{i\in I} \mathcal{I}_i)(\omega) \subseteq E'\} = \{\omega_0\}$. Next, consider the world ω_1. Observe that $\omega_1 \in LK(E')$ and $p(E \mid \mathcal{I}_{Alice}(\omega_1)) = 1 \neq \frac{2}{3} = p(E \mid \mathcal{I}_{Bob}(\omega_1))$. Thus, $CK(E') \neq \emptyset$, and world ω_1 satisfies both conditions $\omega_1 \in LK(E')$ as well as $p(E \mid \mathcal{I}_{Alice}(\omega_1)) \neq p(E \mid \mathcal{I}_{Bob}(\omega_1))$. ♣

Observe that the structure of the two agents' possibility partitions in the proof of Theorem 1 and Example 1 is similar to the structure of the partitions in Rubinstein's [22] electronic mail game. Indeed, the resemblance lies in the existence of an infinite chain-type pattern which consecutively links information cells of the two agents by a single world in the intersection of the respective two cells, and where each information cell only contains two worlds. In our framework, such a pattern ensures that the sequence of iterated mutual knowledge is strictly shrinking, which is a necessary condition for limit knowledge to differ from common knowledge.

More generally, observe that all possible ways of limit-agreeing to disagree can actually be classified into three mutually exclusive cases.

First of all, disagreement on the agents' actual posteriors may be induced by limit knowledge of already distinct posteriors, while common knowledge of these posteriors is empty. Such agreeing to disagree is illustrated by the interactive situation depicted in the proof of Theorem 1.

Secondly, disagreement on the agents' actual posteriors can be induced by limit knowledge of identical posteriors, while common knowledge of these posteriors is non-empty. Such agreeing to disagree is illustrated in Example 1.

Thirdly, disagreement on the agents' actual posteriors can also be induced by limit knowledge of identical agents' posteriors, while common knowledge of these posteriors is empty. To see this, consider the Aumann structure given in the proof of Theorem 1, the event $E = \{\omega_0\}$, the world ω_2, and the topology on the event space $\mathcal{T} = \{O \subseteq \mathcal{P}(\Omega) : \{\omega_0, \omega_1\} \notin O\} \cup \{\mathcal{P}(\Omega)\}$. It thus follows that $p(E \mid \mathcal{I}_{Alice}(\omega_2)) = p(E \mid \mathcal{I}_{Bob}(\omega_2)) = 0$ and $E' = \bigcap_{i\in I} \{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\omega_2))\} = \Omega \setminus \{\omega_0, \omega_1\}$. Then, $CK(E') = \emptyset$ and $LK(E') = \{\omega_0, \omega_1\}$, as well as both $\omega_1 \in LK(E')$ and $p(E \mid \mathcal{I}_{Alice}(\omega_1)) = \frac{2}{3} \neq 0 = p(E \mid \mathcal{I}_{Bob}(\omega_1))$.

The fourth possibility of a disagreement on the agents' actual posteriors based on limit knowledge of already distinct posteriors and with non-emptiness of common knowledge of these posteriors is excluded by Aumann's agreement theorem.

Besides, in the epistemic-topological situations described in the proof of Theorem 1 as well as in Example 1, limit knowledge is not factive, i.e. the relation $LK(E') \subseteq E'$ does not hold for the considered event E'. However, the possibility to limit-agree to disagree established in Theorem 1 does not directly follow from the non-factiveness of limit knowledge. We now show that agents can limit-agree to disagree with factive limit knowledge, and that in this case, the distinct actual posteriors are induced by limit knowledge of already distinct posteriors as well as common knowledge of these posteriors being empty.

Lemma 2

There exist an Aumann structure $\mathcal{A} = (\Omega, (\mathcal{I}_i)_{i \in I}, p)$ equipped with a topology $\mathcal{T}$ on the event space $\mathcal{P}(\Omega)$, an event $E \subseteq \Omega$, and worlds $\omega, \hat{\omega} \in \Omega$ such that $LK(\bigcap_{i \in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}) \subseteq \bigcap_{i \in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}$, $\omega \in LK(\bigcap_{i \in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\})$, as well as $p(E \mid \mathcal{I}_i(\omega)) \neq p(E \mid \mathcal{I}_j(\omega))$ for some agents $i, j \in I$. In this case, $CK(\bigcap_{i \in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}) = \emptyset$ and $p(E \mid \mathcal{I}_i(\hat{\omega})) \neq p(E \mid \mathcal{I}_j(\hat{\omega}))$ for some agents $i, j \in I$.

Proof. First, the existence of an Aumann structure and a topology satisfying the postulated properties is established. Consider the Aumann structure $\mathcal{A}$ and the events E and E' given in the proof of Theorem 1. Let the topology $\mathcal{T}$ on the event space $\mathcal{P}(\Omega)$ be defined by $\mathcal{T} = \{O \subseteq \mathcal{P}(\Omega) : \{\omega_2\} \notin O\} \cup \{\mathcal{P}(\Omega)\}$. It follows that $LK(E') = \lim_{m \to \infty} K^m(E') = \{\omega_2\} \subseteq \Omega \setminus \{\omega_0, \omega_1\} = E'$, and $p(E \mid \mathcal{I}_{Alice}(\omega_2)) = \frac{2}{3} \neq \frac{1}{3} = p(E \mid \mathcal{I}_{Bob}(\omega_2))$. Taking $\omega = \hat{\omega} = \omega_2$ concludes the first part of the proof. Next, it is shown that if an Aumann structure and a topology satisfy the postulated properties, then the corresponding posteriors are distinct and common knowledge of these posteriors is thus empty. Consider some Aumann structure $\mathcal{A} = (\Omega, (\mathcal{I}_i)_{i \in I}, p)$, some topology $\mathcal{T}$ on the event space $\mathcal{P}(\Omega)$, some event E, and some worlds $\omega, \hat{\omega} \in \Omega$ satisfying the postulated conditions. Let the event $\bigcap_{i \in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}$ be denoted by E'. Since $\omega \in LK(E')$ and $LK(E') \subseteq E'$ both hold by the postulated conditions, it follows that $\omega \in E'$, i.e. $p(E \mid \mathcal{I}_i(\omega)) = p(E \mid \mathcal{I}_i(\hat{\omega}))$, for all $i \in I$. Moreover, as the postulated conditions ensure that $p(E \mid \mathcal{I}_i(\omega)) \neq p(E \mid \mathcal{I}_j(\omega))$ for some agents $i, j \in I$, it is the case that $p(E \mid \mathcal{I}_i(\hat{\omega})) \neq p(E \mid \mathcal{I}_j(\hat{\omega}))$ also obtains for some agents $i, j \in I$. Now, the contraposition of Aumann's Agreement Theorem (Version 1) directly implies that $CK(E') = \emptyset$. ■

