The Hierarchical Approach to Modeling Knowledge and Common Knowledge*

Ronald Fagin
IBM Almaden Research Center
San Jose, CA 95120
fagin@almaden.ibm.com
`http://www.almaden.ibm.com/cs/people/fagin`

John Geanakoplos
Cowles Foundation
Yale University
New Haven, CT 06520
john.geanakoplos@yale.edu

Joseph Y. Halpern†
Computer Science Dept.
Cornell University
Ithaca, NY 14853
halpern@cs.cornell.edu
`http://www.cs.cornell.edu/home/halpern`

Moshe Y. Vardi‡
Dept. of Computer Science
Rice University
Houston, TX 77005-1892
vardi@cs.rice.edu
`http://www.cs.rice.edu/~vardi`

Abstract

One approach to representing knowledge or belief of agents, used by economists and computer scientists, involves an infinite hierarchy of beliefs. Such a hierarchy consists of an agent's beliefs about the state of the world, his beliefs about other agents' beliefs about the world, his beliefs about other agents' beliefs about other agents' beliefs about the world, and so on. (Economists have typically modeled belief in terms of a probability distribution on the uncertainty space. In contrast, computer scientists have modeled belief in terms of a set of worlds, intuitively, the ones the agent considers possible.) We consider the question of when a countably infinite hierarchy completely describes the uncertainty of the agents. We provide various necessary and sufficient conditions for this property. It turns out that the probability-based approach can be viewed as satisfying one of these conditions, which explains why a countable hierarchy suffices in this case. These conditions also show that whether a countable hierarchy suffices may depend on the "richness" of the states in the underlying state space. We also consider the question of whether a countable hierarchy suffices for "interesting" sets of events, and show that the answer depends on the definition of "interesting".

*An extended abstract of this paper appeared under the title "The Expressive Power of the Hierarchical Approach to Modeling Knowledge and Common Knowledge" in *Proc. of Fourth Conf. on Theoretical Aspects of Reasoning about Knowledge*, 1992, pp. 229–244.

†Much of this work was carried out while the author was at the IBM Almaden Research Center. IBM's support is gratefully acknowledged. The work was also supported in part by NSF under grant IRI-96-25901, and by the Air Force Office of Scientific Research under contract F49620-91-C-0080 and grant F49620-96-1-0323.

‡Much of this work was carried out while the author was at the IBM Almaden Research Center. IBM's support is gratefully acknowledged. The work was also supported in part by NSF grants CCR-9628400 and CCR-9700061.

1 Introduction

Reasoning about the knowledge of agents and their knowledge of each other's knowledge has now been recognized as a fundamental concern in game theory, computer science, artificial intelligence, and philosophy (see [FHMV95]). The importance of finding good formal models that can represent the knowledge of the agents has also been long recognized.

The original approach to representing knowledge and common knowledge in the game-theory literature is due to Aumann [Aum76]. Consider a situation with n agents. To model this, Aumann considers structures of the form $A = (W, \mathcal{K}_1, ..., \mathcal{K}_n)$, where W is a set of states of the world, and each $\mathcal{K}_i$ is a partition of W. We henceforth call these *Aumann structures*.[1] An agent "knows" about *events*, which are identified with subsets of W. Agent i's knowledge is modeled by $\mathcal{K}_i$, his *information partition*. Given a state $s \in W$, we use $\mathcal{K}_i(s)$ to denote the set of states in the same element of the partition as s; these are the states that agent i considers to be possible in state s. Agent i is said to *know* an event E at the state s if $\mathcal{K}_i(s)$ is a subset of E. The intuition behind this is that in state s, agent i cannot distinguish between any of the worlds in $\mathcal{K}_i(s)$. Thus, agent i knows E in state s if E holds at all the states that i cannot distinguish from s. Using this intuition, we define an operator K_i from events to events. Given an event E, the event $\mathsf{K}_i(E)$ (intuitively, the event "agent i knows E") is identified with the set of states where agent i knows E according to our definition. We also define the event $\mathsf{O}(E)$ ("everyone knows E") to be the intersection of the events $\mathsf{K}_i(E)$, over all agents $i = 1, \ldots, n$. Finally, we define the event $\mathsf{C}(E)$ ("E is common knowledge") to be the intersection of the events $\mathsf{O}(E)$, $\mathsf{O}(\mathsf{O}(E))$, and so on.

There is, unfortunately, a philosophical difficulty with this approach (cf. [Gil88, TW88, Aum89]). The problem is that it is not *a priori* clear what the relation is between a state in an Aumann structure—which is, after all, just an element of a set—and the rather complicated reality that this state is trying to model. If we think of a state as a complete description of the world, then it must capture all of the agents' knowledge. Since the agents' knowledge is defined in terms of the partitions, the state must include a description of the partitions. This seems to lead to circularity, since the partitions are defined over the states, but the states contain a description of the partitions. One particularly troubling issue, already mentioned in Aumann's original paper, is how the states can be used to capture knowledge about the model itself, such as the fact that the partitions are common knowledge. (See [BD93] for discussion about the importance of this assumption.) Again, there seems to be some circularity here, since the state must describe the model, which therefore includes a description of itself.

Partly in response to these concerns, an alternative approach to modeling knowledge was investigated in a number of economics papers [BE79, MZ85, TW88, BD93]. This approach, which involves an infinite hierarchy of beliefs, takes its cue from the work of Harsanyi [Har68]. We start with a set S of *states of nature*, which we take to be descriptions of certain facts about the world, such as the possible outcomes of a game, and the associated payoffs. Each agent has beliefs about the state of nature, where these beliefs are modeled by a probability distribution over S. These beliefs are clearly highly relevant to the agent's choice of strategy. But agents also have beliefs about other agents' beliefs, and beliefs about other agents' beliefs about their beliefs, and so on. Pursuing this line, one is naturally led to associate with each agent a hierarchy of beliefs. We can build up this hierarchy level by level: at the 0^{th} level is the state of nature; the first-order beliefs of

[1]The reader with a background in modal logic will recognize that an Aumann structure is nothing more than a *Kripke frame for S5* [Kri59, HC68, HM92]. In [Aum76], Aumann assumes that there is a probability distribution on W. Since the probability function plays no role in our discussion of knowledge in Aumann structures, we have decided to drop it here. This is consistent with Aumann's own discussion of knowledge in later papers (see [Aum89]), and with the presentation of Aumann's framework in, for example, [Wer89].

agent i are modeled by a probability distribution on the possible states of nature; for each natural number $m \geq 1$, the $(m+1)^{\text{st}}$-order beliefs of agent i are modeled by a probability distribution on the possible states of nature and the other agents' m^{th}-order beliefs (together with some consistency conditions described in [MZ85, BD93]). An agent's *type* is his infinite hierarchy of beliefs. We define a *belief structure* to consist of a state of nature and a description of each agent's type. Given a set S of states of the world, we take $\mathcal{B}(S)$ to be the set of belief structures where S is the set of states of nature.

In belief structures, knowledge is identified with "belief with probability 1". That is, roughly speaking, agent i is said to *know* an event $E \subseteq S$ in a given belief structure $\mathbf{b}$ if, according to agent i's type in $\mathbf{b}$, event E is assigned probability 1 at level 1 of agent i's hierarchy. Similarly, agent i knows that agent j knows E if the event "agent j knows E" is assigned probability 1 at level 2 of agent i's type hierarchy. Finally, we say that E is common knowledge if all agents know E, all agents know that they know E, and so on.

We would like to think of a belief structure as describing a state of the world. It is not clear, however, that a belief structure is an adequate description of a state of the world. Even if we accept the doctrine that a state of the world can be adequately described by describing the actual state of nature and each agent's uncertainty about the state of nature and other agents' uncertainty (at all levels), it is not clear that the infinite hierarchy just described completely exhausts an agent's uncertainty. After all, an agent may have uncertainty as to the type of other agents. Harsanyi essentially assumed that there is an exogenously given probability distribution that describes each agent's probability distribution on the state of nature and the other agents' types. The key result proved in [BE79, MZ85] is that the hierarchy described above does exhaust an agent's beliefs: an agent's type determines a unique probability distribution on the states of nature and the other agents' types.

This result also suggests that we can view the belief structures in $\mathcal{B}(S)$ as the states in an Aumann structure, since each one completely describes a state of the world. If we take that view, then we might hope that the definitions of knowledge and common knowledge in Aumann structures and belief structures coincide. Unfortunately, this is not quite the case. Nevertheless, Brandenburger and Dekel [BD93] show that these notions do coincide if we interpret knowledge in Aumann structures probabilistically. Thus, we view $\mathcal{B}(S)$ as an Aumann structure, with the information partitions being determined by the type (so that two belief structures $\mathbf{b}$ and $\mathbf{b}'$ are in the same equivalence class of $\mathcal{K}_i$ iff agent i has the same type in $\mathbf{b}$ and $\mathbf{b}'$). In addition, we endow $\mathcal{B}(S)$ with probability measures μ_i (one probability for each agent i) based on information in the individual belief structures (for more details on the construction, see [BD93]). Suppose we identify the event $E \subseteq S$ with the subset of $\mathcal{B}(S)$ consisting of all belief structures for which the state of nature is in E. We then take the event "agent i knows E" to hold in state s if $\mu_i(E \mid \mathcal{K}_i(s)) = 1$; similar modifications are necessary for common knowledge. Brandenburger and Dekel then show that an event $E \subseteq S$ is common knowledge in a state $\mathbf{b}$ in the (probabilistically endowed) Aumann structure $\mathcal{B}(S)$ iff E is common knowledge in the belief structure $\mathbf{b}$. A complementary result is proved in [TW88], where it is shown that given an Aumann structure $A = (S, \mathcal{K}_1, \ldots, \mathcal{K}_n)$ and $s_0 \in S$, there is a belief structure $\mathbf{b} \in \mathcal{B}(S)$ such that an event $E \subseteq S$ is common knowledge at s_0 iff E is common knowledge in $\mathbf{b}$.

This may seem to pretty much complete the picture: the hierarchical approach provides the answer to the problem of circularity in Aumann structures, since the above results seem to indicate that belief structures are adequate for modeling the states in Aumann structures. Unfortunately, the situation is somewhat more complicated than these results suggest. The fundamental problem with these results is that they are trying to relate two incomparable concepts of knowledge:

the information-theoretic concept in Aumann structures and the probability-theoretic concept in belief structures (which is why Brandenburger and Dekel had to recast Aumann's framework in a probabilistic setting). The probabilistic framework masks some of the subtleties in the issue of the adequacy of the hierarchical approach. Thus, we examine the issue of the adequacy of the hierarchical approach here in a non-probabilistic setting.

A non-probabilistic setting for the hierarchical approach is described in [FHV91]. (A precursor to this approach is described in [EGS80].) We again start with a set S of states of nature (at "level 0") and build a hierarchy, level by level. In this case, the first-order knowledge of agent i is a set of states of nature (which intuitively corresponds to the set of states the agent considers possible); the $(m+1)^{\text{st}}$-order knowledge of agent i (for $m \geq 1$) is modeled by a set of possibilities, each of which is a description of a state of nature and each agent's m^{th}-order knowledge (again, certain consistency conditions must be satisfied). Intuitively, whatever is in the subset is considered to be possible, and whatever is not in the subset is known to be impossible. Note that there is no probability distribution, just a set of possibilities. A *knowledge structure* consists of a state of nature and, for each agent, a hierarchy consisting of that agent's m^{th}-order knowledge, for each finite $m \geq 1$. We take $\mathcal{F}(S)$ to be the set of knowledge structures, where S is the set of states of nature.

Knowledge and common knowledge are defined in knowledge structures in an information-theoretic fashion, as in Aumann structures. That is, agent i is said to know $E \subseteq S$ in a given knowledge structure if the set of states that i considers possible at level 1 is a subset of E; agent j knows that agent i knows E if the set of sequences of length 2 that j considers possible at level 2 is a subset of the set of sequences of length 2 where i knows E. Common knowledge is again defined in the standard way in terms of knowledge.

In [FHV91], results connecting knowledge structures and Aumann structures analogous to those of [BD93] and [TW88] are proved. Namely, it is shown that we can view $\mathcal{F}(S)$ as an Aumann structure, where the partitions are determined by the agents' types, and an event $E \subseteq S$ is common knowledge in a knowledge structure $\mathbf{f} \in \mathcal{F}(S)$ according to Aumann's definition iff E is common knowledge at $\mathbf{f}$ according to the knowledge-structure definition. Moreover, it is shown that given an Aumann structure $A = (S, \mathcal{K}_1, \ldots, \mathcal{K}_n)$ and a state $s_0 \in S$, there is a knowledge structure $\mathbf{f} \in \mathcal{F}(S)$ such that an event $E \subseteq S$ is common knowledge at s_0 iff E is common knowledge in $\mathbf{f}$.

This seems to confirm the results of [TW88, BD93] and suggest that the hierarchical approach does address the circularity problem. Unfortunately, it is also shown in [FHV91] that knowledge structures are in general *not* an adequate description of the world, since they do not completely describe an agent's uncertainty. In particular, an agent's type does not determine what other types the agent considers possible. The problem is that the hierarchy in knowledge structures (as well as in belief structures) contains only ω levels, when in general we need to consider transfinite hierarchies.[2] In fact, Fagin [Fag94] and Heifetz and Samet [HS93, HS98] show independently that in general, no ordinal level in the hierarchy is sufficiently large to describe completely an agent's uncertainty. We say more about this result in Section 7.

Why are knowledge structures not an adequate description of an agent's knowledge while belief structures are? And how do we reconcile the inadequacy of knowledge structures with the results relating knowledge structures to Aumann structures? Our goal in this paper is to address these questions by using the non-probabilistic framework of knowledge structures to examine the adequacy of hierarchical structures and to make precise how expressive they are.

We start by considering the question of when a knowledge structure does completely characterize

[2] *Transfinite hierarchies* have levels that are indexed by infinite ordinals (or as they are often called, *transfinite ordinals*). For a discussion of transfinite ordinals, see almost any book on set theory and many books on logic, such as [Sho67].

the agents' knowledge. More precisely, we consider (in Section 3) when it is the case that the first ω levels of the hierarchy completely determine the rest of the hierarchy. We provide three necessary and sufficient conditions for this to be the case. One surprising condition is that a knowledge structure completely characterizes the agents' knowledge iff it characterizes the first $\omega + \omega$ levels of knowledge. A consequence of that is that in order to check if the first ω levels of the hierarchy determine the rest of the hierarchy, it suffices to show that they determine the first $\omega + \omega$ levels of the hierarchy. Another consequence of this condition is that the adequacy of knowledge structures may depend on the "richness" of the states in the underlying state space S. If the states of nature are modeled in enough detail, then knowledge structures do characterize the agents' knowledge; otherwise, they may not.

In Section 4, we provide a different analysis of adequacy, one that sheds further light on why we can stop after ω levels in the probabilistic case. This analysis highlights the role of a certain *limit-closure* property, which says that what happens at finite levels determines what happens at the limit. Limit closure can be viewed as a continuity property. The probabilistic analogue to limit closure holds for belief structures, but only because we restrict attention to countably additive measures. If we allow probabilities that are only finitely additive, then the analogue to limit closure does not hold and, as we show, such belief structures do not in general completely characterize the agents' beliefs. That is, the results of [BE79, MZ85] no longer hold once we consider probabilities that are only finitely additive.

Since knowledge structures do not, in general, characterize the agents' knowledge, we next consider the question as to whether knowledge structures characterize the agents' knowledge with respect to "interesting" sets of events.The answer, of course, depends on what is considered to be an "interesting" set of events. It turns out, for example, that if we consider only events that can be defined from "natural events" by knowledge and common knowledge operators, then knowledge structures are adequate. If, on the other hand, we are interested in common knowledge among coalitions of agents (rather than just common knowledge among *all* the agents), then knowledge structures are not adequate. In this case, a transfinite hierarchy is necessary, but ω^2 levels suffice. Note that this result is quite different from that involving the ordinal $\omega + \omega$ mentioned earlier. The later result says that if all we care about are events that can be defined from the base events and operators for "coalition" common knowledge, then ω^2 levels of the hierarchy suffice. The earlier result applies to arbitrary events, not just interesting ones, and shows that the first ω levels determine the whole hierarchy iff they determine the first $\omega + \omega$ levels of the hierarchy.

