Int J Game Theory (1999) 28:367–384

International Journal of Game Theory

The difference between common knowledge of formulas and sets*

Robert Samuel Simon

Institute for Mathematical Stochastics, Lotzestr. 13, D-37083 Göttingen, Germany
(e-mail: simon@math.uni-goettingen.de)

Abstract. Common knowledge can be defined in at least two ways: syntactically as the common knowledge of a set of formulas or semantically, as the meet of the knowledge partitions of the agents. In the multi-agent S5 logic with either finitely or countably many agents and primitive propositions, the semantic definition is the finer one. For every subset of formulas that can be held in common knowledge, there is either only one member or uncountably many members of the meet partition with this subset of formulas held in common knowledge. If there are at least two agents, there are uncountably many members of the meet partition where only the tautologies of the multi-agent S5 logic are held in common knowledge. Whether or not a member of the meet partition is the only one corresponding to a set of formulas held in common knowledge has radical implications for its topological and combinatorial structure.

Key words: Common knowledge, Baire category, Cantor sets

1. Introduction

The knowledge of agents has been understood primarily in two ways, syntactically and semantically. Syntactic knowledge is expressed by sentences of a

* This research was supported by the German Science Foundation (Deutsche Forschungsgemeinschaft), the Institute of Mathematical Economics (Bielefeld), the Center for Rationality and Interactive Decision Theory (Jerusalem), the Department of Mathematics of the Hebrew University (Jerusalem), and the Edmund Landau Center for Research in Mathematical Analysis (Jerusalem), sponsored by the Minerva Foundation (Germany).

language or a logic (for example by the sentence $k_j f$, meaning that agent j knows that the sentence f is true). Semantic knowledge is expressed by subsets of a state space. In game theory and economics, the semantic approach has been the dominant mode to understanding knowledge and common knowledge; for example, witness the semantic definition of common knowledge given by Aumann (1976). Through semantic models for the syntax of a logic, the two approaches to knowledge can be closely related.

We look at a logic with a corresponding canonical semantic model where knowledge is defined by partitions and for which there is an equivalence between syntactic and semantic knowledge. We compare syntactic and semantic common knowledge in direct parallel to the equivalence of syntactic and semantic knowledge. With regard to common knowledge, we show that, unlike with knowledge, there is a large discrepancy between the syntactics and the semantics.

We consider the space of maximally consistent sets of sentences using the multi-agent epistemic logic S5, implicitly equivalent to what are known in the literature as the space of "knowledge structures"; (see Fagin, Halpern, and Vardi 1991). Because of their simplicity, we will call the sentences used in this paper *formulas*. For every agent, we define a partition of the space that we call its *knowledge partition*. Two points belong to the same member of the knowledge partition of an agent if and only if the agent knows the same set of formulas at both points. The knowledge partition is a semantic concept defined in syntactic terms (yet we will see later from a completeness theorem that this definitional equivalence of semantic and syntactic knowledge carries deeper significance).

A formula is common knowledge if all agents know the formula, all agents know that all agents know the formula, ad infinitum. We define two partitions of the space that correspond to syntactic and semantic common knowledge. For the former partition, two points belong to the same partition member if and only if they share the same formulas in common knowledge. For the latter partition, we take the meet of the knowledge partitions for all the agents, the same definition for common knowledge introduced by Aumann (1976).

A member of the meet partition defining semantic common knowledge we call a *cell*.

Throughout this paper we assume that there are denumerably many formulas, which includes the case of *finitely* many agents and primitive propositions.

We prove the following two theorems. **Theorem 1:** The partition defining semantic common knowledge is finer than that defining syntactic common knowledge, with either one cell or uncountably many cells corresponding to any set of formulas held in common knowledge. **Theorem 2:** With at least two agents, there exist uncountably many cells where only the tautologies of the multi-agent S5 logic are known in common.

We prove Theorems 1 and 2 by introducing a topology on the semantic state space and applying the Baire Category Theorem. With this topology, a sequence of points converges to a point if every formula in the syntax true at this limit point is true at all points in the sequence that come after some finite stage. This topology is equivalent to that used previously; see Samet (1990).

We call a cell *centered* if and only if there is no other cell with exactly the same set of formulas held in common knowledge. We define a *possibility set* to be a member of any one agent's knowledge partition.

From any semantic model, there is a canonical map to the state space of maximally consistent sets of formulas. If x in a semantic model maps to y in this state space, it has been shown that the agents' knowledge at y may not represent completely the knowledge of the agents at x; (see Fagin, Halpern, and Vardi 1991, and especially Fagin 1994). We mean by this that there is a semantic model with two points x and x^* mapping both to the same point y, yet there is an agent whose knowledge in the semantic model at x and x^* differs in an essential way; the precise meaning of essentially equivalent knowledge in semantic models concerns distinctions of knowledge defined using transfinite ordinals; (see also Heifetz and Samet 1997). Given there are finitely many agents and primitive propositions, we examine a property of an element in the state space that is equivalent to its representing completely the knowledge of the agents at its inverse image in any semantic model, namely the property that every possibility set in the cell containing this element is finite; (see Fagin, et al, 1997). When a cell has this property, we say that it has *finite fanout*. We consider another property of a cell, namely that every semantic model that maps to this cell does so surjectively. Such a cell we call a *surjective* cell. It is easy to show that all cells with finite fanout are also surjective cells. We prove **Theorem 3:** given denumerably many formulas, if a cell is surjective and is either centered or uncountable, then it has finite fanout. There is a example of a countable, un-centered and surjective cell that does not have finite fanout, (Simon 1998).

In addition to the main results, this paper has four additional contributions.