Finally, agreeing to disagree with limit knowledge can be graphically illustrated for the case of two agents. A particular interactive situation is represented in Figure 4. Again, as in Figure 1, the event space is partitioned in equivalence classes of worlds inducing a same posterior belief profile in some underlying event E. Hence, the event E' denotes the equivalence class of worlds inducing the same posterior beliefs for all agents as at the auxiliary world $\hat{\omega}$. In the considered interactive situation, the event $CK(E')$ is non-empty, and the topology on the event space implies that the event $LK(E')$ is well-defined, distinct from $CK(E')$, and not even included in the event E' itself. Since $CK(E')$ is non-empty, Aumann's agreement theorem ensures that the agents' posterior beliefs induced by the auxiliary world $\hat{\omega}$ coincide, and thus any world in the equivalence class E' also induces identical posterior beliefs for all agents. Therefore, the posterior beliefs that are both common knowledge as well as limit knowledge are identical for all agents. Moreover, note that, as ω_1 lies in the equivalence class E', the agents' posterior beliefs at world ω_1 are the same as the ones induced by the auxiliary world $\hat{\omega}$, and thus all identical. Consequently, since ω_1 is also contained in both $CK(E')$ as well as in $LK(E')$, agents do agree on identical posteriors at ω_1, both with common knowledge and with limit knowledge. Besides, the position of the world ω_2 relative to ω_1 in Figure 4 ensures that *Alice* and

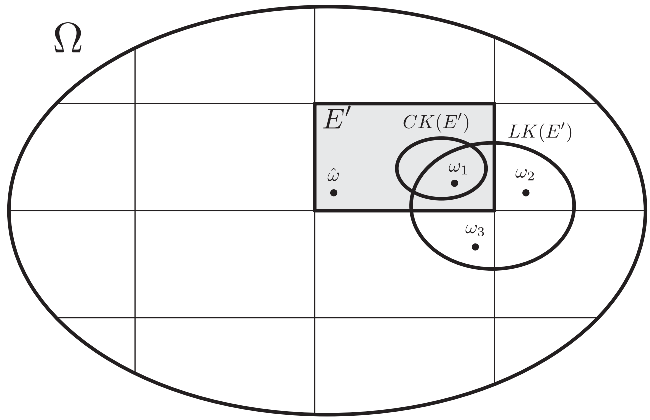

FIGURE 4. Illustration of limit-agreeing to disagree for the case of two agents.

Bob hold distinct posterior beliefs at ω_2. Similarly, the agents have different posterior beliefs at ω_3. Since ω_2 and ω_3 are contained in $LK(E')$, agents actually agree to disagree with limit knowledge at both worlds. Observe that, although being based on limit knowledge of the same posterior beliefs, the disagreements at ω_2 and ω_3 differ. Furthermore, both such disagreements are in fact induced by limit knowledge of equal posteriors, as the posteriors of *Alice* and *Bob* coincide throughout E'.

4 A representative example

The extension of the standard set-based approach to interactive epistemology with a topological dimension has been shown to enable the possibility for agents to limit-agree to disagree on their posterior beliefs. The question then arises whether limit-agreeing to disagree remains possible in interactive situations, where the topologies are based on intuitive properties. A topology describing a specific agents' perception of the event space is now presented. Accordingly, lower iterated mutual knowledge up to some finite level is grasped by the agents in a more refined manner than higher iterations from that level onwards. Such a property seems natural, since real world agents typically only have distinguished cognitive access to iterated knowledge claims up to some finite level, in contrast to the idealized agents that can equally well conceive of any layer of the complete hierarchy of interactive knowledge. For this intuitive topology it is then shown that agreeing to disagree with limit knowledge is possible. Besides, this intuitive topology establishes that limit knowledge is identical to iterated mutual knowledge for some finite level m, i.e. it is equal to almost common knowledge, a concept due to Rubinstein [22]. The example could thus also be seen as a contribution to the literature of bounded reasoning. However, it differs from models of k-level reasoning, which express a specific and different kind of finite level reasoning in the particular context of games.[10]

Towards this purpose, consider an Aumann structure $\mathcal{A} = (\Omega,\ (\mathcal{I}_i)_{i\in I},\ p)$ and an event E. Furthermore, for any world $\omega \in \Omega$, let E'_ω denote the event consisting of all worlds that induce the same posterior beliefs in E for all agents as at ω, i.e. $E'_\omega = \bigcap_{i\in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\omega))\}$.

[10]Loosely speaking, k-level reasoning restricts a belief hierarchy in the particular context of games such that random play is assumed at the first level, and only best responses to the respective preceding levels are admitted at the iterated levels up to some finite level k. In fact, Crawford et al. [12] provide a recent overview on the literature of k-level reasoning and other theories of finite reasoning.