This discussion gives the impression that the only issue underlying the adequacy of the hierarchical approach is that of the "length" of the hierarchy. But it is easy to see that knowledge structures are also deficient in a way that no transfinite hierarchy can remedy. Aumann structures contain information about all *conceivable* states, even states that are commonly known not to hold. Thus, Aumann structures enable *counterfactual* reasoning, such as "If Ron Fagin were the President, then he would not have stopped the war against Iraq so soon." A counterfactual statement can be viewed as a statement about a world commonly known not to be possible. (It is presumably common knowledge that Ron Fagin is not the President.) Knowledge structures, on the other hand, do not enable such reasoning, since situations commonly known to be impossible never appear as prefixes in knowledge structures.

It turns out that this deficiency is not inherent in the hierarchical approach, but rather is the result of the manner in which this approach was used in knowledge structures. Knowledge structures were designed to model knowledge; no more, no less. As we show, the hierarchical approach can also be used to define structures that do capture information about conceivable states. These results suggest that hierarchical structures can always serve as adequate models of the world. In general, however, we may need to capture more than just knowledge and we may need to continue the hierarchy into the transfinite ordinals, in order to completely capture the agents' uncertainty. What we choose to capture and how far into the ordinals we need to go depends on the events that we are interested in capturing. Thus, the question of whether knowledge or belief structures as defined are adequate models depends both on what features of the world we are trying to model, and on the events we are interested in describing.

In Section 2, we review knowledge structures and belief structures. In Section 3, we define what it means for a knowledge structure to characterize the agents' knowledge at a given level, and in particular for a knowledge structure to completely characterize the agents' knowledge. We give three necessary and sufficient conditions for a knowledge structure to completely characterize the agents' knowledge, including the result that a knowledge structure completely characterizes the agents' knowledge iff it characterizes the first $\omega + \omega$ levels of knowledge. We also give a simple sufficient condition, which arises naturally in practice, that guarantees that a knowledge structure completely characterizes the agents' knowledge. In order to better understand how the characterization of knowledge in knowledge structures relates to the characterization of beliefs in belief structures, we present in Section 4 an alternative way of capturing the intuition of when a knowledge structure characterizes the agents' knowledge, in terms of the limit-closure condition mentioned above. We also show that belief structures no longer characterize the agents' beliefs if we consider probability measures that are only finitely additive. In Section 5, we consider whether it really is a problem when knowledge structures do not characterize an agents' knowledge, and show that this depends on the set of events we are interested in. In Section 6, we discuss how to modify knowledge structures to model counterfactual statements. In Section 7, we discuss some results related to those in this paper. In Section 8, we give our conclusions. In Appendix A, we give proofs of some theorems from Section 3, and in Appendix B, we give proofs of some theorems from Section 5.

2 Knowledge structures and belief structures: a review

In this section we review the definitions of knowledge structures and belief structures. We begin with knowledge structures. The following material is largely taken from [FHV91], slightly modified to be consistent with the rest of our presentation here. For the sake of generality, and since we will need these definitions later, we define not just knowledge structures, but the more general "λ-worlds" for ordinals λ; knowledge structures are the special case where $\lambda = \omega$.

We start with a set S of *states (of nature)* and a fixed finite set $\{1, \dots, n\}$ of *agents*. For each ordinal $\lambda > 1$ (finite or infinite), we now define *λ-worlds*, by induction on λ. A *0th-order knowledge assignment* f_0 is a member of S, that is, a state of nature (which, intuitively, corresponds to the "real world"). We call $\langle f_0 \rangle$ a *1-world* (since its length is 1). Assume inductively that κ-worlds have been defined for all κ with $1 \leq \kappa < \lambda$. Let W_κ be the set of all κ-worlds, for $\kappa < \lambda$. If $\kappa \geq 1$, then a *κth-order knowledge-assignment* f_κ is a function that associates with each agent i a set $f_\kappa(i) \subseteq W_\kappa$ of "possible κ-worlds"; we think of the worlds in $f_\kappa(i)$ as "possible" for agent i and the worlds in $W_\kappa - f_\kappa(i)$ as "impossible" for agent i. A λ-world is a sequence $\mathbf{f} = \langle f_0, f_1, \dots \rangle$ of length λ such that for each $\kappa < \lambda$, we have that f_κ is a κth-order knowledge assignment and each *κ-prefix* (i.e., prefix of length κ) is a κ-world. If λ is a limit ordinal, there are no further conditions on λ-worlds. If $\lambda = \lambda' + 1$ is a successor ordinal, there are further conditions. Note that in this case, a λ-world is a sequence $\mathbf{f} = \langle f_0, f_1, \dots, f_{\lambda'} \rangle$. Let us use $\mathbf{f}_{<\kappa}$ to denote the κ-prefix of $\mathbf{f}$. The conditions are:

K1. Correctness: $\mathbf{f}_{<\lambda'} \in f_{\lambda'}(i)$.

K2. Introspection: If $\langle g_0, g_1, \ldots\rangle \in f_{\lambda'}(i)$, then $g_\kappa(i) = f_\kappa(i)$ for all κ with $0 < \kappa < \lambda'$.

K3. Extendibility: If $0 < \kappa < \lambda'$, then $\mathbf{g} \in f_\kappa(i)$ iff there is some $\mathbf{h} \in f_{\lambda'}(i)$ such that $\mathbf{g} = \mathbf{h}_{<\kappa}$.

These conditions enforce some intuitive properties of knowledge. Intuitively, K1 says that each agent correctly takes the actual world to be one of the worlds he considers possible. By contrast, for belief, as opposed to knowledge, an agent can (incorrectly) believe that the actual world is not a possibility. K2 implies that agents are introspective about their own knowledge; at each level, they know exactly what they know and what they do not know at lower levels. Finally, K3 says that the different levels of knowledge describing a knowledge world are consistent with each other.

Let $\mathbf{f} = \langle f_0, f_1, \ldots\rangle$ be a λ-world. Define agent i's *type in* $\mathbf{f}$, denoted $\pi_i(\mathbf{f})$, to be the sequence $\langle f_1(i), f_2(i), \ldots\rangle$. We write $\mathbf{f} \sim_i \mathbf{f}'$ if $\pi_i(\mathbf{f}) = \pi_i(\mathbf{f}')$, that is, if i has the same type in $\mathbf{f}$ and $\mathbf{f}'$. Define *i's view (at* $\mathbf{f}$*)*, denoted $\mathbf{f}^{\sim_i}$, to be $\{\mathbf{g} \,|\, \mathbf{f} \sim_i \mathbf{g}\}$. Intuitively, i's view at $\mathbf{f}$ consists of the λ-worlds where i has the same knowledge as in $\mathbf{f}$.

We are in particular interested in ω-worlds, which we refer to (following [FHV91]) as *knowledge structures*. Thus, a knowledge structure describes knowledge of arbitrary finite depth. We use $\mathcal{F}(S)$ to denote the set of knowledge structures over S. We now define knowledge in knowledge structures. Let w be a k-world. We say that agent i *considers w possible* in a knowledge structure $\mathbf{f} = \langle f_0, f_1, \ldots\rangle$ if $w \in f_k(i)$. A *k-ary event* (or *k-event*, for short) is a set of k-worlds. Thus, a 0-event is an assertion about the state of nature, as it is essentially a subset of the set S; a 1-event is an assertion about the state of nature and the agents' knowledge of the state of nature; a 2-event is an assertion about the state of nature, the agents' knowledge of the state of nature, and the agents' knowledge of the agents' knowledge of the state of nature; and so on. Agent i *knows* a k-event E in $\mathbf{f}$ if all the k-worlds agent i considers possible in $\mathbf{f}$ are in E, that is, if $f_k(i) \subseteq E$. This definition of knowledge has the same information-theoretic flavor as the definition of knowledge in Aumann structures given in the introduction.

Belief structures are defined along similar lines. We briefly sketch the definition here, and refer the reader to [MZ85, TW88, BD93] for more details. We start with S, which we assume is endowed with a topology that makes it a compact metric space.[3] Given a compact metric space X, let $\Delta(X)$ denote the set of Borel probability measures on X. If we endow $\Delta(X)$ with the topology of weak convergence of measures, then $\Delta(X)$ is also a compact metric space. Define a sequence of spaces X_k, for $k = 0, 1, 2, \ldots$, inductively, by taking $X_0 = S$ and $X_{k+1} = X_k \times \Delta(X_k)^n$. Thus,

$$X_{k+1} = X_0 \times \Delta(X_0)^n \times \Delta(X_1)^n \times \ldots \times \Delta(X_k)^n.$$

A belief structure $\mathbf{b}$ is a sequence $\langle b_0, b_1, \ldots\rangle$ such that $b_0 \in S$, and $b_k \in \Delta(X_{k-1})^n$ for each $k > 0$. This means that, for $k > 0$, we can view b_k as a function such that for each agent i, we have $b_k(i) \in \Delta(X_{k-1})$. We have consistency conditions B1 and B2 on belief structures that correspond to K2 and K3:

B1. For all $k > 1$, the probability measure $b_k(i)$ assigns probability 1 to the subspace of X_{k-1} consisting of sequences $\langle c_0, \ldots, c_{k-1}\rangle$ with $c_{k-1}(i) = b_{k-1}(i)$. This says that agent i knows his own probability assignment.

B2. For all $k > 1$, the probability measure $b_{k-1}(i)$ is the marginal of $b_k(i)$ on X_{k-2}.

[3] The assumption that S is a compact metric space is made in [TW88]. Variants of this assumption were used in [BE79, BD93, MZ85].

The assumption and all its variants are trivially true if S is finite, which is often a reasonable assumption in practice.

3 When are knowledge structures adequate models of knowledge?

It is shown in [FHV91] that, in a precise sense, knowledge structures are not adequate to fully capture all of an agent's knowledge. An agent's type does not necessarily determine the set of knowledge structures that he considers possible. To make this precise, we need to make clear what we mean by "the knowledge structures that agent i considers possible".

Two definitions are given in [FHV91] for when an agent considers a world possible; these are then shown to be equivalent. We already saw one definition. Let w be a k-world. Recall that agent i considers w possible in a knowledge structure $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$ if $w \in f_k(i)$. There is, however, another notion of possibility. We say that agent i *considers w possible'* in $\mathbf{f}$ if w is a prefix of some knowledge structure $\mathbf{f}'$ such that $\mathbf{f} \sim_i \mathbf{f}'$; i.e., w is the prefix of a knowledge structure that agent i cannot distinguish from $\mathbf{f}$. The following theorem assures us that the two notions of "possible world" are identical.

Theorem 3.1: [FHV91] *Agent i considers a k-world w possible in a knowledge structure* $\mathbf{f}$ *iff agent i considers w possible' in* $\mathbf{f}$.

The notion "possible'" can be thought of as an *external* notion of possibility. It says that we consider each of the knowledge structures $\mathbf{f}'$ that i considers possible (that is, each knowledge structure $\mathbf{f}'$ in i's view $\mathbf{f}^{\sim_i}$) and take its k-prefix. The other notion ("possible") is an *internal* notion: we consider every k-world that i considers possible, by "looking inside" the knowledge structure (at level k). Theorem 3.1 tells us that the external and internal notions coincide. Consequently, agent i knows a k-event E in $\mathbf{f}$ precisely when $\mathbf{f}^{\sim_i}$ is consistent with E; that is, the k-prefix of every knowledge structure in $\mathbf{f}^{\sim_i}$ is in E. In other words, it does not matter whether we define knowledge in terms of possible worlds or in terms of possible' worlds.

Now consider an $(\omega + 1)$-world $\mathbf{f}' = \langle f_0, f_1, \ldots, f_\omega \rangle$, extending the knowledge structure $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$. We frequently abuse notation in such situations by writing $\langle \mathbf{f}, f_\omega \rangle$ as an abbreviation for $\mathbf{f}'$. As before, there are two ways that we can define "the knowledge structures that agent i considers possible in $\mathbf{f}'$". One way is to say that agent i considers the knowledge structure $\mathbf{g}$ possible in $\mathbf{f}'$ precisely if $\mathbf{g} \in f_\omega(i)$. Another way is to say that agent i considers the knowledge structure $\mathbf{g}$ possible' in $\mathbf{f}'$ precisely if $\mathbf{g}$ is the prefix of some $(\omega + 1)$-world $\mathbf{g}'$ such that $\mathbf{f}' \sim_i \mathbf{g}'$. It is shown in [FHV91] that these two ways are *not* equivalent; the set of knowledge structures that agent i considers possible' in $\mathbf{f}'$ is precisely $\mathbf{f}^{\sim_i}$; this is always a superset of $f_\omega(i)$, but equality need not hold. In fact, knowledge structures do not fully describe the agents' knowledge; there are distinct $(\omega + 1)$-worlds that agree on the first ω levels (an example, taken from [FHV91], is given in Example 3.10).

This "discrepancy" can also be described in terms of knowledge of ω-events, which are sets of ω-worlds. We can define knowledge of an ω-event E in the $(\omega + 1)$-world $\mathbf{f}'$, in two ways. We say that agent i *knows* E in $\mathbf{f}'$ if every ω-world $\mathbf{g}$ that agent i considers possible in $\mathbf{f}'$ is in E, i.e., $f_\omega(i) \subseteq E$. We say that agent i *knows'* E in $\mathbf{f}'$ if every ω-world $\mathbf{g}$ that agent i considers possible' in $\mathbf{f}'$ is in E, i.e., $\mathbf{f}^{\sim_i} \subseteq E$. Note that there is possibly a difference between knowing and knowing'; if $f_\omega(i)$ is a proper subset of $\mathbf{f}^{\sim_i}$ then agent i knows but does not know' the ω-event $f_\omega(i)$. In this sense, knowledge structures may not fully describe the agents' knowledge.

The fact that knowledge structures may not fully describe the agents' knowledge should be contrasted with the situation for belief structures, which completely describe the agents' beliefs. Thus, in the case of belief worlds, the first ω levels of the hierarchy completely describe the agents' beliefs, which is not the case for knowledge worlds. To understand this difference better, the first question we want to examine here is when knowledge structures completely describe the agents' knowledge.

3.1 Three characterizations of adequacy

To answer this question of when knowledge structures completely describe the agents' knowledge, we first need to formalize it. If $\lambda \geq \omega$ is an ordinal, then we say that a knowledge structure $\mathbf{f}$ *characterizes the agents' λ-knowledge* if there is a unique extension of $\mathbf{f} = \langle f_0, f_1, f_2, \ldots \rangle$ to a $(\lambda+1)$-world $\langle f_0, f_1, f_2, \ldots, f_\lambda \rangle$. In particular, $\mathbf{f}$ characterizes the agents' ω-knowledge if the "next" level f_ω is uniquely determined. We say that a knowledge structure $\mathbf{f}$ *(completely) characterizes the agents' knowledge* if it characterizes the agents' λ-knowledge for every $\lambda \geq \omega$, that is, if all extensions of $\mathbf{f}$ are determined. This definition captures the intuition that the first ω levels determine the agents' knowledge. As we have already observed, the result of [BE79, MZ85] implies that all belief structures characterize the agents' beliefs in this sense. An example is given in Remark 3.11 where a knowledge structure characterizes the agents' ω-knowledge, but not the agents' knowledge.