First, the proofs of the three theorems reveal a radical difference in the topological and combinatorial structure of cells depending on whether they are centered or un-centered. As part of **Proposition 2**, we prove that a centered cell is an open set relative to its closure. (On the other hand, a uncentered cell is nowhere dense in its closure.) We say that there is *n-mutual knowledge* of a formula f if for every string of agents $j_1, j_2, \ldots, j_n$, of length n it holds that agent j_1 knows that agent j_2 knows ... that agent j_n knows that f is true. Also as part of Proposition 2, we prove that for every compact cell C, there is an n such that if there is n-mutual knowledge at any $z \in C$ that a formula f is true it follows that there is common knowledge in C that f is true. For a syntactic characterization of a compact cell, see Feinberg (1997).

Second, Theorem 1 of this paper introduces two difficult open problems (also stated in the 6th section). Without the denumerability restriction on the number of formulas, we have failed to exclude any cardinality, for example the number 2, for the number of cells that can hold the same set of formulas in common knowledge. The denumerability of the formulas is essential to most of the important arguments in this paper. Also we have found no way to prove Theorem 1 without the use of Baire Category, and therefore also no way to strengthen Theorem 1 from a conclusion of uncountably many un-centered cells to a continuum cardinality of un-centered cells (given that we do not assume the Continuum Hypothesis).

Third, the main results of this paper bring us to what seems to be a paradox. Insofar as our agents have knowledge of their possibility sets, they have knowledge of a cell. From an intuitive notion of the meaning of knowledge and common knowledge, it seems that prior to the acquisition of their knowledge we should be able to identify this cell to our agents without revealing to them the rest of their knowledge; from Theorems 1 and 2 we cannot do so without at least telling them considerably more than the formulas in common knowledge. There is a way to identify a cell without revealing more information; but some may consider it to be a kind of cheating. Assuming the axiom of choice, we can well order any set and then identify a cell from a selection of a representative member point for each cell. However, we would prefer a more impartial method of identification, one based upon a well ordering of the *formulas* and some canonical rule. For example, if a game of incomplete information has a structure of states similar to that of a cell, we would prefer to describe the game and its rules without stating explicitly a possible situation (or all possible situations) which can arise in that game. The set of formulas held in common knowledge can be identified canonically from a well ordering of the formulas, simply by listing them. We would like to have a similarly straightforward method to identify any cell, something which may be impossible.

Fourth, our discovery of a discrepancy between syntactic and semantic common knowledge is largely independent of the discrepancies between syntactic and semantic knowledge that have to do directly with cardinality. With denumerably many formulas of the logic, an agent may have uncountably many different possibility sets and some possibility sets with uncountably many members. First, uncountably many possibility sets implies that the knowledge of an agent cannot always be implied logically by a single formula (or equivalently by finitely many formulas). Second, if an agent's possibility set is uncountable, the agent could select alternatively as possible only a proper subset of this uncountable set without changing the subset of formulas it considers to be possible; (see for example Heifetz 1999). This second possible discrepancy is responsible directly for the need to consider hierarchically constructed semantic models employing transfinite ordinals to represent the knowledge of the agents. (See Fagin 1994, Heifetz and Samet 1997, and especially Fagin, et al, 1999). However, Theorem 1 is a result that has cardinality consequences but is not *due* to cardinality, since both an element in the state space and syntactic common knowledge are defined by a subset of formulas. There are, however, relations between our topic and the *number* and *size* of possibility sets. Part of this relation finds expression in Theorem 3, mentioned above, and we explore it further in Section 6.1. Some degree of independence between our topic and the size of possibility sets is demonstrated by **Proposition 1:** With denumerably many formulas and at least two agents, there exists at least one cell with finite fanout such that only the tautologies are held in common knowledge. Because there are other cells containing infinite possibility sets such that only the tautologies are held in common knowledge, Proposition 1 implies that the set of formulas held in common knowledge does not determine whether a cell has finite fanout; (see above the relationship between finite fanout and the representation of knowledge in semantic models.) Therefore there is at least one important property of common knowledge that is determined by membership in a cell yet not determined by the set of formulas held in common knowledge.

The rest of this paper is organized as follows. In the second section, we introduce the background definitions and results that allow us to state precisely and prove the main result, Theorem 1. In the third section, we prove Theorem 1. The fourth section is devoted to the proof of Proposition 1 and Theorem 2. The fifth section explores the radical difference in structure between centered and un-centered cells. In this section we prove Proposition 2 and Theorem 3. The sixth section examines possibilities for further research, and also contains a short survey of how the centered property relates to other well known properties of cells, including properties concerning cardinality.

2. Background

Construct the set $\mathscr{L}(X,J)$ of formulas using the denumerable sets X and J in the following way:

1) If $x \in X$ then $x \in \mathscr{L}(X,J)$,
2) If $g \in \mathscr{L}(X,J)$ then $(\neg g) \in \mathscr{L}(X,J)$,
3) If $g,h \in \mathscr{L}(X,J)$ then $(g \wedge h) \in \mathscr{L}(X,J)$,
4) If $g \in \mathscr{L}(X,J)$ then $k_j g \in \mathscr{L}(X,J)$ for every $j \in J$,
5) Only formulas constructed through application of the above four rules are members of $\mathscr{L}(X,J)$.

We write simply $\mathscr{L}$ if there is no ambiguity. We define $g \vee h$ to be $\neg(\neg g \wedge \neg h)$ and $g \to h$ to be $\neg g \vee h$. For every finite subset $K \subseteq J$ $E_K(f) = E_K^1(f)$ is defined to be $\bigwedge_{j \in K} k_j f$, $E_K^0(f) := f$, and for $i \geq 1$, $E_K^i(f) := E_K(E_K^{i-1}(f))$.