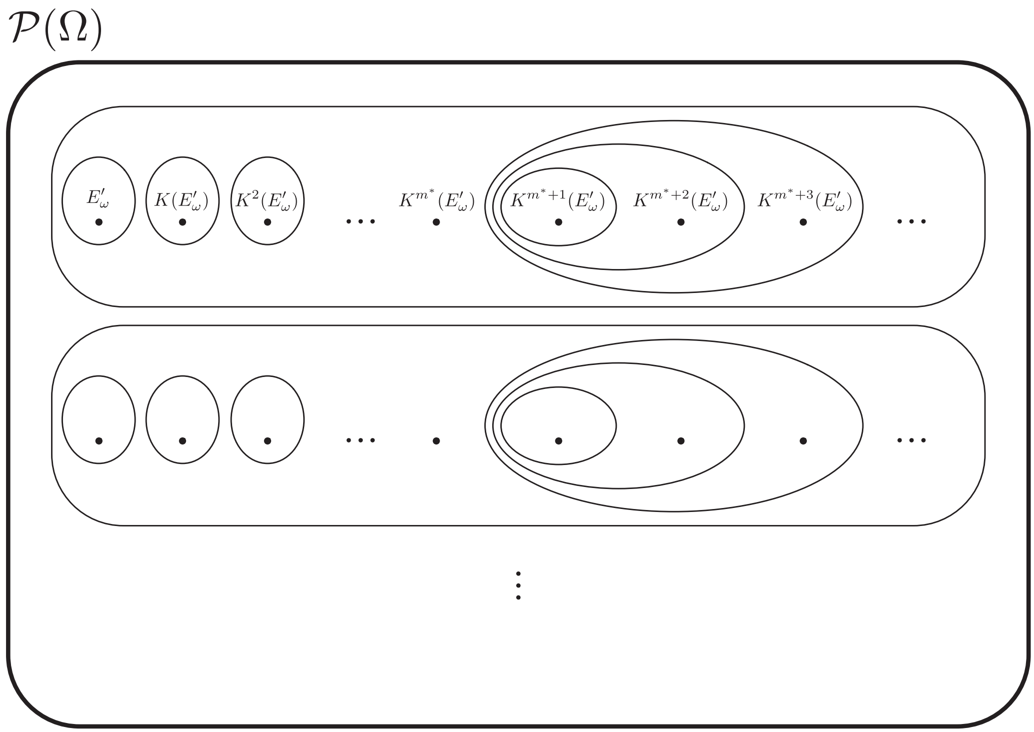

FIGURE 5. Illustration of the topology $\mathcal{T}_{E,m^*}$.

Note that constancy of the agents' posterior beliefs in E yields an equivalence relation on the set of possible worlds, and hence every E'_ω represents an equivalence class of worlds. Consequently, the collection $\mathcal{C} = \{E'_\omega : \omega \in \Omega\}$ of all equivalence classes of worlds that induce a same posterior belief profile forms a partition of Ω. Given some event E and some index $m^* > 0$, the epistemically-based topology $\mathcal{T}_{E,m^*}$ is defined as the topology on the event space $\mathcal{P}(\Omega)$ generated by the subbase

$$\begin{aligned}
&\{\{K^m(E'_\omega) : m \geq 0\} : \omega \in \Omega\} \\
\cup\ &\{\mathcal{P}(\Omega) \setminus \{K^m(E'_\omega) : m \geq 0 \text{ and } \omega \in \Omega\}\} \\
\cup\ &\{\{K^m(E'_\omega)\} : 0 \leq m < m^* \text{ and } \omega \in \Omega\} \\
\cup\ &\{\{K^{m^*+j}(E'_\omega) : 0 < j \leq n\} : n > 0 \text{ and } \omega \in \Omega\}.
\end{aligned}$$

The topology $\mathcal{T}_{E,m^*}$ is illustrated in Figure 5, where the infinite sequence $(K^m(E'_\omega))_{m \geq 0}$ is represented by a horizontal sequence of points for each $\omega \in \Omega$, and open sets of the subbase by circle-type shapes around these points. In this topology, the closeness relation between events is represented by means of the T_0 and T_2 separation properties.[11]

The topology $\mathcal{T}_{E,m^*}$ reveals a specific agent perception of the event space, according to which the agents express a more refined distinction between the m^* first iterated mutual knowledge of their

[11] Given a topological space $(A, \mathcal{T})$, two points in A are called T_2-*separable*, if there exist two disjoint $\mathcal{T}$-open neighbourhoods of these two points. Moreover, two points in A are called T_0-*separable*, if there exists a $\mathcal{T}$-open set containing precisely one of these two points. Note that T_2-separability implies T_0-separability. In fact, two events that are T_0-separable but not T_2-separable can be said to be closer to each other than two events that are both T_0-separable as well as T_2-separable.

posterior beliefs in E than between the remaining ones. This specific perception is formally reflected by two separation properties satisfied by the topology $\mathcal{T}_{E,m^*}$.

First, given two events X and Y, if X and Y are two distinct terms of the same sequence $(K^m(E'_\omega))_{m>0}$, such that $X = K^{m_1}(E'_\omega)$ and $Y = K^{m_2}(E'_\omega)$, for some $\omega \in \Omega$ and $m_1, m_2 < m^*$, then X and Y are T_2-separable, and therefore also T_0-separable. Secondly, if X and Y are two different elements of the same sequence $(K^m(E'_\omega))_{m>0}$, such that $X = K^{m_1}(E'_\omega)$ and $Y = K^{m_2}(E'_\omega)$, for some $\omega \in \Omega$ and $m_1, m_2 > m^*$, then X and Y are T_0-separable, yet not T_2-separable. According to these two separation properties, agents have access to a more refined distinction between the m^* first iterated knowledge claims of their posterior beliefs in E than between the iterated mutual knowledge claims of order strictly larger than m^*. In other words, iterated mutual knowledge claims are only precisely discerned up to a given amount of iterations, and thereafter the higher iterations become less distinguishable for the agents. Also, from a bounded rationality point of view, the agents' perception of higher-order mutual knowledge due to the topology $\mathcal{T}_{E,m^*}$ reflects that people typically lose track from some iteration level onwards when reasoning about higher-order mutual knowledge.