There is a very simple case where a knowledge structure characterizes the agents' knowledge: namely, when there is only one agent. In fact, in this case, the first two levels (f_0 and f_1) completely characterize the agent's knowledge:

Proposition 3.2: *Assume that there is only one agent. Let $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$ and $\mathbf{g} = \langle g_0, g_1, \ldots \rangle$ be λ-worlds, where $\lambda \geq 2$. If $f_0 = g_0$ and $f_1 = g_1$, then $\mathbf{f} = \mathbf{g}$.*

Proof: Assume that $f_0 = g_0$ and $f_1 = g_1$. We shall show that $f_2 = g_2$. The proof that $f_\theta = g_\theta$ for every $\theta < \lambda$ is very similar. Suppose the only agent is agent 1. We now show that $f_2(1) = \{\langle h_0, f_1 \rangle \mid h_0 \in f_1(1)\}$. If $\langle h_0, h_1 \rangle \in f_2(1)$, then condition K3 tells us that $h_0 \in f_1(1)$, and condition K2 tells us that $h_1(1) = f_1(1)$, so $h_1 = f_1$ (since agent 1 is the only agent). Conversely, if $h_0 \in f_1(1)$, then condition K3 tells us that there is some h_1 such that $\langle h_0, h_1 \rangle \in f_2(1)$, and condition K2 tells us that $h_1(1) = f_1(1)$, so again $h_1 = f_1$. We have shown that $f_2(1) = \{\langle h_0, f_1 \rangle \mid h_0 \in f_1(1)\}$. Similarly, $g_2(1) = \{\langle h_0, g_1 \rangle \mid h_0 \in g_1(1)\}$. So, since $f_1 = g_1$, it follows that $f_2(1) = g_2(1)$, and so $f_2 = g_2$. ∎

We shall shortly provide a necessary and sufficient condition for a knowledge structure to characterize the agents' knowledge when there are two or more agents. First, we give a necessary and sufficient condition for a knowledge structure to characterize the agents' ω-knowledge.

Assume that agent i considers the k-world w possible in the knowledge structure $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$, that is, $w \in f_k(i)$. By Theorem 3.1, it follows that there is a knowledge structure $\mathbf{g}$ such that w is a prefix of $\mathbf{g}$ and $\mathbf{f} \sim_i \mathbf{g}$. We say that w is *i-uniquely extendible w.r.t. (with respect to)* $\mathbf{f}$ if there is a *unique* such knowledge structure $\mathbf{g}$.

We need another definition before we prove our next theorem. Let $\mathbf{f}$ be a λ-world and let i be an agent. Define the *(one-step) no-information extension* $\mathbf{f}^+$ *of* $\mathbf{f}$ to be the $(\lambda+1)$-world $\langle \mathbf{f}, f_\lambda \rangle$ extending $\mathbf{f}$ such that $f_\lambda(i) = \mathbf{f}^{\sim_i}$ for each agent i. By results in [FHV91], the no-information extension is indeed a $(\lambda+1)$-world. Intuitively, the one-step no-information extension $\mathbf{f}^+$ describes what each agent knows at depth λ, assuming that "all that each agent knows" is already described by $\mathbf{f}$. Thus, in this case $f_\lambda(i)$ is the set of all λ-worlds that are compatible with i's lower-depth knowledge.

Theorem 3.3: *A knowledge structure* $\mathbf{f}$ *characterizes the agents' ω-knowledge iff for each agent i and each knowledge structure $\mathbf{g} \neq \mathbf{f}$ such that $\mathbf{f} \sim_i \mathbf{g}$, some finite prefix of $\mathbf{g}$ is i-uniquely extendible w.r.t. $\mathbf{f}$.*

Proof: ($\Rightarrow$): Assume that there is some agent i and some knowledge structure $\mathbf{g}$ different from $\mathbf{f}$ such that $\mathbf{f} \sim_i \mathbf{g}$, but no finite prefix of $\mathbf{g}$ is i-uniquely extendible w.r.t. $\mathbf{f}$. Therefore, for every finite prefix of $\mathbf{g}$, there is some knowledge structure $\mathbf{h}$ with that prefix such that $\mathbf{f} \sim_i \mathbf{h}$ and $\mathbf{h} \neq \mathbf{g}$.

Let $\mathbf{f}^+ = \langle \mathbf{f}, f_\omega \rangle$ be the one-step no-information extension of $\mathbf{f}$. Thus, $f_\omega(i) = \mathbf{f}^{\sim_i}$. In particular, $\mathbf{g} \in f_\omega(i)$. Define f'_ω by letting $f'_\omega(i) = f_\omega(i) - \{\mathbf{g}\}$, and $f'_\omega(j) = f_\omega(j)$ if $j \neq i$. It is easy to check that $\mathbf{f}' = \langle \mathbf{f}, f'_\omega \rangle$ is an $(\omega + 1)$-world: the correctness condition K1 holds, since $\mathbf{g} \neq \mathbf{f}$; the introspection condition K2 is immediate; and the extendibility condition K3 holds, since for every finite prefix of $\mathbf{g}$ there is some $\mathbf{h}$ with that prefix, such that $\mathbf{h} \sim_i \mathbf{f}$ and $\mathbf{h} \neq \mathbf{g}$. Since $\mathbf{f}' \neq \mathbf{f}^+$, it follows that $\mathbf{f}$ does not characterize the agents' ω-knowledge.

($\Leftarrow$): Let $\langle \mathbf{f}, f_\omega \rangle$ be an arbitrary $(\omega + 1)$-world extending $\mathbf{f}$. We shall show that $f_\omega(i) = \mathbf{f}^{\sim_i}$ for each agent i. We first show that $f_\omega(i) \subseteq \mathbf{f}^{\sim_i}$. Assume that $\mathbf{h} = \langle h_0, h_1, \ldots \rangle \in f_\omega(i)$; we must show that $\mathbf{h} \in \mathbf{f}^{\sim_i}$. By K2, we have that $h_k(i) = f_k(i)$ for all $k \geq 1$. It follows that $\mathbf{h} \sim_i \mathbf{f}$. That is, $\mathbf{h} \in \mathbf{f}^{\sim_i}$, as desired.

Conversely, assume that $\mathbf{g} \in \mathbf{f}^{\sim_i}$; we must show that $\mathbf{g} \in f_\omega(i)$. If $\mathbf{g} = \mathbf{f}$, then $\mathbf{g} \in f_\omega(i)$, by condition K1. So assume that $\mathbf{g} \neq \mathbf{f}$. By assumption, some finite prefix w of $\mathbf{g}$ is i-uniquely extendible w.r.t. $\mathbf{f}$. This tells us that $\mathbf{g}$ is the unique member of $\mathbf{f}^{\sim_i}$ with prefix w. By condition K3, there must be some $\mathbf{g}' \in f_\omega(i)$ with prefix w. Since $f_\omega(i) \subseteq \mathbf{f}^{\sim_i}$, it follows that $\mathbf{g}' \in \mathbf{f}^{\sim_i}$. Since $\mathbf{g}' \in \mathbf{f}^{\sim_i}$ and $\mathbf{g}'$ has prefix w, it follows by uniqueness that $\mathbf{g}' = \mathbf{g}$, and that so $\mathbf{g} \in f_\omega(i)$, as desired.

Thus, $f_\omega(i) = \mathbf{f}^{\sim_i}$. We have shown that f_ω is uniquely determined by $\mathbf{f}$, since $f_\omega(i) = \mathbf{f}^{\sim_i}$ for each agent i. So $\mathbf{f}$ characterizes the agents' ω-knowledge. ∎

Theorem 3.3, which will turn out to be quite useful, gives a sense in which knowledge at finite levels determines when the agents' ω-knowledge is "forced" to a unique value.

We might hope that if a knowledge structure characterizes the agents' ω-knowledge, then it completely characterizes the agents' knowledge. Unfortunately, this is not the case. For example, there is a knowledge structure with two agents that characterizes the agents' ω-knowledge and has two extensions to $(\omega + 1)$-worlds: Roughly speaking, in one of these, agent 2 knows that agent 1's ω-knowledge is characterized, and in the other extension agent 2 does not know this. Another example of a knowledge structure that characterizes the agents' ω-knowledge but does not characterize the agents' knowledge is given in Remark 3.11. As the next result shows, a knowledge structure characterizes the agents' knowledge iff it is *common knowledge* that the first ω levels characterizes the agents' ω-knowledge. To make this precise, we need some more definitions.

Let $\mathbf{f}$ and $\mathbf{g}$ be knowledge structures. We say that $\mathbf{g}$ is *reachable* from $\mathbf{f}$ (by a path of length r) if there are knowledge structures $\mathbf{h}_0, \ldots, \mathbf{h}_r$ such that $\mathbf{f} = \mathbf{h}_0$, $\mathbf{g} = \mathbf{h}_r$, and for all $j < r$, we have $\mathbf{h}_j \sim_i \mathbf{h}_{j+1}$ for some agent i. There is a close connection between reachability and common knowledge. For example, it can be shown that an event $E \subseteq S$ is common knowledge in $\mathbf{f}$ iff E holds at each knowledge structure reachable from $\mathbf{f}$. (See [Aum76, HM92] for analogous results in the context of Aumann structures, and [TW88] for an analogous result in the context of belief structures.)

The following two theorems give necessary and sufficient conditions for a knowledge structure to characterize the agents' knowledge.

Theorem 3.4: *A knowledge structure* $\mathbf{f}$ *characterizes the agents' knowledge iff every knowledge structure reachable from* $\mathbf{f}$ *characterizes the agents'* ω*-knowledge.*

Proof: See Appendix A. ∎

It is not hard to provide examples of knowledge structures that do and knowledge structures that do not characterize the agents' knowledge. For example, given a k-world w, define (as in [FHV91]) the *no-information extension* w^* of w by repeatedly taking one-step no-information extensions. Informally, w^* is the knowledge structure where all each agent knows is what is already described by w. It can be shown from the construction of the one-step no-information extension that w^* does not characterize the agents' ω-knowledge. We shall see another example later (Example 3.10) where the knowledge structure does not characterize the agents' ω-knowledge. An example of a knowledge structure that characterizes the agents' knowledge is one where the state of nature is common knowledge. This is a knowledge structure $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$ where every $f_k(i)$ is a singleton set. We leave to the reader the straightforward verification, using Theorem 3.4, that such a knowledge structure characterizes the agents' knowledge.

As we noted, there exist knowledge structures that characterize the agents' ω-knowledge, but do not completely characterize the agents' knowledge. Rather surprisingly, it turns out that if a knowledge structure $\mathbf{f}$ characterizes the agents' knowledge through the first $\omega + \omega$ levels (that is, if $\mathbf{f}$ characterizes the agents' $(\omega + k)$-knowledge for every natural number k), then $\mathbf{f}$ completely characterizes the agents' knowledge.

Theorem 3.5: *A knowledge structure characterizes the agents' knowledge iff it characterizes the agents' knowledge through the first $\omega + \omega$ levels.*

Proof: See Appendix A.

We now provide another characterization of knowledge structures that characterize the agents' knowledge, in the case where the state space S is finite.

Theorem 3.6: *Assume that there are only finitely many states of nature. A knowledge structure* $\mathbf{f}$ *characterizes the agents' knowledge iff* $\mathbf{g}^{\sim_i}$ *is finite for every knowledge structure* $\mathbf{g}$ *reachable from* $\mathbf{f}$ *and every agent i.*

Proof: See Appendix A. ∎

If $\mathbf{f}$ is a knowledge structure, then let $G_{\mathbf{f}}$ be a graph whose nodes are all knowledge structures reachable from $\mathbf{f}$, such that there is an edge between two nodes $\mathbf{g}$ and $\mathbf{h}$ iff $\mathbf{g} \sim_i \mathbf{h}$ for some agent i. Then Theorem 3.6 says that $\mathbf{f}$ characterizes the agents' knowledge iff $G_{\mathbf{f}}$ has finite fanout at every node. This is closely related to Theorem 5.7 of [Fag94], which gives a similar finite fanout characterization for structures like knowledge structures, except that they do not satisfy condition K1.

3.2 A sufficient condition for characterizing the agents' knowledge

To gain a better understanding of the issue of characterization of knowledge, we now consider a simple sufficient condition on knowledge structures that guarantees characterization of the agents' knowledge. Let $\mathbf{f}$ be a knowledge structure. A world *is reachable from* $\mathbf{f}$ if it is a prefix of a knowledge structure that is reachable from $\mathbf{f}$. Intuitively, a world w is reachable from $\mathbf{f}$ if some agent considers it possible that some agent considers it possible $\cdots$ that some agent considers w possible. We say that it is *common knowledge in* $\mathbf{f}$ *how the state of nature determines the agents' knowledge* if whenever $w = \langle g_0, \ldots, g_r \rangle$ and $w' = \langle g'_0, \ldots, g'_r \rangle$ are reachable from $\mathbf{f}$, and $g_0 = g'_0$, then $w = w'$. Intuitively, it is common knowledge in $\mathbf{f}$ how the state of nature determines the agents' knowledge if there is a "commonly-known algorithm" for determining each agent's finite levels of knowledge from the state of nature. It can be easily shown that it is common knowledge in $\mathbf{f}$ how the state of nature determines the agents' knowledge precisely if whenever $\mathbf{g}$ and $\mathbf{g}'$ are reachable from $\mathbf{f}$, and the state of nature is the same in $\mathbf{g}$ and $\mathbf{g}'$, then $\mathbf{g} = \mathbf{g}'$. The next lemma follows easily from this characterization.

Lemma 3.7: *Assume that it is common knowledge in* $\mathbf{f}$ *how the state of nature determines the agents' knowledge. Assume also that* $\mathbf{h}$ *is reachable from* $\mathbf{f}$*. Then it is common knowledge in* $\mathbf{h}$ *how the state of nature determines the agents' knowledge.*

The next proposition gives us a simple sufficient condition on a knowledge structure that guarantees that it characterizes the agents' knowledge.

Proposition 3.8: *Assume that it is common knowledge in the knowledge structure* $\mathbf{f}$ *how the state of nature determines the agents' knowledge. Then* $\mathbf{f}$ *characterizes the agents' knowledge.*

Proof: By Theorem 3.4, it suffices to show that if $\mathbf{h}$ is reachable from $\mathbf{f}$, then $\mathbf{h}$ characterizes the agents' ω-knowledge. Theorem 3.3 tells us that to show this, we need only show that for each agent i and each knowledge structure $\mathbf{g} = \langle g_0, g_1, \ldots \rangle \neq \mathbf{h}$ such that $\mathbf{h} \sim_i \mathbf{g}$, some finite prefix w of $\mathbf{g}$ is i-uniquely extendible w.r.t. $\mathbf{h}$. Let w be the prefix $\langle g_0 \rangle$. By Lemma 3.7, it is common knowledge in $\mathbf{h}$ how the state of nature determines the agents' knowledge. Therefore, if $\mathbf{g}'$ is a knowledge structure such that $\mathbf{h} \sim_i \mathbf{g}'$, and $\mathbf{g}'$ has prefix w, then $\mathbf{g}' = \mathbf{g}$. Hence, w is i-uniquely extendible w.r.t. $\mathbf{h}$, as desired. ∎

The interest in Proposition 3.8 comes from the fact that the way an agent determines what states are possible (or, in the case of belief structures, the way an agent determines how to assign probabilities) clearly ultimately depends on circumstances external to the agent, including perhaps what the agent has observed, the agent's upbringing, and a myriad of other influences. In many applications, the most natural way to model the state of nature will capture these external circumstances, and therefore it is common knowledge how the state of nature determines the agents' knowledge. The following simple example, based on the coordinated attack problem discussed in [HM90] (and later modified as the electronic mail game by Rubinstein [Rub89]), may clarify this.

Example 3.9: There are three agents, 1, 2, and 3. Consider a fact p such as "the price of IBM stock is over \$100". Suppose agents 1 and 3 discover whether or not p holds, and agent 2 does not. If p does not hold, then nothing happens. If p holds, then agents 1 and 2 start to communicate about p over an unreliable channel. First agent 1 tells agent 2 that p holds. If agent 2 receives the message, he sends an acknowledgment. If agent 1 receives the acknowledgment, he acknowledges the acknowledgment, and so on. If at any point a message is not received, there is no further communication. There is never any communication between agent 3 and the other two agents. We consider the system at some time after agent 1 discovers p. We also assume that agent 3 has no idea how much time has passed, so that, if p holds, he has no upper bound on the number of messages that may have been received by agents 1 and 2. We can thus take S to consist of $\bar{p}$ (the state where the negation $\bar{p}$ of p holds) and pairs of the form (p, k), $k \geq 0$; intuitively, these are the states where p holds, k messages were received by 1 and 2, and a $(k+1)^{\text{st}}$ message was sent by the recipient of the k^{th} message (or by agent 1 if $k = 0$), but not received.