A formula $f \in \mathscr{L}(X,J)$ is common knowledge in a subset of formulas $A \subseteq \mathscr{L}(X,J)$ if $E_K^n f \in A$ for every $n < \infty$ and every finite subset $K \subseteq J$.

Throughout this paper, the multi-agent epistemic logic $S5$ will be assumed. For a discussion of the $S5$ logic, see Cresswell and Hughes (1968); and for the multi-agent variation, see Halpern and Moses (1992) and also Bacharach, et al, (1997). Briefly, the $S5$ logic is defined by two rules of inference, modus ponens and necessitation, and five types of axioms. Modus ponens means that if f is a theorem and $f \to g$ is a theorem, then g is also a theorem. Necessitation means that if f is a theorem then $k_i f$ is also a theorem for all $j \in J$. The axioms are the following, for every $f, g \in \mathscr{L}(X,J)$ and $j \in J$:

1) all formulas resulting from theorems of the propositional calculus through substitution,
2) $(k_j f \wedge k_j(f \to g)) \to k_j g$,
3) $k_j f \to f$,
4) $k_j f \to k_j(k_j f)$,
5) $\neg k_j f \to k_j(\neg k_j f)$.

A set of formulas $\mathscr{A} \subseteq \mathscr{L}(X,J)$ is called *complete* if for every formula $f \in \mathscr{L}(X,J)$ either $f \in \mathscr{A}$ or $\neg f \in \mathscr{A}$. A set of formulas is called *consistent* if no finite subset of this set leads to a logical contradiction, meaning a deduction of f and $\neg f$ for some formula f. We define

$$\Omega(X,J) := \{S \subseteq \mathscr{L}(X,J) \mid S \text{ is complete and consistent}\}.$$

Any consistent set of formulas can be extended to a complete and consistent set of formulas, a property we call the *Extension Property*, proven by applying Lindenbaum's Lemma. A *tautology* of $\Omega(X,J)$ is a formula f in $\mathscr{L}(X,J)$ such that f is contained in every member of $\Omega(X,J)$. A formula is *possible* if its negation is not a tautology.

For every agent $j \in J$ we define its knowledge partition $\mathscr{Q}^j(X,J)$ to be the partition of $\Omega(X,J)$ generated by the inverse images of the function $\beta^j : \Omega \to 2^{\mathscr{L}(X,J)}$, the set of subsets of $\mathscr{L}(X,J)$, defined by

$$\beta^j(z) := \{f \in \mathscr{L}(X,J) \mid k_j f \in z\}.$$

Due to the fifth set of axioms $\beta^j(z) \subseteq \beta^j(z')$ implies that $\beta^j(z) = \beta^j(z')$. We will write $\mathscr{Q}^j$ if there is no ambiguity. A possibility set is defined to be a member of $\mathscr{Q}^j$ for some $j \in J$ and a cell is defined to be a member of the meet partition $\mathscr{Q} := \bigwedge_{j \in J} \mathscr{Q}^j$.

All the components of the proof of the following lemma can be found in other papers (Lemma 4.1 of Halpern and Moses 1992, Aumann 1989):

Lemma 1. *For any cell C of $\Omega(X,J)$ $\{f \in \mathscr{L}(X,J) \mid f$ is common knowledge in z for some $z \in C\} = \{f \in \mathscr{L}(X,J) \mid f$ is common knowledge in z for all $z \in C\} = \{f \in \mathscr{L}(X,J) \mid f \in z$ for all $z \in C\}$.*

Proof: The difficult direction is that starting with the assumption that $f \in z$ for all $z \in C$. By induction, it suffices to show for a given $z \in C$ that $k_j f \in z$ for all $j \in J$. By the Extension Property and the property that $\beta^j(z) \subseteq \beta^j(z')$ implies $\beta^j(z) = \beta^j(z')$ we need only that $f \notin \beta^j(z)$ implies that $\{\neg f\} \cup \{k_j g \mid g \in B^j(z)\}$ is consistent. We leave the rest as an exercise. The key step is to use that $(k_j f) \wedge (k_j g) \to k_j(f \wedge g)$ is a theorem. $\square$

For any subset $L \subseteq J$ the *L-adjacency distance* between any two points z, z' of Ω is defined to be $\rho_L(z,z') := \min\{d \mid$ there is a sequence $z = z_0, \dots, z_d = z'$, a function $a : \{1, \dots, d\} \to L$ and sequence of sets $D_i \in \mathscr{Q}^{a(i)}$ such that z_i and z_{i-1} both belong to $D_i\}$, with $\rho_L(z,z) = 0$ and $\rho_L(z,z') = \infty$ if there is no such sequence from z to z'. The *adjacency distance* is defined to be the J-adjacency distance.

We define a topology for Ω, the same as in Samet (1990). For every $f \in \mathscr{L}$ define $\alpha(f) := \{z \in \Omega \mid f \in z\}$. Let $\{\alpha(f) \mid f \in \mathscr{L}\}$ be the base of open sets of Ω. (A topology is defined by the fact that $\alpha(f) \cap \alpha(g) = \alpha(f \wedge g)$.) The topology of a subset A of Ω will be the relative topology for which the open sets of A are $\{A \cap O \mid O$ is an open set of $\Omega\}$. For any subset $D \subseteq \Omega$, $\overline{D}$ will stand for the closure of D.

Due to Lemma 1, we have a map F from the meet partition $\mathscr{Q}$ to subsets of formulas defined by $F(C) := \{f \mid f$ is common knowledge in any (equivalently all) members of C$\}$. Now we investigate this map.