Furthermore, the topology $\mathcal{T}_{E,m^*}$ notably satisfies the following epistemic-topological property: for any event E'_ω, if the sequence $(K^m(E'_\omega))_{m>0}$ is strictly shrinking, then $LK(E'_\omega) = K^{m^*}(E'_\omega)$. Indeed, suppose that the sequence $(K^m(E'_\omega))_{m>0}$ is strictly shrinking. Then, by definition of $\mathcal{T}_{E,m^*}$, the only open neighbourhoods of $K^{m^*}(E'_\omega)$ are $\mathcal{P}(\Omega)$ and $\{K^m(E'_\omega) : m \geq 0\}$. Since both sets contain all terms of the sequence $(K^m(E'_\omega))_{m>0}$, it follows that $K^{m^*}(E'_\omega)$ is a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. Moreover, it can be shown that this limit point is actually unique.[12] Therefore, $LK(E'_\omega) = \lim_{m\to\infty} K^m(E'_\omega) = K^{m^*}(E'_\omega)$. Furthermore, since the sequence $(K^m(E'_\omega))_{m>0}$ is strictly shrinking, $CK(E'_\omega) = \bigcap_{m>0} K^m(E'_\omega) \subsetneq K^{m^*}(E'_\omega)$, and hence $CK(E'_\omega) \neq LK(E'_\omega)$.

Note that if the event space is equipped with the topology $\mathcal{T}_{E,m^*}$, the epistemic-topological event $LK(E')$ actually coincides with Rubinstein's [22] notion of almost common knowledge if m^* is large, i.e. $LK(E') = K^{m^*}(E')$. Thus, agents with topological mental states according to $\mathcal{T}_{E,m^*}$ who have limit knowledge actually reason in line with almost common knowledge. More precisely, if agents can only accurately conceive of higher-order interactive knowledge up to some fixed level, then they reason with limit knowledge if and only if they reason with almost common knowledge up to that level only. In fact, the connection between *cognitive-topologically* perceiving iterated knowledge distinctly only up to some finite level and *epistemically* reasoning in line with almost common knowledge up to that level does seem natural. Thus, the topology $\mathcal{T}_{E,m^*}$ provides an intuitive topological foundation for Rubinstein's almost common knowledge in terms of the agents' cognitive perception of the event space.

Finally, the following example describes an interactive situation, in which the intutive topology $\mathcal{T}_{E,m^*}$ provides a possibility for the agents to agree to disagree on their posterior beliefs with limit knowledge.

Example 2

Consider the Aumann structure $\mathcal{A} = (\Omega, (\mathcal{I}_i)_{i\in I}, p)$, where $\Omega = \{\omega_n : n \geq 0\}$, $I = \{Alice, Bob\}$, $I_{Alice} = \{\{\omega_0\}, \{\omega_1, \omega_2\}, \{\omega_3, \omega_4, \omega_5, \omega_6\}, \{\omega_7, \omega_8, \omega_9\}\} \cup \{\{\omega_{2n}, \omega_{2n+1}\} : n \geq 5\}$, $I_{Bob} = \{\{\omega_0, \omega_1, \omega_2, \omega_3, \omega_4\}, \{\omega_5, \omega_6, \omega_7, \omega_8\}\} \cup \{\{\omega_{2n+1}, \omega_{2n+2}\} : n \geq 4\}$, and $p : \Omega \to \mathbb{R}$ is given by $p(\omega_n) = \frac{1}{2^{n+1}}$ for all $n \geq 0$. Also, consider the event $E = \{\omega_1, \omega_5\} \cup \{\omega_{2n} : n \geq 1\}$ and the world ω_{10}. Besides, for sake of notational convenience, let the event $\bigcap_{i\in I}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\omega_{10}))\}$ be denoted by E'. First of all, observe that the computation of the posterior beliefs of *Alice* and *Bob* gives a variety of distinct values for the first ten worlds $\{\omega_0, \omega_1, \ldots, \omega_9\}$, as well as $p(E \mid \mathcal{I}_{Alice}(\omega_n)) = \frac{2}{3}$

[12] A detailed proof of this fact is provided in Appendix.

and $p(E\,|\,\mathcal{I}_{Bob}(\omega_n))=\frac{1}{3}$, for all $n\geq 10$. It follows that $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Alice}(\omega'))=p(E\,|\,\mathcal{I}_{Alice}(\omega_{10}))\}=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_9\}$ and $\{\omega'\in\Omega : p(E\,|\,\mathcal{I}_{Bob}(\omega'))=p(E\,|\,\mathcal{I}_{Bob}(\omega_{10}))\}=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_8\}$, thus $E'=(\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_9\})\cap(\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_8\})=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_9\}$. Moreover, the definitions of the possibility partitions of *Alice* and *Bob* ensure that $K^m(E')=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_{m+9}\}$, for all $m>0$. Consequently, the sequence $(K^m(E'))_{m>0}$ is strictly shrinking and $CK(E')=\bigcap_{m>0}K^m(E')=\emptyset$. Now, let $m^*>0$ be some index and suppose that $\mathcal{P}(\Omega)$ is equipped with the topology $\mathcal{T}_{E,m^*}$. Since the sequence $(K^m(E'))_{m>0}$ is strictly shrinking, the definition of this topology ensures that $LK(E')=K^{m^*}(E')=\Omega\setminus\{\omega_0,\omega_1,\ldots,\omega_{m^*+9}\}$. Consequently, the computations of the posterior beliefs of *Alice* and *Bob* give $p(E\,|\,\mathcal{I}_{Alice}(\omega))=\frac{2}{3}$ and $p(E\,|\,\mathcal{I}_{Bob}(\omega))=\frac{1}{3}$, for all $\omega\in LK(E')$. In other words, for all $\omega\in LK(E')$, it holds that $p(E\,|\,\mathcal{I}_{Alice}(\omega))\neq p(E\,|\,\mathcal{I}_{Bob}(\omega))$. ♣

In the preceding example, a situation is provided where agents with an intuitive perception of the event space do agree to disagree. Yet, as limit knowledge coincides with Rubinstein's almost common knowledge, the example also shows that Aumann's agreement theorem is not robust in the sense that agents having almost common knowledge of posteriors can hold distinct posterior beliefs.