In this situation, it is common knowledge how the state of nature determines the agents' knowledge. Intuitively, this is because once we know how many messages have been received, we can determine each agent's knowledge. For example, suppose that the state of nature is $(p, 2)$, so that

p holds and two messages have been received (thus far) between 1 and 2 (i.e., 2 received 1's initial message, and 1 received 2's acknowledgment). Then at the first level, agent 1 considers the states $(p, 2)$ and $(p, 3)$ possible (since agent 1 does not know whether his acknowledgment to agent 2's last message was received by agent 2) and 2 considers the states $(p, 1)$ and $(p, 2)$ possible (since agent 2 does not know whether agent 1 received the last acknowledgment he sent). Agent 3 considers all states of the form (p, k), $k \geq 0$ possible, since he knows p holds, but has no idea how many messages have passed between agents 1 and 2. It is not hard to see how we can continue this construction in a unique way, level by level. Since it is common knowledge how the state of nature determines the agents' knowledge, it follows from Proposition 3.8 that each knowledge structure that arises in this scenario characterizes the agents' knowledge.

Before leaving this example, let us consider what knowledge the agents have in each of the knowledge structures that arise in this scenario. Let E be the set of states of nature of the form (p, k); intuitively, E corresponds to the event that "p holds". In (the knowledge structure that corresponds to) the state $(p, 0)$, agent 1 knows E but agent 2 does not know that agent 1 knows E; in the state $(p, 1)$, agent 2 knows that agent 1 knows E, but agent 1 does not know that agent 2 knows that agent 1 knows E; and so on. Thus, for none of these states does common knowledge of E ever hold between agents 1 and 2, where agents 1 and 2 are said to have common knowledge of E if both 1 and 2 know that both 1 and 2 know ... that E holds (cf. the discussion of the coordinated attack problem in [HM90]). Now consider agent 3. Informally, in every state agent 3 certainly knows that agents 1 and 2 do not have common knowledge of E (since they never attain common knowledge of E when communicating over an unreliable channel). He considers it possible, however, that agents 1 and 2 have arbitrarily deep knowledge of E (since agent 3 considers all the states $(p, 0), (p, 1), (p, 2), \ldots$ possible). More precisely, if $\mathbf{f}^s$ is the knowledge structure associated with a state $s \in S$, then in the unique extension $\langle \mathbf{f}^s, f^s_\omega \rangle$ of $\mathbf{f}^s$ to an $\omega + 1$-world (the extension is unique because $\mathbf{f}^s$ characterizes the agents' knowledge), $f^s_\omega(3)$ consists of every knowledge structure $\mathbf{f}^{s'}$ for $s' \in S$. Thus, agent 3 knows that E is not common knowledge among the other two agents, and considers it possible that they have arbitrarily deep knowledge. ∎

While in simple examples it does seem reasonable to include enough information in the state of nature so that it is common knowledge how the state of nature determines the agents' knowledge, in more complicated examples this becomes a serious modeling problem. For example, even if we accept that the sum total of an agent's upbringing, together with hereditary factors and all the agent's experience and observations, completely determines the agent's knowledge, it is not clear that we want to include all this information in the state of nature when modeling, say, a simple game. Once we leave it out, however, the knowledge structure may no longer adequately model the agents' knowledge, as the following example shows.

Example 3.10: Suppose we consider the same situation as in Example 3.9, but change the description of the state of nature. Instead of the state of nature describing not only whether p is true, but also how many messages arrive, suppose we simply take the state of nature to describe whether or not p is true. Thus, there are only two states of nature, p and $\bar{p}$. Essentially, all the states of nature of the form (p, k) have been collapsed to one state, p. Thus, there are two 1-worlds, $\langle \bar{p} \rangle$ and $\langle p \rangle$, which we denote $w_{1,-1}$ and $w_{1,0}$, respectively. (The first component of the subscript represents the length of the world; the reason for the choice of the second component should become clearer shortly.) We construct the k-worlds for $k \geq 2$ inductively. There are three 2-worlds:

- $w_{2,-1} = \langle \bar{p}, f_1 \rangle$, where $f_1(1) = f_1(3) = \{\langle \bar{p} \rangle\}$ and $f_1(2) = \{\langle p \rangle, \langle \bar{p} \rangle\}$. This is the world where $\bar{p}$ is true. Both agents 1 and 3 know this, and 2 does not (since 1 sends no messages in this case, and 2 considers it possible that p is the case and 1's message did not arrive).

- $w_{2,0} = \langle p, f'_1 \rangle$, where $f'_1(1) = f'_1(3) = \{\langle p \rangle\}$ and $f'_1(2) = \{\langle p \rangle, \langle \bar{p} \rangle\}$. This is the world where p is true, but 1's message to 2 does not arrive.
- $w_{2,1} = \langle p, f''_1 \rangle$, where $f''_1(1) = f''_1(2) = f''_1(3) = \{\langle p \rangle\}$. This is the world where p is true, and 1's message to 2 does arrive.

Notice that $w_{2,-1}$ corresponds to the unique 2-world in the previous example where the state of nature is $\bar{p}$; $w_{2,0}$ corresponds to the unique 2-world where the state of nature is $(p, 0)$; and $w_{2,1}$ can be viewed as the result of "collapsing" all the 2-worlds where the state of nature is (p, k') for $k' \geq 1$.

For $k > 2$, we have a similar phenomenon. There are precisely $k+1$ distinct k-worlds, which we denote $w_{k,-1}, w_{k,0}, \ldots, w_{k,k-1}$, where $w_{k,-1}$ corresponds to the unique k-world in the previous example where the state of nature is $\bar{p}$; $w_{k,j}$ (for $0 \leq j < k-1$) corresponds to the unique k-world in the previous example where the state of nature is (p, j); and $w_{k,k-1}$ is the result of collapsing all k-worlds (p, j) with $j \geq k-1$ in the previous example. Notice that if $k \geq 2$, $j \geq 0$, and $w_{k,j} = \langle f_0, \ldots, f_{k-1} \rangle$, then $f_{k-1}(3) = \{w_{k-1,0}, \ldots, w_{k-1,k-2}\}$; if the state of nature is p, agent 3 has no idea how many messages passed between agents 1 and 2. If $0 \leq j < k-1$, and j is even (which means that agent 1's last message is in transit or was not delivered), then $f_{k-1}(1) = \{w_{k-1,j}, w_{k-1,j+1}\}$, since agent 1 does not know whether or not his last message was delivered, and $f_{k-1}(2) = \{w_{k-1,j-1}, w_{k-1,j}\}$. The situation is similar if j is odd. Finally, if $j = k-1$, then $f_{k-1}(1) = f_{k-1}(2) = \{w_{k-1,k-2}\}$.

We can denote the knowledge structures that arise in this example as $\mathbf{f}_{-1}, \mathbf{f}_0, \mathbf{f}_1, \ldots$. The knowledge structure $\mathbf{f}_{-1}$ has as prefixes the worlds $w_{j,-1}$, for $j = 1, 2, 3, \ldots$, and corresponds to the unique knowledge structure in the previous example where the state of nature is $\bar{p}$. The knowledge structure $\mathbf{f}_j$ has as prefixes the worlds $w_{k,k-1}$ for $1 \leq k \leq j$ and $w_{k,j}$ for $k > j$, and corresponds to the knowledge structure in the previous example where the state of nature is (p, j). Notice that the knowledge structure $\mathbf{f}_\infty$ with prefixes $w_{1,0}, w_{2,1}, w_{3,2}, \ldots$ does not arise in this situation (although it is easy to check that $\mathbf{f}_\infty$ is indeed a well-defined knowledge structure). Intuitively, $\mathbf{f}_\infty$ corresponds to the situation where infinitely many messages passed between agents 1 and 2, a situation that is commonly known to be impossible. In $\mathbf{f}_\infty$, the event E (where p holds) is common knowledge among agents 1 and 2. Intuitively, it is because $\mathbf{f}_\infty$ is commonly known to be impossible that agent 3 knows that agents 1 and 2 do not have common knowledge of E. Nevertheless, none of the knowledge structures where p holds that arise in this example capture the fact that $\mathbf{f}_\infty$ is (commonly known to be) impossible. Consider any knowledge structure $\mathbf{f}_j$ with $j \geq 0$, and let $\mathbf{f}_j^+ = \langle \mathbf{f}_j, f_\omega \rangle$ be the one-step no-information extension of $\mathbf{f}_j$. It is not hard to see that $\mathbf{f}_\infty \in f_\omega(3)$, so that in $\mathbf{f}_j^+$, agent 3 does *not* know that agents 1 and 2 do not have common knowledge of E. Of course, there is another extension $\mathbf{f}'_j = \langle \mathbf{f}_j, f'_\omega \rangle$ of $\mathbf{f}_j$ such that $\mathbf{f}_\infty \notin f'_\omega(3)$. This shows that $\mathbf{f}_j$ does not characterize the agents' ω-knowledge. This provides the example that we promised after Theorem 3.4 of a knowledge structure that does not characterize the agents' ω-knowledge. By contrast, the knowledge structure that arises in Example 3.9 does characterize the agents' knowledge. ∎

Remark 3.11: Proposition 3.8 can be strengthened in a number of straightforward ways. One is as follows: We say that it is *common knowledge how level k determines the agents' knowledge* if whenever $w = \langle g_0, \ldots, g_k, \ldots, g_r \rangle$ and $w' = \langle g'_0, \ldots, g'_k, \ldots, g'_r \rangle$ are reachable from $\mathbf{f}$, and their prefixes $\langle g_0, \ldots, g_k \rangle$ and $\langle g'_0, \ldots, g'_k \rangle$ are identical, then $w = w'$. Then Proposition 3.8 still holds when we replace "it is common knowledge how the state of nature determines the agents' knowledge" by "for some k, it is common knowledge how level k determines the agents' knowledge."

We can further strengthen Proposition 3.8 by further weakening the hypotheses: Let $\mathbf{f}$ be a knowledge structure, and let k be a fixed natural number. We say that *agent i knows in* $\mathbf{f}$ *that level k determines the agents' knowledge* if whenever $\mathbf{g} = \langle g_0, g_1, \ldots \rangle$ and $\mathbf{g}' = \langle g'_0, g'_1, \ldots \rangle$ are knowledge structures such that (a) $\mathbf{f} \sim_i \mathbf{g}$, (b) $\mathbf{f} \sim_i \mathbf{g}'$, and (c) the prefixes $\langle g_0, \ldots, g_k \rangle$ and $\langle g'_0, \ldots, g'_k \rangle$ are identical, then $\mathbf{g} = \mathbf{g}'$. Intuitively, this says that level k completely determines the knowledge structure, among those knowledge structures that agent i considers possible. We say that it is *common knowledge in* $\mathbf{f}$ *that level k determines the agents' knowledge* if in every knowledge structure reachable from $\mathbf{f}$, every agent knows that level k determines the agents' knowledge. Assume for now that $\mathbf{f}$ is a knowledge structure where for some k, each agent knows that level k determines the agents' knowledge. It turns out that this condition is *not* sufficient to guarantee that $\mathbf{f}$ completely characterizes the agents' knowledge, even if $k = 0$, that is, even if each agent knows that the state of nature determines the agents' knowledge. Nevertheless, we can show that this assumption (that for some k, each agent knows that level k determines the agents' knowledge) is sufficient to guarantee that the knowledge structure characterizes the agents' ω-knowledge. Note that such knowledge structures $\mathbf{f}$ provide an example, as promised before Proposition 3.2, where the knowledge structure characterizes the agents' ω-knowledge but not the agents' knowledge. If $\mathbf{f}$ is a knowledge structure where for some k, it is common knowledge that level k determines the agents' knowledge, then $\mathbf{f}$ characterizes the agents' knowledge. This is because every knowledge structure reachable from $\mathbf{f}$ then characterizes the agents' ω-knowledge, and so by Theorem 3.4, it follows that $\mathbf{f}$ characterizes the agents' knowledge.

Notice that the definition of common knowledge *that* level k determines the agents' knowledge is different from our earlier definition of common knowledge *how* level k determines the agents' knowledge. It is common knowledge in $\mathbf{f}$ that level k determines the agents' knowledge if in every knowledge structure $\mathbf{g}$ reachable from $\mathbf{f}$, every agent knows that level k determines the agents' knowledge. It is possible, however, that there are two different knowledge structures $\mathbf{g}$ and $\mathbf{g}'$, both reachable from $\mathbf{f}$, that have the same prefix through level k. This cannot happen if it is common knowledge *how* level k determines the agents' knowledge. It is not hard to show that "common knowledge how" implies "common knowledge that". ∎

4 An alternative view of adequacy

How does the characterization of knowledge in knowledge structures relate to the characterization of beliefs in belief structures? To answer this question, we now provide another necessary and sufficient condition for when a knowledge structure characterizes the agents' knowledge. This time, we consider when it is the case that there is enough information in a knowledge structure to determine what other knowledge structures each agent considers possible.

Let $\mathbf{f} = \langle f_0, f_1, f_2, \ldots \rangle$ be a knowledge structure. What are the possibilities for the set $\mathcal{P}$ of knowledge structures that agent i considers possible? That is, what are the possible values of $f_\omega(i)$ for extensions $\langle \mathbf{f}, f_\omega \rangle$ of $\mathbf{f}$ to an $(\omega+1)$-world? If $\mathbf{f}$ characterizes the agents' knowledge, then $\mathcal{P}$ would be precisely $\mathbf{f}^{\sim_i} = \{\mathbf{g} \,|\, \mathbf{f} \sim_i \mathbf{g}\}$. On the other hand, if agent i has *more* information than is described in $\mathbf{f}$, then he might consider only some proper subset of $\mathbf{f}^{\sim_i}$ possible. Notice that if $w \in f_k(i)$ for some k, so that agent i considers w possible, then $\mathcal{P}$ should contain some knowledge structure $\mathbf{f}'$ such that w is a prefix of $\mathbf{f}'$. We say that a set $\mathcal{P}$ of knowledge structures is a *coherent set of possibilities for agent i at* $\mathbf{f}$ if

P1. $\mathbf{f} \in \mathcal{P}$.

P2. $\mathcal{P} \subseteq \mathbf{f}^{\sim_i}$.

P3. If $w \in f_k(i)$ for some $k > 0$, then there is a knowledge structure $\mathbf{f}' \in \mathcal{P}$ such that w is a prefix of $\mathbf{f}'$.

Condition P1, which is analogous to the correctness condition K1, says that the agent considers $\mathbf{f}$ as a possibility. Condition P2, which is analogous to condition K2, says that the agent has at least as much information as is contained in $\mathbf{f}$. Condition P3, which is analogous to condition K3, is an extendibility condition.

Let $\mathcal{P}$ be a set of knowledge structures, and let $\mathbf{f}$ and $\mathbf{f}'$ be members of $\mathcal{P}$. It is easy to see that $\mathcal{P}$ is a coherent set of possibilities for agent i at $\mathbf{f}$ iff $\mathcal{P}$ is a coherent set of possibilities for agent i at $\mathbf{f}'$. Therefore, we say that $\mathcal{P}$ is a *coherent set of possibilities for agent i* if it is a coherent set of possibilities for agent i at $\mathbf{f}$, for every $\mathbf{f} \in \mathcal{P}$.

As expected, agent i always has at least one coherent set of possibilities at $\mathbf{f}$, namely $\mathbf{f}^{\sim_i}$.

Lemma 4.1: *$\mathbf{f}^{\sim_i}$ is a coherent set of possibilities for agent i at $\mathbf{f}$.*

Proof: Condition P1 holds, since $\mathbf{f} \sim_i \mathbf{f}$. Condition P2 holds, since $\mathbf{f}^{\sim_i} \subseteq \mathbf{f}^{\sim_i}$. Condition P3 follows immediately from Theorem 3.1. ∎

If there is only one coherent set of possibilities at $\mathbf{f}$ for each agent i, then we might expect that $\mathbf{f}$ characterizes the agents' ω-knowledge. The following result shows that this is indeed the case.