For any subset of formulas $T \subseteq \mathscr{L}$ define $\underline{Ck}(T) := \{f \in \mathscr{L} \mid$ there exists a finite $L \subseteq J$, an $i < \infty$ and a finite set $T' \subseteq T$ with $(\wedge_{t \in T'} E^i_L(t)) \to f$ a tautology$\}$. We define $\mathscr{T}(X,J) = \{\underline{Ck}(T) \mid T \subseteq \mathscr{L}(X,J)\} \backslash \{\mathscr{L}(X,J)\}$, and we say that T generates $\underline{Ck}(T)$. If there is no ambiguity, we can write simply $\mathscr{T}$. $\underline{Ck}(T)$ is the set of formulas whose common knowledge is implied by the common knowledge of the formulas in T.

Lemma 2. *$\underline{Ck}$ is a closure operator on the subsets of $\mathscr{L}$, meaning that $\underline{Ck}(T) = T$ if T is already a member of $\mathscr{T}$. Furthermore, the image of $F : \mathscr{Q} \to 2^{\mathscr{L}}$ is a subset of $\mathscr{T}$.*

Proof: That it is a closure operator follows from the finiteness of the i, T' and L. For the rest, we need only to show for any S which is the set of formulas in common knowledge at any point of Ω that $\underline{Ck}(S) = S$. This follows directly from the axioms and rules of inference of the S5 logic. □

A cell is *centered* if the map $\mathrm{F} : \mathscr{Q} \to \mathscr{T}$ is one-to-one at this cell. The map $F : \mathscr{Q} \to \mathscr{T}$ is not surjective. Consider the set of formulas $T = \{k_j x \vee k_j(\neg x) \mid j \in J, x \in X\}$. $\underline{Ck}(T)$ does not include any $x \in X$, however for every member of $\Omega(X, J)$ with T in common knowledge and every $x \in X$ either x or $\neg x$ will be in common knowledge. If $T \in \mathscr{T}$ is in the image of F, we say that T has *actual* common knowledge. Since we defined $\mathscr{T}$ so that $\mathscr{L}$ is not a member, by the Extension Property all maximal members of $\mathscr{T}$ have actual common knowledge.

For every set of formulas $T \subseteq \mathscr{L}$ define the set

$$\mathbf{Ck}(T) := \{z \in \Omega \mid \text{every member of } T \text{ is common knowledge in } z\}.$$

For any $T \subseteq \mathscr{L}$, $\mathbf{Ck}(T)$ is a closed set, since the $\mathbf{Ck}(T)$ is the intersection of the sets $\alpha(E^l_K f)$ for all $l < \infty$, finite $K \subseteq J$, and all formulas f in T.

Lemma 3. *If C is a cell and $S = F(C)$, then $\bar{C} = \mathbf{Ck}(S)$.*

Proof: Suppose that the cell C has an empty intersection with $\alpha(f)$ for some formula $f \in \mathscr{L}$ such that $\alpha(f) \cap \mathbf{Ck}(S)$ is not empty. Then we can conclude by Lemma 1 that $\neg f$ is common knowledge in the cell C, $\neg f \in S$ and $\neg f$ is true everywhere in $\mathbf{Ck}(S)$, a contradiction. □

By Lemma 3 one would get the equivalent correspondence to F by defining $\mathscr{T}^* := \{\mathbf{Ck}(T) \mid T \subseteq \mathscr{L}, \mathbf{Ck}(T) \neq \emptyset\}$ and for a cell C defining $F^*(C)$ to be $\bar{C}$. $\mathscr{T}^*$ and $\mathscr{T}$ are in one to one correspondence, but in reverse containment order, with $\mathbf{Ck} : \mathscr{T} \to \mathscr{T}^*$ defining the bisection. By Lemma 1 and the above, $\mathscr{T}^*$ is the set of non-empty closed subsets of Ω that are also unions of cells.

3. The central result

For every $d < \infty$, finite subset $L \subseteq J$ and $z \in \Omega$ define $R^d_L(z) := \{z' \mid \rho_L(z, z') \le d\}$. Notice that the cell containing z is the union of the $R^d_L(z)$ for all finite $L \subset J$ and $d < \infty$.

The following lemma is a direct corollary of Lemma 2 of Samet (1990).

Lemma 4. *For every $z \in \Omega$, finite subset $L \subseteq J$, and $0 \le d < \infty$, $R^d_L(z)$ is a closed set.*

Recall that a set is defined to be *meagre* if its closure contains no nonempty open set.

Theorem 1. *The map $F : \mathscr{Q} \to \mathscr{T}$ is everywhere either one-to-one or uncountable-to-one.*

Proof: Let us assume that $S \in \mathscr{T}, F^{-1}(S)$ is not empty, and that for some $z \in C \in F^{-1}(S)$, some finite $L \subseteq J$ and some $d > 0$ the set $R^d_L(z)$ is not meagre in $\mathbf{Ck}(S)$ and therefore contains $\alpha(f) \cap \mathbf{Ck}(S) \neq \emptyset$ for some formula f. Let us assume that z' is a point in $\mathbf{Ck}(S)\backslash C$. It follows that $\alpha(f) \cap R^e_M(z') = \emptyset$ for every $e < \infty$ and every finite $M \subseteq J$, since otherwise z' would be in $R^{d+e}_{L \cup M}(z) \subseteq C$. By Lemma 1 we conclude that $\neg f$ is common knowledge at z' and therefore z' cannot belong to any member of $F^{-1}(S)$.