5 Conclusion

In a topologically extended epistemic model, agents have been shown to be able to limit-agree to disagree. More precisely, if Bayesian agents have a common prior belief as well as limit knowledge of their posteriors beliefs, then their actual posterior beliefs may indeed differ. Our possibility result contrasts with Aumann's impossibility theorem. Actually, limit-agreeing to disagree is possible in interactive situations enriched by some topology that reveals a specific cogent agent perception of the event space. Indeed, our representative example, in which limit knowledge coincides with Rubinstein's [22] almost common knowledge, illustrates the non-robustness of Aumann's agreement theorem in the sense that the posteriors of agents holding approximate common knowledge of posteriors may actually differ.

The possibility of agreeing to disagree with limit knowledge — as finitely iterated mutual knowledge — in a topologically enriched epistemic structure can be seen in the context of the growing literature on k-level reasoning and other theories on bounded reasoning. It would be intriguing for future work to investigate such models of finite thinking from an epistemic-topological point of view. An interesting recent point of departure could be Kets' [16] and [17] theory of finite depth reasoning. Accordingly, the language of game-playing agents is restricted such that they can only reason about higher-order beliefs up to some level k. This is related to our intuitive topology in Section 4, which also restricts agents' reasoning to finitely iterated mutual knowledge. Even though any higher-order mutual knowledge is — in contrast to Kets' models — part of the agents' language, they cannot conceive of it in a precise but only in a 'blurred' way. Hence, a relevant question would be what topological conditions need to generally be invoked on standard epistemic structures for games, such that the players' reasoning remains restricted as in Kets' style type spaces.

Besides, note that it is impossible for agents to limit-agree to disagree in the case of finite Aumann structures as well as in the case of infinite Aumann structures where the sequence of iterated mutual knowledge is not strictly shrinking. Indeed, as shown in Bach and Cabessa [8], in such cases, the epistemic-topological operator limit knowledge necessarily coincides with the purely epistemic operator common knowledge, and consequently, Aumann's impossibility result does apply.

The epistemic-topological operator limit knowledge is based on the notion of set proximity in contrast to the purely epistemic operator common knowledge which is based on set inclusion. Indeed, while the latter notion corresponds to logical implication, the former does not comply with any purely logical concept, but relates to a whole variety of possible cognitive perceptions induced by the respective topology under consideration. Depending on the underlying topology, a given notion of set proximity between events may reflect physical properties of these events, but also mental representations on the event space such as in our representative example in Section 4. Therefore, as opposed to common knowledge which only captures a single epistemic phenomenon, limit knowledge is able to represent a variety of possible epistemic-topological phenomena as a function on the particular topology on the event space. In this sense, limit knowledge can be viewed as some kind of generalized epistemic-topological concept which, for every possible underlying topology, becomes an operator with precise meaning.

Actually, the topological approach to set-based interactive epistemology, in which topologies model agents' perceptions of closeness between events, can be used to describe various agent reasoning patterns that do not only depend on mere epistemic but also on topological features of the underlying interactive situation.

In general, the epistemic-topological approach is of descriptive and not normative character. Due to the heterogeneity of real-world agents and similar to the variety of solution concepts in games, it appears plausible to study distinct topological perception patterns of the event space and their implications for reasoning via the epistemic-topological operator limit knowledge.

Moreover, we envision the construction of a more general epistemic-topological framework — *topological Aumann structures* — comprising topologies not only on the event space but also on the state space. Such an extension permits an explicit consideration of a notion of closeness between events and between worlds, enabling us to model common agent perceptions of the event and state spaces as well as their interconnection.[13] In particular, it might be of distinguished interest to base topologies on first principles such as epistemic axioms or natural closeness properties. In line with this perspective, the topology provided in Section 4 reflects the natural agent perception for which iterated mutual knowledge becomes imprecise from some level onwards.

Besides, in order to model subjective rather than common agent perceptions of the event and state spaces, the epistemic-topological framework envisioned here could be amended by assigning specific and potentially distinct topologies to every agent. A collective topology reflecting a common closeness perception could then be constructed on the basis of the particular agent topologies, and limit knowledge be defined in such a global topological context. For instance, by providing a topology that is coarser than each agent's one, the meet topology could be used as a representative collective topology. Alternatively, an agent-specific operator limit knowledge could be defined with respect to each particular topology, and mutual limit knowledge as their intersection then be considered.

Finally, in a general epistemic-topological framework, various issues can be addressed. For example, the possibility of agents' agreeing to disagree with limit knowledge can be further analysed for other epistemically-based as well as agent specific topologies. Furthermore, analogously to the epistemic program in game theory that attempts to provide epistemic foundations for solution concepts, an epistemic-topological approach could generate epistemic-topological foundations for solution concepts. In addition, it could be attempted to develop a theory of counterfactuals in set-based interactive epistemology founded on a notion of similarity of worlds or events provided by topologies on the state or event space, respectively.

[13] Note that similar considerations also arise in epistemic logical frameworks such as in Dégremont and Roy [13]. Since the plausibility orderings in their framework could not only be defined on the states but also on the propositions, which are events from a semantic point of view, it could be of interest to analyse different — intuitive — ways of deriving plausibility orderings on propositions from the plausibility orderings on the states, or to more generally impose intuitive criteria on such orderings.

Acknowledgements

We would like to thank Richard Bradley, Adam Brandenburger, Jacques Duparc, Andrés Perea, Elias Tsakas, as well as the workshop participants at the 3rd International Workshop on Logic, Rationality and Interaction (LORI-III). Besides, valuable remarks by two anonymous referees at the *Journal of Logic and Computation* have been highly appreciated.