Theorem 4.2: *The knowledge structure $\mathbf{f}$ characterizes the agents' ω-knowledge iff there is only one coherent set of possibilities at $\mathbf{f}$ for each agent i.*

Proof: Assume that there is only one coherent set of possibilities at $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$ for each agent i. Let $\langle f_0, f_1, \ldots, f_\omega \rangle$ be an $(\omega+1)$-world that extends $\mathbf{f}$. It follows easily from the consistency conditions on $(\omega+1)$-worlds that $f_\omega(i)$ is a coherent set of possibilities at $\mathbf{f}$, for each agent i. So by assumption, $f_\omega(i)$ is uniquely determined by $\mathbf{f}$, for each agent i. Therefore, by definition, the knowledge structure $\mathbf{f}$ characterizes the agents' ω-knowledge.

Conversely, assume that there are two distinct coherent sets of possibilities at $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$, for some agent i. Let us denote these two distinct coherent sets by $\mathcal{P}$ and $\mathcal{P}'$. Define f_ω by letting $f_\omega(i) = \mathcal{P}$, and $f_\omega(j) = \mathbf{f}^{\sim_j}$ for $j \neq i$. Similarly, define f'_ω by letting $f'_\omega(i) = \mathcal{P}'$, and $f'_\omega(j) = \mathbf{f}^{\sim_j}$ for $j \neq i$. It is straightforward to verify that $\langle f_0, f_1, \ldots, f_\omega \rangle$ and $\langle f_0, f_1, \ldots, f'_\omega \rangle$ are distinct extensions of $\mathbf{f}$. Therefore, $\mathbf{f}$ does not characterize the agents' ω-knowledge. ∎

Thinking in terms of coherent sets of possibilities gives us some insight into why belief structures do characterize the agents' beliefs. We say that a set $\mathcal{P}$ of knowledge structures is *limit closed* if a knowledge structure $\mathbf{g} = \langle g_0, g_1, \ldots \rangle$ is in $\mathcal{P}$ whenever, for all k, there is a knowledge structure $\mathbf{g}^k \in \mathcal{P}$ such that $\langle g_0, \ldots, g_k \rangle$ is a prefix of $\mathbf{g}^k$. Thus, $\mathcal{P}$ is limit closed if, whenever every finite prefix of a knowledge structure appears in $\mathcal{P}$, then the whole knowledge structure appears in $\mathcal{P}$.

The next result shows that $\mathbf{f}^{\sim_i}$ is limit closed. A coherent set of possibilities need not, however, be limit closed in general. As the next result shows, if it is limit closed, then it must in fact be $\mathbf{f}^{\sim_i}$.

Proposition 4.3: *$\mathcal{P}$ is a limit-closed coherent set of possibilities for agent i at $\mathbf{f}$ iff $\mathcal{P} = \mathbf{f}^{\sim_i}$.*

Proof: By Lemma 4.1, $\mathbf{f}^{\sim_i}$ is a coherent set of possibilities for agent i at $\mathbf{f}$. To show that it is limit closed, let $\mathbf{g} = \langle g_0, g_1, \ldots \rangle$ be a knowledge structure such that for all k, there is a knowledge structure $\mathbf{g}^k \in \mathbf{f}^{\sim_i}$ where $\langle g_0, \ldots, g_k \rangle$ is a prefix of $\mathbf{g}^k$. We want to show that $\mathbf{f} \sim_i \mathbf{g}$. Since $\mathbf{f} \sim_i \mathbf{g}^k$, we have $f_k(i) = g_k(i)$. Since this is true for every k, it follows that $\mathbf{f} \sim_i \mathbf{g}$. So $\mathbf{g} \in \mathbf{f}^{\sim_i}$, as desired.

For the converse, suppose that $\mathcal{P}$ is a limit-closed coherent set of possibilities for agent i at $\mathbf{f}$. Since $\mathcal{P} \subseteq \mathbf{f}^{\sim_i}$ by condition P2, we need only show that $\mathbf{f}^{\sim_i} \subseteq \mathcal{P}$. Assume that $\mathbf{g} = \langle g_0, g_1, \ldots \rangle \in \mathbf{f}^{\sim_i}$, that is, $\mathbf{f} \sim_i \mathbf{g}$. We want to show that $\mathbf{g} \in \mathcal{P}$. By K1, for all k we have $\langle g_0, \ldots, g_k \rangle \in g_{k+1}(i)$. Since $\mathbf{f} \sim_i \mathbf{g}$, we must have $g_{k+1}(i) = f_{k+1}(i)$. Hence, for all k, we have $\langle g_0, \ldots, g_k \rangle \in f_{k+1}(i)$. It follows from P3 that for every k, there is a knowledge structure $\mathbf{g}^k \in \mathcal{P}$ such that $\langle g_0, \ldots, g_k \rangle$ is a prefix of $\mathbf{g}^k$. Since $\mathcal{P}$ is limit closed, we have that $\mathbf{g} \in \mathcal{P}$, as desired. ∎

Proposition 4.3 tells us that if *every* coherent set of possibilities for each agent i at $\mathbf{f}$ is limit closed, then there is only one coherent set of possibilities, namely $\mathbf{f}^{\sim_i}$. It then follows from Theorem 4.2 that $\mathbf{f}$ characterizes the agents' ω-knowledge. Combining Theorem 3.4, Lemma 4.1, Theorem 4.2, and Proposition 4.3, we immediately get the following characterization of adequacy.

Theorem 4.4: *A knowledge structure* $\mathbf{f}$ *characterizes the agents' knowledge iff for every knowledge structure* $\mathbf{g}$ *reachable from* $\mathbf{f}$ *and each agent* i*, every coherent set of possibilities for* i *at* $\mathbf{g}$ *is limit closed.*

Limit closure can be viewed as a continuity condition and, as we have shown, it is essentially this continuity that is necessary for knowledge structures to characterize the agents' knowledge. Since the results of [BE79, MZ85] show that all belief structures characterize the agents' beliefs, we would expect there to be some continuity condition implicit in the construction of belief structures. As Lipman [Lip91] observed, the Δ operator used in constructing belief structures can be viewed as a continuous operator. As we now show, this continuity arises from the fact that probability measures are assumed to be countably additive. (We remark that the view of countable additivity as a continuity condition is quite standard.) Limit closure says that if all finite prefixes of a knowledge structure are considered possible, then so is the knowledge structure itself. Analogously, it follows from countable additivity that the probabilities of the finite prefixes of a belief structure determine the probability of the belief structure.

We now show that without countable additivity of probability measures, it would not necessarily be the case that such (modified) belief structures completely characterize the agents' beliefs. In particular, we give an example where the probability measure is only finitely additive, rather than countably additive, and where the resulting belief structure does *not* completely characterize the agents' beliefs. The definition of belief structures remains unchanged, except that we now allow probability measures that are only finitely additive, and not necessarily countably additive.

Example 4.5: Our example is a variant of Example 3.10. Again we have two possible states of nature, p and $\bar{p}$, and three agents, 1, 2, and 3. Agents 1 and 3 find out whether or not p is true, while 2 does not. Initially, agent 2 considers p and $\bar{p}$ equally likely. If p is true, then agents 1 and 2 start to communicate. Suppose that it is common knowledge that agents 1 and 2 assign probability 1/2 to the $(k+1)^{\text{st}}$ message arriving, given that k messages have arrived, while agent 3 assigns probability 1 to the $(k+1)^{\text{st}}$ message arriving, given that k messages arrive. Intuitively, agent 3's beliefs are incompatible with those of agents 1 and 2. Moreover, we assume (quite unrealistically!) that (it is commonly known that) the time for the k^{th} message to arrive is $1/2^k$. Thus, all communication has ended by time 1. We now consider the agents' beliefs at time 1.

In a fashion analogous to Example 3.10, we can construct finite prefixes of belief structures by induction on length. In fact, there is a one-to-one correspondence between the prefixes of length k of belief structures that now arise and the k-worlds that we constructed in Example 3.10. Again, there are two prefixes of length 1, namely $\langle \bar{p} \rangle$ and $\langle p \rangle$, which we now denote $v_{1,-1}$ and $v_{1,0}$ respectively. There are three possible prefixes of length 2, analogous to the three 2-worlds in

Example 3.10. More generally, for all $k \geq 1$, there are precisely $k+1$ prefixes of length k that form a support for all the probability measures that arise; we denote these $v_{k,-1}, v_{k,0}, \ldots, v_{k,k-1}$. Suppose $v_{k,j} = \langle b_0, \ldots, b_{k-1} \rangle$. If $j = -1$, then $b_{k-1}(1)$ and $b_{k-1}(3)$ both assign probability 1 to $v_{k-1,-1}$, while $b_{k-1}(2)$ places probability 1/2 on each of $v_{k-1,-1}$ and $v_{k-1,0}$. If $0 \leq j \leq k-1$, then $b_{k-1}(3)$ places probability 1 on $v_{k-1,k-2}$ (recall that agent 3 assigns probability 1 to every message arriving). If $0 \leq j < k-1$ and j is even, then $b_{k-1}(1)$ places probability 1/2 on each of $v_{k-1,j}$ and $v_{k-1,j+1}$, while $b_{k-1}(2)$ places probability 1/2 on each of $v_{k-1,j-1}$ and $v_{k-1,j}$. The situation is similar if j is odd. Finally, if $j = k-1$, then $b_{k-1}(1)$ and $b_{k-1}(2)$ both place probability 1 on $v_{k-1,k-2}$.

So far all the measures that have arisen have had finite support. In this case there is no distinction between finitely and countably additive measures. The difference arises when we consider complete belief structures. We have belief structures $\mathbf{b}_{-1}, \mathbf{b}_0, \mathbf{b}_1, \ldots, \mathbf{b}_\infty$ that are the obvious analogues to $\mathbf{f}_{-1}, \mathbf{f}_0, \mathbf{f}_1, \ldots, \mathbf{f}_\infty$. Suppose $j \geq 0$. Is there a unique extension of $\mathbf{b}_j$ of length $\omega + 1$? If we consider only countably additive probability measures, then the results of [BE79, MZ85] tell us that there is. In fact, it is easy to see directly that this is so. For suppose that $\langle \mathbf{b}_j, b_\omega \rangle$ is such an extension. Recall that when we considered extensions of $\mathbf{f}_j$, what caused problems was $f_\omega(3)$. But if we consider countably additive measures, then $b_\omega(3)$ is determined. Since agent 3 is certain that all messages arrive, $b_\omega(3)$ places probability 1 on the belief structure $\mathbf{b}_\infty$, since it must place probability 0 on $\mathbf{b}_j$ for $j \neq \infty$, by the consistency constraints. This is no longer the case if we move to finitely additive measures. For each $\alpha \in [0,1]$, we now show that there is a finitely additive measure that places probability 0 on $\mathbf{b}_j$ for all $j \neq \infty$ and probability α on $\mathbf{b}_\infty$. This follows from the well-known result that there exists a finitely additive probability measure on the integers that assigns probability 0 to each finite subset. For completeness, we sketch the proof here.

Given a set U, a *filter* $\mathcal{F}$ on U is a nonempty set of subsets of U such that

1. $\emptyset \notin \mathcal{F}$,
2. $\mathcal{F}$ is closed under finite intersections, so that if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$,
3. $\mathcal{F}$ is closed under supersets, so that if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$.

An *ultrafilter* on U is a filter that is not a proper subset of any other filter. It is easy to show that if $\mathcal{F}$ is an ultrafilter on U and if $A \subseteq U$, then either A or its complement is in $\mathcal{F}$, but not both [BS74, Lemma 3.1]. It is also easy to see that if $u \in U$, then $\mathcal{F}_u = \{A \subseteq U : u \in A\}$ is an ultrafilter on U. Ultrafilters of the form $\mathcal{F}_u$ are called *principal ultrafilters*. It is well-known that every infinite set has a nonprincipal ultrafilter [BS74, Lemma 3.8], that is, an ultrafilter where no member is a singleton set.

Let $\mathcal{F}$ be a nonprincipal ultrafilter on the set $B = \{\mathbf{b}_{-1}, \mathbf{b}_0, \mathbf{b}_1, \ldots\}$. We define a function μ_α on $B^+ = B \cup \{\mathbf{b}_\infty\}$ as follows. If $A \subseteq B^+$, then

$$\mu_\alpha(A) = \begin{cases} 0 & \text{if } \mathbf{b}_\infty \notin A, \text{ and } A \notin \mathcal{F}, \\ 1-\alpha & \text{if } \mathbf{b}_\infty \notin A \text{ and } A \in F, \\ \alpha & \text{if } \mathbf{b}_\infty \in A \text{ and } A - \{\mathbf{b}_\infty\} \notin F, \\ 1 & \text{if } \mathbf{b}_\infty \in A \text{ and } A - \{\mathbf{b}_\infty\} \in F. \end{cases}$$

We leave it to check that μ_α is indeed a finitely additive probability measure. Clearly $\mu_\alpha(\{\mathbf{b}_\infty\}) = \alpha$ and $\mu_\alpha(\{\mathbf{b}_j\}) = 0$ for $j \neq \infty$. Finally, it is easy to see that there is an extension $\langle \mathbf{b}_j, b_\omega^\alpha \rangle$ of $\mathbf{b}_j$ such that $b_\omega^\alpha(3) = \mu_\alpha$. In particular, it follows that if we consider finitely additive probabilities, then the first ω levels of a belief structure do *not* characterize the agents' beliefs. ∎

5 Adequacy revisited

We have seen that, in general, knowledge structures do not characterize the agents' knowledge. How serious a problem is this? That depends on the events we are interested in. As shown in [FHV91], if we are interested only in common knowledge of events, then knowledge structures are indeed adequate, even if they do not characterize the agents' knowledge. But having the same common knowledge is not the same as having the same information. For more complicated events, we need to go further out in the hierarchy. These questions are addressed in [FHV91] in a logic-theoretic framework; we reconsider them here in an event-based setting. Our results also give us a better understanding of the relationship between Aumann structures and knowledge structures.

One way to approach the adequacy issue is to consider an Aumann structure with $\mathcal{F}(S)$, the set of knowledge structures over S, the set of states of nature, as its state space. (Brandenburger and Dekel [BD93] use an analogous construction, except, for them, the state space of the Aumann structure is the set of belief structures over S.) Given, however, that knowledge structures do not completely describe the agents' knowledge, it does not seem right to take the state space to be $\mathcal{F}(S)$. Instead, we consider a more general framework. Let us consider an Aumann structure with state space T such that every state $t \in T$ is associated with a knowledge structure $\mathbf{f}^t \in \mathcal{F}(S)$. Intuitively, we can think of the knowledge structure $\mathbf{f}^t$ as defining the agents' knowledge at state t, through the first ω levels. Let $\tau : T \rightarrow \mathcal{F}(S)$ be the mapping such that $\tau(t) = \mathbf{f}^t$. We allow a knowledge structure to be associated with more than one state; since, as we have shown, knowledge structures do *not* in general completely characterize the agents' knowledge, there may be two states of the world where the agents' knowledge through the first ω levels are identical, although the agents' knowledge differ in the two states. We say that a partition $\mathcal{K}_i$ of T is *coherent* (with respect to τ) if, for every state $t \in T$, the set of knowledge structures associated with the states in $\mathcal{K}_i(t)$ form a coherent set of possibilities for agent i. Intuitively, since the knowledge structures associated with the states describe the finite levels of knowledge of the agents, we would expect the partitions to respect this knowledge and therefore be coherent. We say that $A = (T, \mathcal{K}_1, \ldots, \mathcal{K}_n)$ is a *coherent Aumann structure based on* (S, T, τ) if each of $\mathcal{K}_1, \ldots, \mathcal{K}_n$ is coherent with respect to τ. We may still have a lot of freedom in defining partitions in a coherent Aumann structure. We now examine the effect of defining different partitions. Our goal is to understand whether defining different partitions of T can affect the knowledge of the agents in the resulting Aumann structures.