Now we assume that $F^{-1}(S)$ is not empty and for every $z \in C \in F^{-1}(S)$, every d and every finite $L \subseteq J$ that the set $R^d_L(z)$ is meagre in $\mathbf{Ck}(S)$. Since every member of $F^{-1}(S)$ is dense in $\mathbf{Ck}(S)$, if $g \notin S$ then the closed set $\mathbf{Ck}(S \cup \{g\})$ is either empty or meagre in $\mathbf{Ck}(S)$. Since $\mathscr{L}$ is countable, $\bigcup_{g \in \mathscr{L}\backslash S} \mathbf{Ck}(S \cup \{g\}) = \mathbf{Ck}(S)\backslash(\bigcup_{C \in F^{-1}(S)} C)$ is a countable union of meagre closed sets of $\mathbf{Ck}(S)$. By Lemma 4 and our assumption, every $C \in F^{-1}(S)$ is a countable union of meagre closed sets of $\mathbf{Ck}(S)$. Therefore the Baire Category Theorem implies that $F^{-1}(S)$ is uncountable. q.e.d.

4. Existence

Before we prove Theorem 2 we need some preliminaries concerning the semantic models for our logic, called Kripke structures.

In this paper, a Kripke structure is a quintuple $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$ where J is a set of agents, for each $j \in J$ $\mathscr{P}^j$ is a partition of the set S, X is a set of primitive propositions, and $\psi : X \to 2^S$ is a map from X to the subsets of S, such that for every $x \in X$ the set $\psi(x)$ is interpreted to be the subset of S where x is true. (The usual definition of a Kripke structure is more general, but this more restricted usage applies to the $S5$ logic.) We define a map $\alpha^\mu : \mathscr{L}(X, J) \to 2^S$ inductively on the structure of the formulas in the following way:

Case 1: $f = x \in X : \alpha^\mu(x) := \psi(x)$.

Case 2: $f = \neg g : \alpha^\mu(f) := S\backslash\alpha^\mu(g)$,

Case 3: $f = g \wedge h : \alpha^\mu(f) := \alpha^\mu(g) \cap \alpha^\mu(h)$,

Case 4: $f = k_j(g) : \alpha^\mu(f) := \{s \mid s \in P \in \mathscr{P}^j \Rightarrow P \subseteq \alpha^\mu(g)\}$.

We define a map $\phi^\mu : S \to \Omega(X, J)$ (see Fagin, Halpern, and Vardi 1991) by

$$\phi^\mu(s) := \{f \in \mathscr{L}(X, J)) \mid s \in \alpha^\mu(f)\}.$$

We are justified in using again the notation α for the following reason. Consider the map $\bar{\psi} : X \to 2^\Omega$ defined by $\bar{\psi}(x) := \{z \in \Omega \mid x \in z\}$. We have a Kripke structure $\Omega = (\Omega; J; \mathscr{Q}^1, \ldots, \mathscr{Q}^n; X; \bar{\psi})$. (Due to its canonical nature, we index this Kripke structure with Ω.)

Theorem. *For every $f \in \mathscr{L}(X, J)$, f is a theorem of the multi-agent S5 logic if and only if f is a tautology. Furthermore, $\phi^\Omega(z) = z$ for every $z \in \Omega$.*

For a proof of the first part of this theorem, see Halpern and Moses (1992) and Cresswell and Hughes (1968), and for how the second part follows from the first part see Aumann (1999). We will call this result the "Completeness Theorem."

We define the *depth* of a formula inductively on the structure of the formulas. If $x \in X$, then depth $(x) := 0$. If $f = \neg g$ then depth $(f) :=$ depth (g); if $f = g \wedge h$ then depth $(f) :=$ max (depth (g), depth (h)); and if $f = k_j(g)$ then depth $(f) :=$ depth $(g) + 1$.

For a Kripke structure $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$, if $s \in \alpha^\mu(f)$, or equivalently $f \in \phi^\mu(s)$, we say that f is true at s with respect to μ. We say that f is valid with respect to the Kripke structure μ if f is true at s with respect to μ for every $s \in S$. The Kripke structure is *connected* if the meet partition $\wedge_{j \in J} \mathscr{P}^j$ is a singleton (equal to $\{S\}$). We define a *connected component* of the Kripke structure to be a member of this meet partition. Two points $s, s' \in S$ are *adjacent* if they share some member of $\mathscr{P}^j$ for some $j \in J$. We define the *adjacency distance* between any two points s and s' in S as $\min\{d \mid$ there is a sequence $s = s_0, \ldots, s_d = s'$, a function $a : \{1, \ldots, d\} \to J$ and sequence of sets $D_1 \in \mathscr{P}^{a(1)}, \ldots, D_d \in \mathscr{P}^{a(d)}$ such that for all $1 \leq i \leq d$ s_i and s_{i-1} both belong to $D_i\}$, with zero distance between any point and itself and infinite distance if there is no such sequence from s to s'. By Lemma 4.1 of Halpern and Moses (1992), the truth of any formula of depth d at a point $s \in S$ is determined by the map ψ and the agents' partition members that are of adjacency distance of no more than d from s.