References

[1] R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, **4**, 1236–1239, 1976.
[2] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, **55**, 1–18, 1987.
[3] R. J. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, **8**, 6–19, 1995.
[4] R. J. Aumann. Interactive epistemology I: knowledge. *International Journal of Game Theory*, **28**, 263–300, 1999a.
[5] R. J. Aumann. Interactive epistemology II: probability. *International Journal of Game Theory*, **28**, 301–314, 1999b.
[6] R. J. Aumann. Musings on information and knowledge. *Econ Journal Watch*, **2**, 88–96, 2005.
[7] C. W. Bach and J. Cabessa. Limit knowledge of rationality. In *TARK 2009: Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 34–40. ACM, 2009.
[8] C. W. Bach and J. Cabessa. Common knowledge and limit knowledge. *Theory and Decision*, **73**, 423–440, 2012.
[9] C. W. Bach and A. Perea. Agreeing to disagree with lexicographic prior beliefs. *Mathematical Social Sciences*, **66**, 129–133, 2013.
[10] M. Bacharach. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, **37**, 167–190, 1985.
[11] G. Bonanno and K. Nehring. Agreeing to disagree: a survey. Mimeo, University of California, 1997.
[12] V. P. Crawford, M. A. Costa-Gomes and N. Iriberri. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. *Journal of Economic Literature*, **51**, 5–62, 2013.
[13] C. Dégremont and O. Roy. Agreement theorems in dynamic-epistemic logic. *Journal of Philosophical Logic*, **41**, 735–764, 2012.
[14] L. Demey. Agreeing to disagree in probabilistic dynamic epistemic logic. *Synthese*, **191**, 409–438, 2014.
[15] J. Geanakoplos and M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, **26**, 363–390, 1982.
[16] W. Kets. Finite depth of reasoning and equilibrium play in games with incomplete information. Working Paper, 2014a.
[17] W. Kets. Bounded reasoning and higher-order uncertainty. Working Paper, 2014b.
[18] D. Lewis. *Counterfactuals*. Blackwell, 1973.
[19] Ménager L. Agreeing to disagree: a review. Working Paper, 2012.
[20] P. Milgrom and N. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, **26**, 17–27, 1982.
[21] D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, **1**, 170–190, 1989.

[22] A. Rubinstein. The electronic mail game: strategic behavior under "almost common knowledge". *The American Economic Review*, **79**, 385–391, 1989.
[23] D. Samet. Ignoring ignorance and agreeing to disagree. *Journal of Economic Theory*, **52**, 190–207, 1990.
[24] D. Stalnaker. A theory of conditionals. *Studies in Logical Theory, American Philosophical Quarterly*, **2**, 98–112, 1968.

Received 16 June 2015

Appendix

Proof of Lemma 1. Firstly, we show that $CK^{meet}(E) \subseteq CK^{inter}(E)$. By definition, $CK^{meet}(E)$ is the union of the cells of the meet $\bigwedge_{i\in I}\mathcal{I}_i$ that are included in E. Since the meet $\bigwedge_{i\in I}\mathcal{I}_i$ is coarser than each agent's possibility partition, the event $CK^{meet}(E)$ can thus be written as a union of information cells for each agent $i\in I$, i.e. for all $i\in I$, there exists a set $F_i\subseteq\Omega$ such that $CK^{meet}(E)=\bigcup_{\omega'\in F_i}\mathcal{I}_i(\omega')$. It follows that $K_i(CK^{meet}(E)) = K_i(\bigcup_{\omega'\in F_i}\mathcal{I}_i(\omega')) = \bigcup_{\omega'\in F_i}\mathcal{I}_i(\omega') = CK^{meet}(E)$, for all $i\in I$. Hence, $K(CK^{meet}(E)) = \bigcap_{i\in I} K_i(CK^{meet}(E)) = \bigcap_{i\in I} CK^{meet}(E) = CK^{meet}(E)$. It follows by induction that $K^m\ (CK^{meet}\ (E)) = CK^{meet}\ (E)$, for all $m>0$. Hence, $CK^{inter}\ (CK^{meet}(E)) = \bigcap_{m>0} K^m(CK^{meet}(E)) = \bigcap_{m>0} CK^{meet}\ (E) = CK^{meet}\ (E)$, which shows that $CK^{meet}(E)$ is a fixed point of the CK^{inter} operator. Finally, since the operator CK^{inter} is monotone with respect to set inclusion and since $CK^{meet}(E)\subseteq E$, it follows that $CK^{meet}(E)=CK^{inter}(CK^{meet}(E))\subseteq CK^{inter}(E)$. Secondly, we show that $CK^{inter}(E)\subseteq CK^{meet}(E)$. As a preliminary claim, we prove by induction on $n\in\mathbb{N}$ that, for any $\omega,\omega^*\in\Omega$, if both $\omega^*\notin E$ and there exists a sequence of $n+1$ possibility cells $(\mathcal{I}_k)_{k=0}^n$ such that $\omega\in\mathcal{I}_0$, $\omega^*\in\mathcal{I}_n$, and $\mathcal{I}_k\cap\mathcal{I}_{k+1}\neq\emptyset$ for all $0\leq k<n$, then $\omega\notin K^{n+1}(E)$. First of all, for the case $n=0$, let $\omega,\omega^*\in\Omega$ and $\mathcal{I}_0$ be a possibility cell such that $\omega,\omega^*\in\mathcal{I}_0$, and $\omega^*\notin E$. Since $\omega\in\mathcal{I}_0$, the cell $\mathcal{I}_0$ can be written as $\mathcal{I}_{i_0}(\omega)$, where $i_0\in I$ denotes the agent to whom the cell $\mathcal{I}_0$ belongs. Also, $\omega^*\in\mathcal{I}_0=\mathcal{I}_{i_0}(\omega)$ and $\omega^*\notin E$ imply that $\mathcal{I}_{i_0}(\omega)\not\subseteq E$. It follows that $\omega\notin K_{i_0}(E)$ and hence $\omega\notin\bigcap_{i\in I}K_i(E)=K(E)=K^1(E)$. Now, assume that the claim holds true for some $n=m$. Let $\omega,\omega^*\in\Omega$ and let $(\mathcal{I}_k)_{k=0}^{m+1}$ be a sequence of $m+2$ possibility cells such that $\omega^*\notin E$, $\omega\in\mathcal{I}_0$, $\omega^*\in\mathcal{I}_{m+1}$, and $\mathcal{I}_k\cap\mathcal{I}_{k+1}\neq\emptyset$ for all $0\leq k<m+1$. Since $\mathcal{I}_0\cap\mathcal{I}_1\neq\emptyset$, there exists some world $\omega'\in\mathcal{I}_0\cap\mathcal{I}_1$. Now consider the sequence of $m+1$ possibility cells $(\mathcal{J}_k)_{k=0}^m$ defined by $\mathcal{J}_k=\mathcal{I}_{k+1}$ for all $0\leq k\leq m$. This sequence satisfies $\omega'\in\mathcal{J}_0$, $\omega^*\in\mathcal{J}_m$, and $\mathcal{J}_k\cap\mathcal{J}_{k+1}\neq\emptyset$ for all $0\leq k<m$. By the induction hypothesis, it follows that $\omega'\notin K^{m+1}(E)$. Moreover, since $\omega'\in\mathcal{I}_0$, the cell $\mathcal{I}_0$ can be written as $\mathcal{I}_{i_0}(\omega')$, where $i_0\in I$ denotes the agent to whom the cell $\mathcal{I}_0$ belongs. Hence, $\omega'\in\mathcal{I}_0=\mathcal{I}_{i_0}(\omega)$ and $\omega'\notin K^{m+1}(E)$ imply that $\mathcal{I}_{i_0}(\omega)\not\subseteq K^{m+1}(E)$. It follows that $\omega\notin K_{i_0}(K^{m+1}(E))$ and hence $\omega\notin\bigcap_{i\in I}K_i(K^{m+1}(E))=K(K^{m+1}(E))=K^{m+2}(E)$, which completes the proof of the preliminary claim. Now, let $\omega\notin CK^{meet}(E)$. By the definition of $CK^{meet}(E)$, it holds that $(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\not\subseteq E$. Since the meet $\bigwedge_{i\in I}\mathcal{I}_i$ is the finest partition coarser than all agents' possibility partitions, the cell $(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$ can be written as a union of consecutively intersecting distinct cells, i.e. there exists an index set $K\subseteq\mathbb{N}$ and a sequence of possibility cells $(\mathcal{I}_k)_{k\in K}$ such that $(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)=\bigcup_{k\in K}\mathcal{I}_k$, as well as $\omega\in\mathcal{I}_0$ and $\mathcal{I}_{k-1}\cap\mathcal{I}_k\neq\emptyset$, for all $k\in K\setminus\{0\}$. Note that any two consecutive terms of the sequence $(\mathcal{I}_k)_{k\in K}$ actually belong to distinct agents, since $\mathcal{I}_{k-1}\cap\mathcal{I}_k\neq\emptyset$ ensures that $\mathcal{I}_{k-1}$ and $\mathcal{I}_k$ are not in a same agent's possibility partition. As $(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\not\subseteq E$, there exists a world ω^* such that $\omega^*\in(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$ and $\omega^*\notin E$. Moreover, since $\omega^*\in(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)=\bigcup_{k\in K}\mathcal{I}_k$, there exists an index $l\in K$ such that $\omega^*\in\mathcal{I}_l$. Therefore, $\omega^*\notin E$, and the sequence of $l+1$ possibility cells $(\mathcal{I}_k)_{k=0}^l$ satisfies $\omega\in\mathcal{I}_0$, $\omega^*\in\mathcal{I}_l$, as well as $\mathcal{I}_k\cap\mathcal{I}_{k+1}\neq\emptyset$, for all $0\leq k<l$. By the preliminary claim, it thus follows that $\omega\notin K^{l+1}(E)$. Hence, $\omega\notin\bigcap_{m>0}K^m(E)=CK^{inter}(E)$, which concludes the proof. ■