Note that there are two state spaces involved in Aumann structures where the states are associated with knowledge structures from $\mathcal{F}(S)$: the state space S for the knowledge structures and the state space T for the Aumann structure. We identify an event $E \subseteq S$ with the set of all states $t \in T$ such that the state of nature in $\mathbf{f}^t$ is in E.

Let $A_1 = (T, \mathcal{K}_1, \ldots, \mathcal{K}_n)$ and $A_2 = (T, \mathcal{K}'_1, \ldots, \mathcal{K}'_n)$ be two coherent Aumann structures based on (S, T, τ). Assume $E \subseteq S$. *A priori*, the event $\mathsf{C}(E)$ could be different in A_1 and A_2, since in A_1, we use the partition $\mathcal{K}_i$ to determine the K_i operator, whereas in A_2 we use $\mathcal{K}'_i$. Intuitively, since knowledge structures do not characterize the knowledge of the agents, different partitions may result in different common knowledge by the agents. We use the notation C^A (and, similarly, K_i^A) when we want to emphasize that we are considering the operators C and K_i determined by the partitions in Aumann structure A. The next theorem shows that if A_1 and A_2 are both coherent (and use the same association of states to knowledge structures), then $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$.

Theorem 5.1: *If A_1 and A_2 are coherent Aumann structures based on (S, T, τ) and $E \subseteq S$ then $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$.*

Proof: This result follows immediately from Theorem 5.2 below. ∎

Theorem 5.1 can be viewed as saying that there is a precise sense in which knowledge structures do completely characterize the common knowledge that agents have regarding events defined by subsets of S.

We are often interested, however, not just in common knowledge of events defined by subsets of S, but in common knowledge of more complicated events. For example, we might be interested in the fact that it is common knowledge that agent 3 does not know that a message was sent from agent 1 to agent 2. If the state space S is defined by events of the form "a message was sent from agent i to agent j", then typically the event "agent 3 does not know that a message was sent from agent 1 to agent 2" is not an event in S, so Theorem 5.1 does not apply. Furthermore, common knowledge is just one aspect of an agent's information. Agent 1 might know that agent 2 knows that a message arrived, without this fact being common knowledge. Nevertheless, this could well be an important piece of information. We can strengthen the previous result so that it deals with common knowledge of events that are not necessarily defined by subsets of S, and also deals with knowledge that is not common knowledge.

Suppose A is a coherent Aumann structure based on (S, T, τ). We can define the *ck-events over S in A*, denoted $ck_A(S)$, as the result of starting with the events defined by subsets of S, and then closing off under complementation, finite intersection, and the knowledge and common knowledge operators.

Theorem 5.2: *If A_1 and A_2 are coherent Aumann structures based on (S, T, τ) then $ck_{A_1}(S) = ck_{A_2}(S)$. Moreover, if $E \in ck_{A_1}(S)$ and i is an agent, then $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$ and $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$.*

Proof: See Appendix B. ∎

Theorem 5.1 tells us that knowledge structures characterize common knowledge that agents have regarding events defined by subsets of S. Theorem 5.2 tells us even more; knowledge structures in fact characterize knowledge and common knowledge of more complicated events, obtained from events that are subsets of S and closing under complementation, finite intersection, and the knowledge and common knowledge operators. As was suggested in Example 3.10, the situation changes when we consider common knowledge among coalitions of agents. We can define a *coalition common knowledge operator* C_G in Aumann structures, for every coalition G of agents, along the same lines as we defined the common knowledge operator. Namely, we define the operator O_G ("everyone in coalition G knows") on events by taking $\mathsf{O}_G(E)$ to be the intersection over $i \in G$ of the events $\mathsf{K}_i(E)$. The event $\mathsf{C}_G(E)$ is then the intersection of the events $\mathsf{O}_G(E)$, $\mathsf{O}_G(\mathsf{O}_G(E))$, and so on. The common knowledge operator C is the special case where G is taken to be all the agents.

Given an Aumann structure A_1 as above, we can define the *cck-events of S in A_1*, denoted $cck_{A_1}(S)$, to be the result of closing off the sets of events also under the coalition common knowledge operators. As Example 3.10 suggests, Theorem 5.2 fails if we replace the ck-events by the cck-events.[4]

We can get an analogue to Theorem 5.2 if we carry the construction of the hierarchy somewhat further into the ordinals. That is, we must consider λ-worlds, for $\lambda > \omega$. As we now show, if all we care about are the cck-events, then it suffices to take $\lambda = \omega^2$.

[4]As the proof of Theorem 5.2 in Appendix B shows, the problem lies in the combination of complementation and coalition common knowledge. If we did not close off the ck and cck events under complementation, then Theorem 5.2 would hold even with coalition common knowledge. In particular, if we are interested in a statement that agent 3 knows that it is common knowledge between agents 1 and 2 that E holds (as opposed to agent 3 knowing that it is *not* common knowledge between agents 1 and 2 that E holds, as is the case in Example 3.10), then the analogue of Theorem 5.2 does hold.

Let $\mathcal{F}^{\omega^2}(S)$ consist of all ω^2-worlds over S. We say that $\mathcal{P} \subseteq \mathcal{F}^{\omega^2}(S)$ is an *ω^2-coherent set of possibilities* for agent i at $\mathbf{f} \in \mathcal{F}^{\omega^2}(S)$ if, as before, $\mathbf{f} \in \mathcal{P}$ and $\mathcal{P} \subseteq \mathbf{f}^{\sim_i}$, and the obvious extension of P3 to level ω^2 holds: namely, if $w \in f_\alpha(i)$ for some $\alpha < \omega^2$, then there is an ω^2-world $\mathbf{f}' \in \mathcal{P}$ such that w is a prefix of $\mathbf{f}'$. We now consider Aumann structures each of whose states is associated with a knowledge structure in $\mathcal{F}^{\omega^2}(S)$. We say that $A = (T, \mathcal{K}_1, \ldots, \mathcal{K}_n)$ is an *ω^2-coherent Aumann structure based on* (S, T, τ) if each of $\mathcal{K}_1, \ldots, \mathcal{K}_n$ is ω^2-coherent (using the obvious definition of ω^2-coherent for partitions).

Theorem 5.3: *If A_1 and A_2 are ω^2-coherent Aumann structures based on (S, T, τ) then $cck_{A_1}(S) = cck_{A_2}(S)$. Moreover, if $E \in cck_{A_1}(S)$, if i is an agent, and if G is a group of agents, then $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$ and $\mathsf{C}_G^{A_1}(E) = \mathsf{C}_G^{A_2}(E)$.*

Proof: See Appendix B. ∎

We remark that it can be shown that we actually need to consider structures of length ω^2 in order to get a result such as Theorem 5.3 (cf. Theorem 5.14 in [FHV91]); that is, no smaller length suffices.

The results of this section help explain the apparent inconsistency mentioned in the introduction that, in spite of the connection between knowledge structures and Aumann structures in terms of knowledge and common knowledge, it still happens that knowledge structures are not an adequate description of an agent's knowledge. The point is that there are richer notions than simply knowledge and common knowledge, such as coalition common knowledge, that knowledge structures do not capture.

6 Counterfactual information

The focus so far has been on the issue of how much knowledge is captured by knowledge structures. As we observed in the introduction, however, knowledge structures also seem to be deficient in another manner, since they capture only worlds that are commonly known to be possible, while omitting worlds that are merely conceivable (such as ones where Ron Fagin is President). Clearly, this deficiency is orthogonal to the issue of the length of the hierarchy; it is a function of the definition of knowledge structures, and not of the hierarchical approach. We now show how a generalization of knowledge structures can capture counterfactual information.

There has been a great deal of work done on modeling counterfactuals [Lew73]. We present here a somewhat naive version of the standard approach. Our goal is not to provide a sophisticated model of counterfactuals, but to show that counterfactuals can be dealt with using the hierarchical approach.

The basic idea is to augment the definition of knowledge assignments. As before, a *0th-order extended knowledge assignment* f_0 is a member of S, that is, a state of nature (which, intuitively, corresponds to the "real world"). We call $\langle f_0 \rangle$ an *extended 1-world.* Assume inductively that extended κ-worlds have been defined for all κ with $1 \leq \kappa < \lambda$. Let U_κ be the set of all extended κ-worlds, for $\kappa < \lambda$. If $1 \leq \kappa < \lambda$, a *κth-order extended knowledge assignment* is a pair $f_\kappa = (f_\kappa^p, f_\kappa^c)$ of functions that associates with each agent i a set $f_\kappa^p(i) \subseteq U_\kappa$ of "possible" extended κ-worlds, and a set $f_\kappa^c(i) \subseteq U_\kappa$ of "conceivable" extended κ-worlds such that $f_\kappa^p(i) \subseteq f_\kappa^c(i)$. Intuitively, the conceivable worlds include not only the possible worlds, but also those that the agent does not consider possible (such as a world where Ron Fagin is President). If λ is a limit ordinal, an extended λ-world is a sequence $\mathbf{f} = \langle f_0, f_1, \ldots \rangle$ of length λ such that for each $\kappa < \lambda$, we have that

f_κ is a κth-order extended knowledge assignment and each *κ-prefix* (i.e., prefix of length κ) is an extended κ-world. If $\lambda = \lambda' + 1$, there are again some consistency conditions that $f_{\lambda'}$ must satisfy, which extend the consistency conditions that knowledge worlds are required to obey.

What are the consistency conditions? Since f^p_κ is now playing essentially the same role as f_κ did before, we require the following analogues of the original consistency conditions K1–K3:

K1′. $\mathbf{f}_{<\lambda'} \in f^p_{\lambda'}(i)$.

K2′. If $\langle g_0, g_1, \ldots\rangle \in f^p_{\lambda'}(i)$, then $g_\kappa(i) = f^p_\kappa(i)$ for all κ with $0 < \kappa < \lambda'$.

K3′. If $0 < \kappa < \lambda'$, then $\mathbf{g} \in f^p_\kappa(i)$ iff there is some $\mathbf{h} \in f^p_{\lambda'}(i)$ such that $\mathbf{g} = \mathbf{h}_{<\kappa}$.

We also require, as we stated above, that

K4′. $f^p_{\lambda'}(i) \subseteq f^c_{\lambda'}(i)$.

What about the analogues of conditions K1′–K3′ for $f^c_{\lambda'}$? The analogue of condition K1′ holds automatically, since $f^p_{\lambda'}(i) \subseteq f^c_{\lambda'}(i)$. We require that the following analogue of condition K3′ hold:

K3″. If $0 < \kappa < \lambda'$, then $\mathbf{g} \in f^c_\kappa(i)$ iff there is some $\mathbf{h} \in f^c_{\lambda'}(i)$ such that $\mathbf{g} = \mathbf{h}_{<\kappa}$.

The reason we require condition K3″ is that we think of each level as giving a finer and finer description. We do not necessarily require that the analogue of condition K2′ hold. Intuitively, condition K2′ says that the agents are introspective, and we do not require an agent to be introspective when considering conceivable worlds. Of course, we could easily impose such conditions on extended knowledge assignments, as well as further conditions to capture more sophisticated counterfactual information.

In this extended setting, we can again ask how far out into the ordinals we need to go. And, just as before, this will depend on the events of interest. There are cases when ω levels suffice, and others where we need to go much farther out into the ordinals.

7 Related results

There are a number of results of Fagin [Fag94] and Heifetz and Samet [HS93, HS98] that are related to ones proved here; we briefly describe them in this section.

It is shown in [FHV91] that for every Aumann structure[5] A and state s of A, and for every ordinal λ, there is a λ-world that in a precise sense captures the knowledge of the agents through level λ at the state s. We say that the state s is *represented* by this λ-world. Fagin [Fag94] defines the *distinguishing ordinal* of an Aumann structure A to be the least ordinal γ such that whenever s and t are states of A that are represented by the same γ-world, then s and t are represented by the same λ-world for every λ. Heifetz and Samet refer to the distinguishing ordinal as the *order of the partition space* in [HS93] and as the *rank of the partition space* in [HS98]. Roughly speaking, we can think of the distinguishing ordinal of A as describing how far out in the knowledge hierarchy we need to go to completely describe the knowledge of agents in a state of A.

Fagin and, independently, Heifetz and Samet, showed the following result. Let γ be an infinite ordinal with cardinality κ (for example, if γ is a countable ordinal, then κ is $\aleph_0$). Then there is

[5] The paper actually uses Kripke structures rather than Aumann structures, as does [Fag94], but in the S5 case, the results convert easily into results about Aumann structures.

an Aumann structure with at most κ states and with distinguishing ordinal γ. For example, if γ is a countable ordinal, then the corresponding Aumann structure can be taken to have a countable state space. Heifetz and Samet prove the result by giving an elegant explicit example, the "Sobers-Drunks Example". This result shows that, in general, there is no bound on how far in the hierarchy we have to go to describe the agents' knowledge.

Describing the agents' knowledge in a state of an Aumann structure is not the same as characterizing it in the sense defined in this paper. Let A be an Aumann structure and let s be a state of A. Define the *uniqueness ordinal*[6] of the Aumann structure A to be the least ordinal μ such that if s is a state of A and $\mathbf{f}$ is the μ-world that represents s, then $\mathbf{f}$ characterizes the agents' knowledge. *A priori*, it is not at all clear that such an ordinal μ exists. However, Fagin and, independently, Heifetz and Samet, showed that indeed, every Aumann structure has a uniqueness ordinal. In fact, if γ is the distinguishing ordinal and μ is the uniqueness ordinal, then $\gamma \leq \mu \leq \gamma + \omega$.

Since the distinguishing ordinal may be arbitrarily large, and since the uniqueness ordinal is at least as big as the distinguishing ordinal, it follows that the uniqueness ordinal may be arbitrarily large. This implies immediately the result stated (informally) in the introduction that no ordinal level of knowledge is sufficiently large to describe completely an agent's uncertainty.

8 Concluding Remarks

As we have seen, the question of how far we have to extend the hierarchy to capture the agents' knowledge is a somewhat subtle one. Although the results of [BE79, MZ85] show that ω levels suffice for belief structures, this result depends on countable additivity of probability functions, and does not hold if we consider knowledge rather than belief defined probabilistically.[7] On the other hand, our results show that, even if we need to go possibly far beyond ω levels to (completely) characterize the agents' knowledge, for many events of interest, ω (or ω^2) levels suffice.

It could be argued that knowledge structures and knowledge worlds as defined here are perhaps not the closest non-probabilistic analogue to belief structures.[8] A somewhat closer analogue would result if we replaced the correctness requirement K1 by the much weaker requirement $f_{\lambda'}(i) \neq \emptyset$. This would result in a notion closer to the traditional philosopher's notion of belief.[9] The arguments given here apply without change to show that countable hierarchies still do not suffice if we use this nonprobabilistic notion of belief. Indeed, the arguments of [Fag94] show that, in general, we need to again go to arbitrarily large ordinals to characterize the beliefs of agents. The key point is that without some sort of continuity condition (which probability gives us), or states of nature rich enough to determine the agents' knowledge, or other equally strong conditions, we need to go well beyond the first ω levels in general to characterize an agent's knowledge or belief.

A Proofs for Section 3

Before we prove Theorems 3.4 and 3.5, we need another lemma. We say that two $(\lambda+1)$-worlds $\langle f_0, \ldots, f_\lambda \rangle$ and $\langle g_0, \ldots, g_\lambda \rangle$ *differ on agent* i if $g_\lambda(i) \neq f_\lambda(i)$.

[6] This is the term used by Fagin. Heifetz and Samet define this notion in terms of "knowledge morphisms".

[7] We remark that [Lip91] gives yet another example of a context in which we need to go beyond ω levels.

[8] We thank one of the referees of the paper for bringing this point to our attention.

[9] Technically, it would satisfy the axioms of the modal logic KD45 rather than S5; see [FHMV95] for a discussion of these notions.