Lemma 5. a) *Let $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$ be a Kripke structure. If P is a member of $\mathscr{P}^j$ for some $j \in J$ then $\phi^\mu(P)$ is a dense subset of F for some $F \in \mathscr{Q}^j$.*

b) *If additionally the Kripke structure μ is connected, then $\phi^\mu(S)$ is contained and dense in a cell of Ω.*

Proof:

a) By the definition of $\phi^\mu, \phi^\mu(P)$ must be contained in a single member F of $\mathscr{Q}^j$. Let us suppose for the sake of contradiction that there is a formula f with $\alpha(f) \cap F \neq \emptyset$ and $\alpha(f) \cap \phi^\mu(P) = \emptyset$. The former implies that $\neg k_j \neg f \in z$ for every $z \in F$ and the latter implies that $P \cap \alpha^\mu(f) = \emptyset$ and $k_j \neg f \in z$ for every $z \in \phi^\mu(P)$, a contradiction.

b) The containment in a cell follows directly from the definition of the map ϕ^μ, since two points adjacent in μ must also have adjacent images. Assume that the cell C contains $\phi^\mu(S)$ and that there exists a formula f with $\emptyset \neq C \cap \alpha(f) \subseteq C \backslash \phi^\mu(S)$. This means that $\neg f$ is valid with respect to μ, and hence by Lemma 4.1 of Halpern and Moses (1992) that $\neg f$ is common knowledge in $\phi^\mu(s)$ for every $s \in S$, a contradiction to Lemma 1, because it would imply that $\neg f$ is common knowledge in C. □

We say that a cell has *finite fanout* if every possibility set contained in the cell is a finite set.

Proposition 1. *With at least two agents there exists a cell of finite fanout that is dense in Ω.*

Proof: Since the set of formulas is denumerable, list all of the possible formulas $\{f_1, f_2, \ldots\}$ in $\mathscr{L}(X, J)$. We leave as an exercise that one can create finite connected Kripke structures $\mu_1, \mu_2, \ldots$ defined on sets $S_1, S_2, \ldots$ such that for every $i < \infty f_i$ is true with respect to μ_i at some point $s_i \in S_i$ and there is a point $t_i \in S_i$ of adjacency distance from s_i of at least depth $(f_i) + 2$. We assume that all the sets S_i are mutually disjoint, and we define the disjoint union $S := \bigcup_{i=1}^{\infty} S_i$. We choose any two agents j_i and j_2 and create a new connected Kripke structure μ with the set S. For all $j \in J$ let F_i^j be the member of $\mathscr{P}_i^j$ containing t_i, where $\mathscr{P}_i^j$ is the agent j's partition of S_i. We define

$$\mathscr{P}^{j_1} = \bigcup_{i=1}^{\infty} \mathscr{P}_i^{j_1} \backslash \{F_i^{j_1}\} \bigcup_{i \text{ is odd}} \{F_i^{j_1} \cup F_{i+1}^{j_1}\}$$

and

$$\mathscr{P}^{j_2} = \mathscr{P}_1^{j_2} \cup \bigcup_{i=2}^{\infty} \mathscr{P}_i^{j_2} \backslash \{F_i^{j_2}\} \bigcup_{i \text{ is even}} \{F_i^{j_2} \cup F_{i+1}^{j_2}\}.$$

For j not equal to j_1 or j_2, we define $\mathscr{P}^j = \bigcup_{i=1}^{\infty} \mathscr{P}_i^j$. We let $\psi : X \to 2^S$ be defined by $\psi(x) = \bigcup_{i=1}^{\infty} \psi_i(x)$. Due to the fact that within an adjacency distance of depth (f_i) from s_i nothing has changed from the Kripke structure μ_i, f_i remains true at s_i with respect to μ; therefore our connected Kripke structure μ is mapped to a cell that intersects $\alpha^{\mu}(f_i)$ for every i, and hence this cell is dense in Ω. We must show that μ maps to this cell surjectively. By Lemma 5a, every possibility set intersecting the image of ϕ^{μ} is finite and is mapped to surjectively. The rest follows by the definition of a cell as a member of the meet partition. □

For finite X and J, one can observe that the dense cells defined by the infinite repetition of the "no-information extension" of Fagin, Halpern, and Vardi (1991) have some possibility sets that are topologically equivalent to Cantor sets. Therefore, in the context of finite X and J, we have shown already that there exists are least two distinct dense cells, and therefore, by Theorem 1, uncountably many. However, we prefer to have a proof of Theorem 2 that does not make use of the "no-information extension" and is valid for countably many agents or primitive propositions.

Proof of Theorem 2: By Proposition 1 there exists at least one cell dense in Ω. By Theorem 1, suppose for the sake of contradiction that there is only one such cell C. By the proof of Theorem 1 C must contain an open set of Ω, hence also $\alpha(f)$ for some possible f. Since for every possible formula $f \in \mathscr{L}$ there is a finite Kripke structure such that f is true at some point of this Kripke structure, (Halpern and Moses 1992), by Lemma 5b there is a finite cell intersecting $\alpha(f)$. By assumption, this finite cell must be C; and by the density of C in Ω, Ω must be finite and equal to C! This can be contradicted in several ways, the easiest being a demonstration that there must be at least two cells in Ω, for example the one resulting from the common knowledge of x for all $x \in X$ and another resulting from the common knowledge of $\neg x$ for all $x \in X$. q.e.d.

5. Related results

5.1. The structure of cells

In this subsection, we expand on the topological and combinatorial characterizations of centeredness and un-centeredness.

Corollary 1. *A cell C is centered if and only if there exists some $d < \infty$, finite subset $L \subseteq J$, and $z \in C$ such that $R_L^d(z)$ is not meagre relative to C.*

Proof: By the proof of Theorem 1, it suffices to prove for $d < \infty$ and finite $L \subseteq J$ that $R_L^d(z)$ is not meagre in C if and only if it is not meagre in $\bar{C}$. This follows directly from the fact that C is dense in $\bar{C}$ and $R_L^d(z)$ is a closed set contained in C. □

Notice from the finite adjacency distance between any two points of a cell that if C is a cell and $R_L^l(z)$ contains an open set of $\bar{C}$ for some $z \in C$, finite $L \subseteq J$ and $l < \infty$ then for all $z' \in C$ there exists an $l' < \infty$ and a finite subset $M \subseteq J$ with $R_M^{l'}(z')$ containing this open set of $\bar{C}$.