PROOF OF AUMANN'S AGREEMENT THEOREM (VERSION 1). Let $\hat{\omega} \in \Omega$ be such that $CK(\bigcap_{i=1}^{n}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}) \neq \emptyset$, and for sake of notational convenience let the event $\bigcap_{i=1}^{n}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}$ be denoted by E'. Consider a world $\omega \in CK(E')$ and some agent $i^* \in \{1,2,\ldots,n\}$. As the meet $\bigwedge_{i=1}^{n}\mathcal{I}_i$ is coarser than agent's i^* possibility partition, the cell $\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega)$ can be written as a disjoint union of information cells of i^*, namely there exists a set $A_{i^*} \subseteq \Omega$ such that $\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega) = \bigcup_{\omega' \in A_{i^*}}\mathcal{I}_{i^*}(\omega')$, and for all $\omega_1, \omega_2 \in A_{i^*}$, if $\omega_1 \neq \omega_2$, then $\mathcal{I}_{i^*}(\omega_1) \neq \mathcal{I}_{i^*}(\omega_2)$. Since $\omega \in CK(E')$, the meet definition of common knowledge ensures that $\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega) \subseteq E'$. By the definition of the event E' it then follows that $p(E \mid \mathcal{I}_{i^*}(\hat{\omega})) = p(E \mid \mathcal{I}_{i^*}(\omega'))$, for all $\omega' \in A_{i^*}$. Thus, $p(E \cap \mathcal{I}_{i^*}(\omega')) = p(E \mid \mathcal{I}_{i^*}(\hat{\omega})) \cdot p(\mathcal{I}_{i^*}(\omega'))$, for all $\omega' \in A_{i^*}$. Summing over the worlds in A_{i^*} yields the following equation of sums $\sum_{\omega' \in A_{i^*}} p(E \cap \mathcal{I}_{i^*}(\omega')) = p(E \mid \mathcal{I}_{i^*}(\hat{\omega})) \cdot \sum_{\omega' \in A_{i^*}} p(\mathcal{I}_{i^*}(\omega'))$. By countable additivity of the probability measure p, pairwise disjointness of the events $E \cap \mathcal{I}_{i^*}(\omega')$ for all $\omega' \in A_{i^*}$, and distributivity of intersection, it follows that $\sum_{\omega' \in A_{i^*}} p(E \cap \mathcal{I}_{i^*}(\omega')) = p(\bigcup_{\omega' \in A_{i^*}}(E \cap \mathcal{I}_{i^*}(\omega'))) = p(E \cap \bigcup_{\omega' \in A_{i^*}}\mathcal{I}_{i^*}(\omega')) = p(E \cap \bigwedge_{i=1}^{n}\mathcal{I}_i(\omega))$ and $\sum_{\omega' \in A_{i^*}} p(\mathcal{I}_{i^*}(\omega')) = p(\bigcup_{\omega' \in A_{i^*}}\mathcal{I}_{i^*}(\omega')) = p(\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega))$. The equation of sums can then be written as $p(E \cap \bigwedge_{i=1}^{n}\mathcal{I}_i(\omega)) = p(E \mid \mathcal{I}_{i^*}(\hat{\omega})) \cdot p(\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega))$, thence $p(E \mid \mathcal{I}_{i^*}(\hat{\omega})) = p(E \cap \bigwedge_{i=1}^{n}\mathcal{I}_i(\omega)) / p(\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega))$. Since i^* has been arbitrarily chosen, the latter equality holds for every agent $i \in \{1,2,\ldots,n\}$. Therefore, $p(E \mid \mathcal{I}_1(\hat{\omega})) = p(E \mid \mathcal{I}_2(\hat{\omega})) = \cdots = p(E \mid \mathcal{I}_n(\hat{\omega})) = p(E \cap \bigwedge_{i=1}^{n}\mathcal{I}_i(\omega)) / p(\bigwedge_{i=1}^{n}\mathcal{I}_i(\omega))$, which concludes the proof. ■