Lemma A.1: *Let i and j be distinct agents. Let $\mathbf{f}$ and $\mathbf{g}$ be λ-worlds (not necessarily distinct) such that $\mathbf{f} \sim_i \mathbf{g}$. Assume that there is a $(\lambda+1)$-world $\langle \mathbf{g}, g_\lambda \rangle$ extending $\mathbf{g}$ that differs from $\mathbf{g}^+$ on agent j. Then there is a $(\lambda+2)$-world $\langle \mathbf{f}^+, f'_{\lambda+1} \rangle$ extending $\mathbf{f}^+$ that differs from $\mathbf{f}^{++} = (\mathbf{f}^+)^+$ on agent i.*

Proof: Suppose $\mathbf{g}' = \langle \mathbf{g}, g_\lambda \rangle$ is a $(\lambda+1)$-world such that $g_\lambda(j)$ is not the j-no-information extension of $\mathbf{g}$. Without loss of generality, we can assume that $g_\lambda(i)$ is $\mathbf{g}^{\sim_i}$. Let $\mathbf{f}^+ = \langle \mathbf{f}, f_\lambda \rangle$ and let $\mathbf{f}^{++} = \langle \mathbf{f}, f_\lambda, f_{\lambda+1} \rangle$. Clearly $\mathbf{f}^+ \sim_i \mathbf{g}'$. Therefore, $\mathbf{g}' \in f_{\lambda+1}(i)$, since $f_{\lambda+1}(i) = (\mathbf{f}^+)^{\sim_i}$. Furthermore, $\mathbf{g}' \neq \mathbf{f}^+$. Define $f'_{\lambda+1}$ so that $f'_{\lambda+1}(j) = f_{\lambda+1}(j)$ for $j \neq i$, and $f'_{\lambda+1}(i) = f_{\lambda+1}(i) - \{\mathbf{g}'\}$. We now show that $\mathbf{f}' = \langle \mathbf{f}^+, f'_{\lambda+1} \rangle$ is a $(\lambda+2)$-world. The correctness condition K1 holds, since $\mathbf{g}' \neq \mathbf{f}^+$; the introspection condition K2 is immediate; and the extendibility condition K3 holds, since $\mathbf{g}^+$ is in $f'_{\lambda+1}(i)$ and has the same λ-prefix $\mathbf{g}$ as $\mathbf{g}'$. Clearly $\mathbf{f}'$ differs from $\mathbf{f}^{++}$ on agent i. ∎

We can now prove Theorems 3.4 and 3.5. We prove the results simultaneously. We repeat the statements for the convenience of the reader.

Theorem 3.4: *A knowledge structure $\mathbf{f}$ characterizes the agents' knowledge iff every knowledge structure reachable from $\mathbf{f}$ characterizes the agents' ω-knowledge.*

Theorem 3.5: *A knowledge structure characterizes the agents' knowledge iff it characterizes the agents' knowledge through the first $\omega + \omega$ levels.*

Proof of Theorems 3.4 and 3.5: We shall show that the following are equivalent:

1. $\mathbf{f}$ characterizes the agents' knowledge.
2. $\mathbf{f}$ characterizes the agents' knowledge through the first $\omega + \omega$ levels.
3. Every knowledge structure $\mathbf{g}$ reachable from $\mathbf{f}$ characterizes the agents' ω-knowledge.

(1) $\Rightarrow$ (2) is immediate.

(2) $\Rightarrow$ (3): We shall show that if $\mathbf{e}$ and $\mathbf{h}$ are knowledge structures with $\mathbf{e} \sim_i \mathbf{h}$, and if $\mathbf{h}$ does not characterize the agents' $(\omega+k)$-knowledge for some natural number k, then $\mathbf{e}$ does not characterize the agents' $(\omega+k+2)$-knowledge. It follows easily that if $\mathbf{g}$ is reachable from $\mathbf{f}$ by a path of length m and $\mathbf{g}$ does not characterizes the agents' ω-knowledge, then $\mathbf{f}$ does not characterize the agents' $(\omega+2m)$-knowledge. This is sufficient to prove that (2) $\Rightarrow$ (3).

So assume that $\mathbf{e} \sim_i \mathbf{h}$, and $\mathbf{h}$ does not characterize the agents' $(\omega+k)$-knowledge for some natural number k. Without loss of generality, we can assume that k is minimal with this property. Since $\mathbf{h}$ does not characterize the agents' $(\omega+k)$-knowledge, we know that there are two distinct $(\omega+k+1)$-worlds $\mathbf{h}_1$ and $\mathbf{h}_2$ extending $\mathbf{h}$. Now one $(\omega+k+1)$-world extending $\mathbf{h}$ is obtained from $\mathbf{h}$ by applying the one-step no-information extension $k+1$ times; without loss of generality, we can assume that $\mathbf{h}_1$ is this $(\omega+k+1)$-world. Since k is minimal, there is an $(\omega+k)$ world $\mathbf{h}'$ extending $\mathbf{h}$ such that $\mathbf{h}_1 = \langle \mathbf{h}', h_{\omega+k} \rangle$, $\mathbf{h}_2 = \langle \mathbf{h}', h'_{\omega+k} \rangle$, and some agent j such that $h_{\omega+k}(j) \neq h'_{\omega+k}(j)$.

Let $\mathbf{e}'$ be $\mathbf{e}$ if $k=0$, and the result of applying to $\mathbf{e}$ the one-step no-information extension k times if $k>0$. Notice that $\mathbf{e}' \sim_i \mathbf{h}'$. We note for later use that, therefore, $(\mathbf{e}')^+ \sim_i (\mathbf{h}')^+ = \mathbf{h}_1$. There are two cases, depending on whether or not $j = i$.

Assume first that $j \neq i$. By Lemma A.1, where the roles of $i, j, \mathbf{f}, \mathbf{g}$ are played by $i, j, \mathbf{e}', \mathbf{h}'$ respectively, it follows that there are two distinct $(\omega+k+2)$-worlds extending $(\mathbf{e}')^+$. Therefore, $\mathbf{e}$

does not characterize the agents' $(\omega+k+1)$-knowledge. Hence, $\mathbf{e}$ does not characterize the agents' $(\omega+k+2)$-knowledge, which was to be shown.

Now assume that $j=i$. We can assume that there are at least two agents, since otherwise, by Proposition 3.2, the knowledge structure is completely determined by its first two levels. Let ℓ be some agent other than agent i. We apply Lemma A.1, where the roles of $i, j, \mathbf{f}, \mathbf{g}$ are played by $\ell, i, \mathbf{h}', \mathbf{h}'$ respectively, and find that $(\mathbf{h}')^{+}=\mathbf{h}_1$ has two extensions to $(\omega+k+2)$-worlds that differ on agent ℓ. We apply Lemma A.1 again, where the roles of $i, j, \mathbf{f}, \mathbf{g}$ are played by $i, \ell, (\mathbf{e}')^{+}, (\mathbf{h}')^{+}$ respectively, and find that $(\mathbf{e}')^{++}$ has two extensions to $(\omega+k+3)$-worlds that differ on agent i. Hence, $\mathbf{e}$ does not characterize the agents' $(\omega+k+2)$-knowledge, as desired.

(3) $\Rightarrow$ (1): Let C be the set of knowledge structures reachable from $\mathbf{f}$. We now show that every member of C characterizes the agents' knowledge. If not, then let λ be the minimal infinite ordinal such that there is some $\mathbf{g} \in C$ that does not characterize the agents' λ-knowledge. By assumption, $\lambda > \omega$. Let $\mathbf{g}' = \langle \mathbf{g}, g_\omega, \ldots, g_\lambda \rangle$ be a $(\lambda+1)$-world extending $\mathbf{g}$. Let i be an arbitrary agent. We must have $g_\omega(i) = \mathbf{g}^{\sim_i}$, or else there would not be a unique extension of $\mathbf{g}$ to level ω. By extendibility, $g_\lambda(i)$ must contain an extension (of the appropriate length) of each knowledge structure $\mathbf{e} \in \mathbf{g}^{\sim_i}$, and only such extensions. But there is at most one such extension for each $\mathbf{e} \in \mathbf{g}^{\sim_i}$; this follows by definition of λ and the fact that $\mathbf{e} \in C$ (since $\mathbf{e} \sim_i \mathbf{g}$). So $g_\lambda(i)$ is uniquely determined. Since this is true for each agent i, it follows that g_λ is uniquely determined, a contradiction. ∎

The next lemma is a refinement of Theorem 3.3. The proof is obtained in a straightforward manner from the proof of Theorem 3.3.

Lemma A.2: *A knowledge structure* $\mathbf{f}$ *characterizes agent* i*'s* ω*-knowledge iff for each knowledge structure* $\mathbf{g} \neq \mathbf{f}$ *such that* $\mathbf{f} \sim_i \mathbf{g}$*, some finite prefix of* $\mathbf{g}$ *is* i*-uniquely extendible w.r.t.* $\mathbf{f}$.

We now prove Theorem 3.6, which is a third characterization of knowledge structures that characterize the agents' knowledge, in the case where the state space S is finite. We first need a definition that slightly refines the notion of a knowledge structure characterizing the agents' ω-knowledge. Let us say that a knowledge structure $\mathbf{f}$ *characterizes agent* i*'s* ω*-knowledge* if whenever $\langle \mathbf{f}, f_\omega \rangle$ and $\langle \mathbf{f}, f'_\omega \rangle$ are extensions of $\mathbf{f}$ to an $(\omega+1)$-world, then $f_\omega(i) = f'_\omega(i)$. Intuitively, this says that there is a unique possible value for $f_\omega(i)$. Clearly, a knowledge structure characterizes the agents' ω-knowledge iff it characterizes agent i's ω-knowledge for each agent i. The next lemma will be useful in our new characterization of when a knowledge structure characterizes the agents' This lemma is due to R. Simon (personal communication).

Lemma A.3: *Assume that there are only finitely many states. Let* $\mathbf{f}$ *be a knowledge structure and* i *an agent. Then* $\mathbf{f}^{\sim_i}$ *is finite iff every member of* $\mathbf{f}^{\sim_i}$ *characterizes agent* i*'s* ω*-knowledge.*

Proof: Assume first that $\mathbf{f}^{\sim_i}$ is finite, and that $\mathbf{g} \in \mathbf{f}^{\sim_i}$. We must show that $\mathbf{g}$ characterizes agent i's ω-knowledge. By Lemma A.2, it suffices to show that for each knowledge structure $\mathbf{h} \neq \mathbf{g}$ such that $\mathbf{g} \sim_i \mathbf{h}$ (that is, such that $\mathbf{h} \in \mathbf{f}^{\sim_i}$), some finite prefix of $\mathbf{h}$ is i-uniquely extendible w.r.t. $\mathbf{g}$. Since $\mathbf{f}^{\sim_i}$ is finite, there is some positive integer k such that no two distinct members of $\mathbf{f}^{\sim_i}$ have the same k-prefix. Therefore, if $\mathbf{h} \in \mathbf{f}^{\sim_i}$, then the k-prefix of $\mathbf{h}$ is i-uniquely extendible w.r.t. $\mathbf{g}$. This was to be shown.

Conversely, assume that $\mathbf{f}^{\sim_i}$ is infinite; we must show that some member of $\mathbf{f}^{\sim_i}$ does not characterize agent i's ω-knowledge. Since $\mathbf{f}^{\sim_i}$ is infinite and S is finite, it follows by a König's Lemma argument that there is a sequence $w_1, w_2, w_3, \ldots$ of worlds, where for each k:

1. w_k is a k-world;
2. w_k is a prefix of w_{k+1}; and
3. w_k has infinitely many distinct extensions to knowledge structures in $\mathbf{f}^{\sim_i}$.

Let $\mathbf{g}$ be the knowledge structure whose k-prefix is w_k for each k. Clearly, $\mathbf{g} \in \mathbf{f}^{\sim_i}$. Since $\mathbf{f}^{\sim_i}$ is infinite, there is some member $\mathbf{h}$ of $\mathbf{f}^{\sim_i}$ such that $\mathbf{h} \neq \mathbf{g}$. We complete the proof by showing that $\mathbf{h}$ does not characterize agent i's ω-knowledge. By Lemma A.2, it suffices to show that no finite prefix of $\mathbf{g}$ is i-uniquely extendible w.r.t. $\mathbf{h}$. But this follows almost immediately from the definition of $\mathbf{g}$. ∎

Theorem 3.6: *Assume that there are only finitely many states. A knowledge structure* $\mathbf{f}$ *characterizes the agents' knowledge iff* $\mathbf{g}^{\sim_i}$ *is finite for every knowledge structure* $\mathbf{g}$ *reachable from* $\mathbf{f}$ *and every agent* i.

Proof: Assume first that the knowledge structure $\mathbf{f}$ characterizes the agents' knowledge, that $\mathbf{g}$ is reachable from $\mathbf{f}$, and that i is an agent. Then every member of $\mathbf{g}^{\sim_i}$ is reachable from $\mathbf{f}$. So by Theorem 3.4, every member of $\mathbf{g}^{\sim_i}$ characterizes the agents' ω-knowledge. So by Lemma A.3, it follows that $\mathbf{g}^{\sim_i}$ is finite.

Conversely, assume that $\mathbf{g}^{\sim_i}$ is finite for every knowledge structure $\mathbf{g}$ reachable from $\mathbf{f}$ and every agent i. By Lemma A.3, every member of $\mathbf{g}^{\sim_i}$, and in particular $\mathbf{g}$ itself, characterizes agent i's ω-knowledge. This shows that $\mathbf{g}$ characterizes the agents' ω-knowledge for every knowledge structure $\mathbf{g}$ reachable from $\mathbf{f}$. So by Theorem 3.4, it follows that $\mathbf{f}$ characterizes the agents' knowledge. ∎

B Proofs for Section 5

Theorem 5.2: *If* A_1 *and* A_2 *are coherent Aumann structures based on* (S, T, τ) *then* $ck_{A_1}(S) = ck_{A_2}(S)$. *Moreover, if* $E \in ck_{A_1}(S)$ *and* i *is an agent, then* $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$ *and* $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$.

Proof: We need some preliminary definitions and lemmas.

Definition B.1: Assume $R \subseteq T$. The *k-projection* of R, denoted R^k, consists of all the k-prefixes of the knowledge structures $\mathbf{f}^t$ such that $t \in R$, that is, $R^k = \{\mathbf{f}^t_{<k} \mid t \in R\}$. ∎

The operation of k-projection maps subsets of T to sets of k-worlds. We now define an operation mapping sets of k-worlds to subsets of T.

Definition B.2: If $B \subseteq W_k$, let $B^* = \{t \in T \mid \mathbf{f}^t_{<k} \in B\}$. A set $R \subseteq T$ is a *k-cylinder set* if $R = B^*$ for some set $B \subseteq W_k$. (Note that if R is a k-cylinder set, then we must in fact have $R = (R^k)^*$.) We just say *cylinder set* if we do not need to emphasize the k. ∎

Lemma B.3: *The set of cylinder sets is closed under finite union and complementation (and hence, finite intersection).*

Proof: Note that $\overline{B^*} = (\overline{B})^*$ (where $\overline{B}$ is the complement of B). Hence, the complement of a cylinder set is a cylinder set. To see that cylinder sets are closed under finite unions, first note that k-cylinder sets are closed under arbitrary union, for fixed k, since $(\cup_{j\in J} B_j)^* = \cup_{j\in J} B_j^*$, where J is an arbitrary index set and B_j is a k-cylinder set, for all $j \in J$. Next suppose that B is a k-cylinder set and C is an m-cylinder set, with $k \leq m$. Assume that $B' \subseteq W_m$ consists of all the m-worlds whose prefixes are in B. It is easy to see that $B^* = (B')^*$. Moreover, $B^* \cup C^* = (B' \cup C)^*$. Thus, the union of a k-cylinder set and an m-cylinder set is a $\max(k,m)$-cylinder set. ∎

Next we show that cylinder sets are also closed under applications of the K_i operator.