If one wants to show that a cell is centered, Corollary 1 is easier to use than Theorem 1. If we want to apply Theorem 1, the following is useful.

Proposition 2. *If C is a centered cell of Ω then for every $z, z' \in C$ there exists an open set O, finite $L \subseteq J$ and an $l < \infty$ such that $z' \in O \cap \bar{C} \subseteq R_L^l(z)$. In particular C is open relative to $\bar{C}$.*

Furthermore, given that $S = F(C)$, consider the following properties:
(1) *C is compact*
(2) *C is centered and S is maximal in $\mathscr{T}$,*
(3) *the adjacency diameter of C is finite.*
(1) and (2) are equivalent and they both imply (3). If there are finitely many agents, then (3) implies both (1) and (2).

Proof: Let z be a fixed member of C, a centered cell. Let $L \subseteq J$ and $d < \infty$ be such that $R_L^d(z)$ contains $\alpha(f) \cap \bar{C} \neq \emptyset$ for some formula f in $\mathscr{L}(X, J)$. Letting z' be any other member of C, by the Completeness Theorem there is a finite subset $K \subseteq J$, an m and a sequence of agents $j_1, j_2, \dots, j_m$ in K such that $\neg k_{j_1} \neg (\cdots (\neg k_{j_m} \neg f) \cdots)$ is in z'. Considering any other $z^* \in \bar{C} \cap \alpha (\neg k_{j_1} \neg (\cdots (\neg k_{j_m} \neg f) \cdots)))$, we have $\emptyset \neq R_K^m(z^*) \cap \alpha(f) \subseteq \bar{C}$, and therefore $\rho(z, z^*) \leq m + d$.

Due to Lemma 4 and the definition of centeredness, the only difficult direction to prove is ((1) and (2)) $\Rightarrow$ (3), which follows directly from the above and compactness. q.e.d.

If a cell C is not centered, then from the proof of Theorem 1 and the Baire Category Theorem $\bar{C} \backslash C$ must be dense in $\bar{C}$. Furthermore, for some $S \in \mathscr{T}$ with actual common knowledge it is possible that $\bigcup_{C \in F^{-1}(S)} C$ is meagre in $\mathbf{Ck}(S)$; from the proof of Theorem 2 we see this is true for the cells dense in Ω.

If a cell C is centered but not compact, from Proposition 2 we know that $\bar{C} \backslash C$ is closed and a non-empty union of cells, and therefore, by Lemma 1, is equal to $\mathbf{Ck}(S)$ for some $S \in \mathscr{T}$, a fact that may prove to be very useful. This $S \in \mathscr{T}$ may, however, fail to be a member of actual common knowledge.

5.2. Unique extension and surjectivity

We say that a point in Ω has *unique extension* if it has a unique extension to all canonical semantic models associated with the transfinite ordinals, (see Fagin, 1994. Unique extendibility is what we referred to in the introduction as the ability of an element in Ω to represent completely the knowledge of the agents at any point in a semantic model that maps to this element.) For finitely many agents and primitive propositions, Fagin, et al, (1999) proved that a point of Ω has unique extension if and only if the cell containing this point has finite fanout.

Since the property of unique extension is a property of cells, hence that of some form of common knowledge, it is natural to formulate a question concerning common knowledge in parallel to what has been proven with regard to knowledge. Given that a cell C does not have finite fanout, we know from Fagin, et al, (1999) that there is a $z \in C$, a Kripke structure $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$, and a $P \in \mathscr{P}^j$ for some j such that $\phi^\mu(S) \subseteq C$ and $\phi^\mu(P) \subseteq F \backslash \{z\}$, where $z \in F \in \mathscr{Q}^j$. Does this imply also that there exists a Kripke structure $\hat{\mu} = (\hat{S}; J; (\hat{\mathscr{P}}^j \mid j \in J); X; \hat{\psi})$ with $\phi^{\hat{\mu}}(\hat{S}) \subseteq C \backslash \{z\}$? We call this **Question 1**, for which the general answer is *no*, (Simon 1998). However, if the cell is either uncountable or centered, the answer is *yes*, and this is proven by Theorems 3a and 3b, together composing Theorem 3, (stated in the introduction). First we must characterize those Kripke structures that map injectively into Ω.

We define a Kripke structure $\mu := (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$ to be *small* if the map $\phi^\mu : S \to \Omega(X, J)$ is injective. From now on, when we say that a Kripke structure μ is small we mean that S is a subset of Ω, that the $\mathscr{P}^j$ are partitions of S finer than or equal to $\mathscr{Q}^j|_S := \{F \cap S \mid F \in \mathscr{Q}^j,\ F \cap S \neq \emptyset\}$ and that $\phi^\mu(z) = z$ for every $z \in S$. By Lemma 5a we know that every member of such a $\mathscr{P}^j$ is dense in the member of $\mathscr{Q}^j$ containing it.