PROOF OF AUMANN'S AGREEMENT THEOREM (VERSION 2). Again, for sake of notational convenience, let the event $\bigcap_{i=1}^{n}\{\omega' \in \Omega : p(E \mid \mathcal{I}_i(\omega')) = p(E \mid \mathcal{I}_i(\hat{\omega}))\}$ be denoted by E'. Since $\omega \in CK(E')$, it holds that $\omega \in E'$. Observe that the definition of the event E' ensures that $p(E \mid \mathcal{I}_i(\omega)) = p(E \mid \mathcal{I}_i(\hat{\omega}))$ for all agents $i \in \{1,2,\ldots,n\}$. Also, Theorem 2 implies that $p(E \mid \mathcal{I}_1(\hat{\omega})) = p(E \mid \mathcal{I}_2(\hat{\omega})) = \cdots = p(E \mid \mathcal{I}_n(\hat{\omega}))$. Therefore, $p(E \mid \mathcal{I}_1(\omega)) = p(E \mid \mathcal{I}_2(\omega)) = \cdots = p(E \mid \mathcal{I}_n(\omega))$. ■

PROOF OF THE UNIQUENESS OF THE LIMIT POINT OF THE SEQUENCE $(K^m(E'_\omega))_{m>0}$. To see that this limit point is actually unique, consider $F \in \mathcal{P}(\Omega)$ such that $F \neq K^{m^*}(E'_\omega)$. Then either $F = K^m(E'_{\omega'})$ for some $m < m^*$ and some $\omega' \in \Omega$, or $F = K^m(E'_{\omega'})$ for some $m > m^*$ and some $\omega' \in \Omega$, or $F = K^{m^*}(E'_{\omega'})$ for some $\omega' \neq \omega$, or $F \neq K^m(E'_{\omega'})$ for all $m \geq 0$ and all $\omega' \in \Omega$. These four mutually exclusive cases are now considered in turn. First of all, if $F = K^m(E'_{\omega'})$ for some $m < m^*$ and some $\omega' \in \Omega$, then $\{K^m(E'_{\omega'})\}$ is an open neighbourhood of F. Since the sequence $(K^m(E'_\omega))_{m>0}$ is strictly shrinking, it can then not be the case that the singleton open neighbourhood $\{K^m(E'_{\omega'})\}$ of F contains all terms of the sequence $(K^m(E'_\omega))_{m>0}$ from some index onwards. Therefore F is not a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. Secondly, if $F = K^m(E'_{\omega'})$ for some $m > m^*$ and some $\omega' \in \Omega$, then $\{K^{m^*+j}(E'_{\omega'}) : 0 < j \leq m - m^*\}$ is an open neighbourhood of F. Since the set $\{K^{m^*+j}(E'_{\omega'}) : 0 < j \leq m - m^*\}$ is finite, F cannot be a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. Thirdly, if $F = K^{m^*}(E'_{\omega'})$ for some $\omega' \neq \omega$, then $\{K^n(E'_{\omega'}) : n \geq 0\}$ is an open neighbourhood of F. Moreover, since $K^{m^*}(E'_\omega) \neq K^{m^*}(E'_{\omega'}) = F$, it directly follows that $E'_\omega \neq E'_{\omega'}$. Yet since $\mathcal{C} = \{E'_{\omega''} : \omega'' \in \Omega\}$ is a partition of Ω, it holds that $E'_\omega \cap E'_{\omega'} = \emptyset$. Moreover, as $K^m(E'_\omega) \subseteq E'_\omega$ for all $m \geq 0$, and $K^n(E'_{\omega'}) \subseteq E'_{\omega'}$ for all $n \geq 0$, as well as $E'_\omega \cap E'_{\omega'} = \emptyset$, it follows that $K^m(E'_\omega) \neq K^n(E'_{\omega'})$ for all $m, n \geq 0$. Thus the open neighbourhood $\{K^n(E'_{\omega'}) : n \geq 0\}$ of F contains no term of the sequence $(K^m(E'_\omega))_{m>0}$ whatsoever. Therefore, F is not a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. Fourthly, if $F \neq K^m(E'_{\omega'})$ for all $m \geq 0$ and all $\omega' \in \Omega$, then $\mathcal{P}(\Omega) \setminus \{K^m(E'_\omega) : m \geq 0 \text{ and } \omega \in \Omega\}$ is an open neighbourhood of F. Yet this set contains no term of the sequence $(K^m(E'_\omega))_{m>0}$. Thus F is not a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. To summarize, there consequently exists no $F \neq K^{m^*}(E'_\omega)$ which is a limit point of the sequence $(K^m(E'_\omega))_{m>0}$. Therefore, the limit point $K^{m^*}(E'_\omega)$ of the sequence $(K^m(E'_\omega))_{m>0}$ is unique. ■