Lemma B.4: *If A_1 and A_2 are coherent Aumann structures based on (S,T,τ), and $R \subseteq T$ is a k-cylinder set, then $\mathsf{K}_i^{A_1}(R)$ is a $(k+1)$-cylinder set. Morever, $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_2}(R)$.*

Proof: Suppose $A_1 = (T, \mathcal{K}_1, \ldots, \mathcal{K}_n)$. By definition, $\mathsf{K}_i^{A_1}(R) = \{t \in T \mid \mathcal{K}_i(t) \subseteq R\}$. We claim that

$$\mathsf{K}_i^{A_1}(R) = \{t \in T \mid \mathbf{f}_k^t(i) \subseteq R^k\}. \tag{1}$$

Note first that

$$t \in \mathsf{K}_i^{A_1}(R) \text{ iff } \mathcal{K}_i(t) \subseteq R. \tag{2}$$

We now show that

$$\mathcal{K}_i(t) \subseteq R \text{ iff } \mathcal{K}_i(t)^k \subseteq R^k. \tag{3}$$

The fact that $\mathcal{K}_i(t) \subseteq R$ implies that $\mathcal{K}_i(t)^k \subseteq R^k$ follows from the general fact that if $A \subseteq B$, then $A^k \subseteq B^k$. The opposite implication depends on the fact that R is a k-cylinder set. Thus, assume that $\mathcal{K}_i(t)^k \subseteq R^k$, and that $t' \in \mathcal{K}_i(t)$; we must show that $t' \in R$. Since $t' \in \mathcal{K}_i(t)$, it follows that $\mathbf{f}_{<k}^{t'} \in \mathcal{K}_i(t)^k$, and hence $\mathbf{f}_{<k}^{t'} \in R^k$. Since R is a k-cylinder set, it follows, as noted earlier, that $R = (R^k)^* = \{t \in T \mid (\mathbf{f}^t)_{<k} \in R^k\}$. Hence $t' \in R$, as desired. This proves (3). Next, we show that

$$\mathcal{K}_i(t)^k = \mathbf{f}_k^t(i). \tag{4}$$

Since $\mathcal{A}_1$ is coherent, we have that the set of knowledge structures associated with states in $\mathcal{K}_i(t)$ is coherent, and so is a coherent set of possibilities for agent i at $\mathbf{f}^t$. Then (4) follows fairly easily from properties P2 and P3 of coherency: P2 implies $\mathcal{K}_i(t)^k \subseteq \mathbf{f}_k^t(i)$ and P3 implies $\mathbf{f}_k^t(i) \subseteq \mathcal{K}_i(t)^k$. Finally, (1) follows from (2), (3), and (4). Note that (1) already shows that $\mathsf{K}_i^{A_1}(R)$ is independent of the partition $\mathcal{K}_i$; it immediately follows that $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_2}(R)$. It is also easy to see that $\{t \in T \mid \mathbf{f}_k^t(i) \subseteq R^k\} = \{\langle w_0, \ldots, w_k\rangle \in W_{k+1} \mid w_k(i) \subseteq R^k\}^*$, showing that $\mathsf{K}_i^{A_1}(R)$ is a $(k+1)$-cylinder set. ∎

Although cylinder sets are closed under finite intersection, they are not closed under infinite intersection. Thus, even if R is a cylinder set, $\mathsf{C}^A(R)$ may not be.

Definition B.5: A set $R \subseteq T$ is *closed* iff R can be written as an arbitrary intersection of cylinder sets. ∎

Note that all cylinder sets are trivially closed. The following lemma shows that the set of closed subsets of T is closed under certain operations. (Note that in this lemma we use the coalition common knowledge operator C_G, which is defined immediately after Theorem 5.2.)

Lemma B.6: *Let A_1 and A_2 be coherent Aumann structures based on (S,T,τ).*

(a) *If J is an arbitrary index set, and if R_j is a closed subset of T for each j in J, then $\cap_{j\in J}R_j$ is also closed.*

(b) *If R is a closed subset of T, then $\mathsf{K}_i^{A_1}(R)$ and $\mathsf{C}_G^{A_1}(R)$ are both closed; moreover $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_2}(R)$ and $\mathsf{C}_G^{A_1}(R) = \mathsf{C}_G^{A_2}(R)$.*

Proof: Part (a) is immediate from the fact that closed subsets are (arbitrary) intersections of cylinder sets. For part (b), suppose that R is a closed subset of T. Then $R = \cap_j B_j$, where each B_j is a cylinder set. It is easy to check that $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_1}(\cap_j B_j) = \cap_j \mathsf{K}_i^{A_1}(B_j)$. (The final equality is a general fact, that does not depend on each B_j being a cylinder set; see part (1) of Lemma B.7 below.) By Lemma B.4, we know that $\mathsf{K}_i^{A_1}(B_j)$ is a cylinder set for all j; hence $\mathsf{K}_i^{A_1}(R)$ is closed, as desired. Moreover, Lemma B.4 tells us that $\mathsf{K}_i^{A_1}(B_j) = \mathsf{K}_i^{A_2}(B_j)$ for all j, and hence $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_2}(R)$. Since $\mathsf{O}_G^{A_h}(R) = \cap_{i\in G}\mathsf{K}_i^{A_h}(R)$ and $\mathsf{C}_G^{A_h}(R) = \cap_k(\mathsf{O}_G^{A_h})^k(R)$, for $h = 1, 2$, it easily follows (using part (a)) that $\mathsf{C}_G^{A_1}(R)$ is closed and that $\mathsf{C}_G^{A_1}(R) = \mathsf{C}_G^{A_2}(R)$. ∎

Notice that Lemma B.6 already suffices to prove Theorem 5.1. If we could only extend Lemma B.6 to show that closed sets were closed under complementation, we could then easily prove Theorem 5.2, even for events formed using the operator C_G for an arbitrary subset G of agents. The complement of a closed set is not, however, necessarily a closed set. In fact, Example 3.10 shows that Theorem 5.2 is false if we can use the operator C_G for any arbitrary subset G of agents.

Nevertheless, because it allows only C rather than C_G for arbitrary G, Theorem 5.2 is true. To prove it, we need a collection of sets that is closed under complementation, finite intersection, and the application of C and K_i. Before we define the appropriate notion, we collect a number of well-known general properties of the knowledge and common knowledge operators, the first of which we already used in the course of proving Lemma B.6. We leave the proof of these properties to the reader.

Lemma B.7: *The following properties hold in all Aumann structures:*

1. $\mathsf{K}_i(\cap_{j\in J}E_j) = \cap_{j\in J}\mathsf{K}_i(E_j)$ *for an arbitrary index set J.*
2. $\mathsf{C}_G(\cap_{j\in J}E_j) = \cap_{j\in J}\mathsf{C}_G(E_j)$ *for an arbitrary index set J.*
3. $\mathsf{C}_G(\overline{\mathsf{C}_G(E)}) = \overline{\mathsf{C}_G(E)}$
4. $\mathsf{K}_i((\mathsf{C}(E) \cup E') = \mathsf{C}(E) \cup \mathsf{K}_i(E')$
5. $\mathsf{C}_G(\mathsf{C}_G(E) \cup E') = \mathsf{C}_G(E) \cup \mathsf{C}_G(E')$

Note that part (4) of Lemma B.7 is the only one that holds for C but not for C_G (if i is not in G). The failure of part (4) for C_G is the reason that Theorem 5.2 fails for coalition common knowledge.

Since $\mathsf{K}_i(\emptyset) = \emptyset$ and $\mathsf{C}(\emptyset) = \emptyset$, we obtain as a special case of parts (4) and (5) of Lemma B.7 (taking $E' = \emptyset$) that $\mathsf{K}_i((\mathsf{C}(E)) = \mathsf{C}(E)$ and $\mathsf{C}(\mathsf{C}(E)) = \mathsf{C}(E)$.

Definition B.8: A subset $D \subseteq T$ is a *fixedpoint set* if in every coherent Aumann structure A based on (S, T, τ), we have $\mathsf{C}^A(D) = D$. A subset $E \subseteq T$ is *safe* if it has the form $\cap_{j=1}^k(D_j \cup R_j)$, where $D_1, \ldots, D_k$ are fixedpoint sets and $R_1, \ldots, R_k$ are cylinder sets. ∎

Safe sets give us what we need.

Proposition B.9: *If A_1 and A_2 are coherent Aumann structures based on (S,T,τ), and E, E_1, and E_2 are safe, then so are $E_1 \cap E_2$, $\overline{E}$, $\mathsf{K}_i^{A_1}(E)$, and $\mathsf{C}^{A_1}(E)$. Moreover, $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$ and $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$.*

Proof: Lemma B.3 says that cylinder sets are closed under finite union and complementation. By parts (3) and (5) of Lemma B.7, so are fixedpoint sets. Straightforward manipulation, using standard properties of complementation and union, shows that safe sets are also closed under finite union and complementation.

Suppose $E = \cap_{j=1}^k (D_j \cup R_j)$ is a safe set. By parts (1) and (4) of Lemma B.7, $\mathsf{K}_i^{A_1}(E) = \cap_{j=1}^k \mathsf{K}_i^{A_1}(D_j \cup R_j) = \cap_{j=1}^k (D_j \cup \mathsf{K}_i^{A_1}(R_j))$. By Lemma B.4, $\mathsf{K}_i^{A_1}(R_j)$ is a cylinder set if R_j is, so safe sets are closed under application of $\mathsf{K}_i^{A_1}$. Moreover, since $\mathsf{K}_i^{A_1}(R_j) = \mathsf{K}_i^{A_2}(R_j)$, we have $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$. Finally, note that by parts (2) and (5) of Lemma B.7, $\mathsf{C}^{A_1}(E) = \cap_{j=1}^k (D_j \cup \mathsf{C}^{A_1}(R_j))$. By Lemma B.6 and part (5) of Lemma B.7 (taking $E' = \emptyset$), it follows that $\mathsf{C}^{A_1}(R_j)$ is a fixedpoint set. So $\mathsf{C}^{A_1}(E)$ is a fixedpoint set, and hence a safe set. Thus, safe sets are closed under application of C^{A_1}. Moreover, since $\mathsf{C}^{A_1}(R_j) = \mathsf{C}^{A_2}(R_j)$ by Lemma B.6, we have $\mathsf{C}^{A_1}(E) = \mathsf{C}^{A_2}(E)$. ∎

Theorem 5.2 follows immediately from Proposition B.9, because every set in $ck_{A_1}(S)$ is clearly safe. ∎

Theorem 5.3: *If A_1 and A_2 are ω^2-coherent Aumann structures based on (S,T,τ) then $cck_{A_1}(S) = cck_{A_2}(S)$. Moreover, if $E \in cck_{A_1}(S)$, if i is an agent, and if G is a group of agents, then $\mathsf{K}_i^{A_1}(E) = \mathsf{K}_i^{A_2}(E)$ and $\mathsf{C}_G^{A_1}(E) = \mathsf{C}_G^{A_2}(E)$.*

Proof: The proof is very similar to that of Theorem 5.2, but much simpler, so we just sketch the details here. Given an association of states in T with ω^2-worlds, we define what it means for a set $R \subseteq T$ to be a λ-*cylinder set* in the obvious way, for an arbitrary $\lambda < \omega^2$, and take an *extended cylinder set* to be a λ-cylinder set for some $\lambda < \omega^2$. The same arguments as in Lemma B.3 show that extended cylinder sets are closed under finite union and complementation. Indeed, extended cylinder sets are also closed under a limited form of countable union and intersection: if R_j is a λ_j-cylinder set and there exists $\lambda' < \omega^2$ such that $\lambda_j < \lambda'$ for $j = 1,2,3,\ldots$, then $\cap_j R_j$ and $\cup_j R_j$ are also extended cylinder sets. The argument in Lemma B.4 shows that extended cylinder sets are closed under the application of K_i and that $\mathsf{K}_i^{A_1}(R) = \mathsf{K}_i^{A_2}(R)$ for an extended cylinder set R. Moreover, these arguments can be extended to show that extended cylinder sets are closed under the application of C_G and that $\mathsf{C}_G^{A_1}(R) = \mathsf{C}_G^{A_2}(R)$ for an extended cylinder set R. For if R is a λ-cylinder set for $\lambda < \omega^2$, then $(\mathsf{O}_G^{A_1})^k(R)$ is a $(\lambda + k)$-cylinder set. Since $\lambda + k < \lambda + \omega < \omega^2$, it follows by our earlier observation that $\mathsf{C}_G^{A_1}(R) = \cap_k (\mathsf{O}_C^{A_1})^k(R)$ is an extended cylinder set. Thus, every set in $cck(A_1)$ is an extended cylinder set, and the result follows. ∎

The proof of Theorem 5.3 uses the fact that whenever $\lambda < \omega^2$, then $\lambda + \omega < \omega^2$. In fact, ω^2 is the least ordinal β such that whenever $\lambda < \beta$, then $\lambda + \omega < \beta$. It is because of our use of this property that we cannot replace ω^2 by any smaller ordinal β in Theorem 5.3.

Acknowledgments: We are grateful to Robert Simon for his help with Theorem 3.6; in particular, Lemma A.3 is due to Simon. We are grateful also to Aviad Heifetz, Eddie Dekel, Jean-Francois Mertens, Dov Samet, and Shmuel Zamir for helpful discussions and comments.

References

[Aum76] R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

[Aum89] R. J. Aumann. Notes on interactive epistemology. Cowles Foundation for Research in Economics working paper, 1989.

[BD93] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189–198, 1993.

[BE79] W. Böge and T. Eisele. On solutions of Bayesian games. *International Journal of Game Theory*, 8:193–215, 1979.

[BS74] J. L. Bell and A. B. Slomson. *Models and Ultraproducts: An Introduction (third revised printing)*. North-Holland/American Elsevier, Amsterdam/New York, 1974.

[EGS80] P. van Emde Boas, J. Groenendijk, and M. Stokhof. The Conway paradox: its solution in an epistemic framework. In *Proc. 3rd Amsterdam Montague Symposium*, pages 87–111. Math. Centrum, Amsterdam, 1980. Reprinted by Foris Publications, 1984.

[Fag94] R. Fagin. A quantitative analysis of modal logic. *Journal of Symbolic Logic*, 59(1):209–252, 1994.

[FHMV95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, Mass., 1995.

[FHV91] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A model-theoretic analysis of knowledge. *Journal of the ACM*, 91(2):382–428, 1991. A preliminary version appeared in *Proc. 25th IEEE Symposium on Foundations of Computer Science*, 1984.

[Gil88] I. Gilboa. Information and meta information. In M. Y. Vardi, editor, *Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 227–243. Morgan Kaufmann, San Francisco, Calif., 1988.

[Har68] J. Harsanyi. Games with incomplete information played by 'Bayesian' players, parts I-III. *Management Science*, 14:159–182, 320–334, 486–502, 1968.

[HC68] G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. Methuen, London, 1968.

[HM90] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.

[HM92] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[HS93] A. Heifetz and D. Samet. Universal partition structures. Working Paper 26/93, Israel Institute of Business Research, Tel Aviv University, June 1993.

[HS98] A. Heifetz and D. Samet. Knowledge spaces with arbitrarily high rank. To appear, *Games and Economic Behavior*, 1998.

[Kri59] S. Kripke. A completeness theorem in modal logic. *Journal of Symbolic Logic*, 24:1–14, 1959.

[Lew73] D. K. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Mass., 1973.

[Lip91] B. L. Lipman. How to decide how to decide how to .... *Econometrica*, 59(4):1105–1126, 1991.

[MZ85] J. F. Mertens and S. Zamir. Formulation of Bayesian analysis for games of incomplete information. *International Journal of Game Theory*, 14(1):1–29, 1985.

[Rub89] A. Rubinstein. The electronic mail game: strategic behavior under "almost common knowledge". *American Economic Review*, 79:385–391, 1989.

[Sho67] J. R. Shoenfield. *Mathematical Logic*. Addison-Wesley, Reading, Mass., 1967.

[TW88] T. C.-C. Tan and S. R. da C. Werlang. On Aumann's notion of common knowledge: an alternative approach. CARESS Working Paper 88-09, University of Pennsylvania, 1988.

[Wer89] S. R. da C. Werlang. Common knowledge. In J. Eatwell, M. Milgate, and P. Newton, editors, *Game Theory*, pages 323–340. W. W. Norton, New York and London, 1989.