Lemma 6. *Any subset $S \subseteq \Omega(X, J)$ and partitions $(\mathscr{P}^j \mid j \in J)$ of S with the property that for every j and every $A \in \mathscr{P}^j$ A is a dense subset of some member of $\mathscr{Q}^j$ define a small Kripke structure. (The proof is similar to the main proof in Heifetz 1999)*

Proof: We define $\mu := (S; J; (\mathscr{P}^j \mid j \in J); X; \bar{\psi}|_S)$, where $\bar{\psi}|_S(x) = \bar{\psi}(x) \cap S$. It suffices to show that $\alpha^\mu(f) = \alpha^\Omega(f) \cap S$ for every $f \in \mathscr{L}$. We proceed by induction on the structure of formulas. From $\bar{\psi}(x) = \{z \mid x \in z\}$ the claim is true for the formulas $x \in \mathscr{L}$ with $x \in X$, and from the definition of α^μ if the claim is true for $f \in \mathscr{L}$ and $g \in \mathscr{L}$ then it true for $f \wedge g$ and $\neg f$. We assume the claim is true for $f \in \mathscr{L}$ and consider $k_j f$. We assume that $z \in F \in \mathscr{Q}^j$ and $z \in A \in \mathscr{P}^j$. If $z \in \alpha(k_j f) \cap S$ then F is contained in $\alpha(f)$. $A \subseteq F$ and the induction hypothesis implies that $A \subseteq \alpha^\mu(f)$ and $z \in \alpha^\mu(k_j f)$. On the other hand, if $z \in S \backslash \alpha(k_j f)$ then $F \cap \alpha^\Omega(\neg f)$ is not empty, and by the density of A in F also $A \cap \alpha^\Omega(\neg f)$ is not empty – therefore $A \cap \alpha^\mu(\neg f)$ is not empty by the induction hypothesis and $z \notin \alpha^\mu(k_j f)$. □

A. Heifetz (1999) has proven that, as long as the number of agents is at least two, there are 2^c distinct small Kripke structures for which the map into Ω is also surjective, where c is the cardinality of the continuum. With the characterization of Lemma 6, one can obtain his result by refining only the partition of one agent at only one possibility set F that is topologically equivalent to a Cantor set. This follows because there will be 2^c different ways to partition this set F into two disjoint dense subsets.

Given a Kripke structure $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$ and a subset $A \subseteq S$, we define the Kripke structure $\mathscr{V}^\mu(A) := (A; J; (\mathscr{P}^j|_A \mid j \in J); X; \psi|_A)$ where for all $j \in J$ $\mathscr{P}^j|_A := \{F \cap A \mid F \cap A \neq \emptyset$ and $F \in \mathscr{P}^j\}$ and for all $x \in X$ $\psi|_A(x) = \psi(x) \cap A$. If there is no ambiguity concerning the Kripke structure μ, we can replace $\mathscr{V}^\mu(A)$ by $\mathscr{V}(A)$.

We define a subset $A \subseteq \Omega$ to be *good* if for every $j \in J$ and every $F \in \mathscr{Q}^j$ satisfying $F \cap A \neq \emptyset$ it follows that $F \cap A$ is dense in F. By Lemma 5a and Lemma 6 A is good if and only if for every $z \in A$ $\phi^{\mathscr{V}(A)}(z) = z$. By Lemma 5b we know that a good subset of a cell must also be a dense subset of the cell.

The following lemma shows that to answer Question 1 we can equivalently ask for the existence of some good subset A of C such that $A \subseteq C \backslash \{z\}$.

Lemma 7. *If $\mu = (S; J; (\mathscr{P}^j \mid j \in J); X; \psi)$ is a Kripke structure then $A := \phi^\mu(S)$ is a good subset.*

Proof: If $F \in \mathscr{Q}^j$ intersects $\phi^\mu(S)$ let $s \in S$ be any point of S such that $\phi^\mu(s)$ is in F and let $s \in P \in \mathscr{P}^j$. $\phi^\mu(P)$ is dense in F by Lemma 5a, hence $\phi^\mu(S)$ is dense in F. □

Now we answer Question 1 for uncountable cells by combining Lemma 7 and Theorem 3a.

Theorem 3a. *For every cell C there is a countable good subset $A \subseteq C$.*

Proof: Letting $F \in \mathscr{Q}^j$ for some $j \in J$ and $F \subseteq C$, we choose any countable dense subset $F' \subseteq F$ and let $A_1 := F'$. We define inductively a sequence of countable sets $A_1, A_2, \ldots$ in the following way. For any fixed $k > 1$, $j \in J$, and $z \in A_{k-1}$, let F^* be the member of $\mathscr{Q}^j$ containing z and let $A_k^j(z)$ be any countable subset of F^* such that $A_k^j(z) \cup (F^* \cap A_{k-1})$ is a dense subset of F^*. We define A_k to be $A_{k-1} \bigcup_{j \in J, z \in A_{k-1}} A_k^j(z)$. The countable set $A = \bigcup_{k=1}^\infty A_k$ satisfies the necessary conditions. □

Now we answer Question 1 for countable centered cells.

Lemma 8. *Let C be a cell with no isolated points, let $F(C) = S$, and for every formula $f \in \mathscr{L}$ with $\alpha(f) \cap \mathbf{Ck}(S) \neq \emptyset$ assume that there is at least one countable cell $C' \subseteq \mathbf{Ck}(S)$ with $\alpha(f) \cap C' \neq \emptyset$, (necessarily the case if there is at least one countable cell in $F^{-1}(S)$). It follows that C is not centered.*

Proof: Let us assume, for the sake of contradiction, that C is centered, so that by Theorem 1 there is a formula f such that $C \supseteq \alpha(f) \cap \mathbf{Ck}(S) \neq \emptyset$. By the hypothesis of the lemma, we must assume that C has countably many members. Therefore C is a countable union of closed meagre sets of $\mathbf{Ck}(S)$. By the Baire Category Theorem, C could not have contained $\alpha(f) \cap \mathbf{Ck}(S)$. □

Theorem 3b. *The subset of isolated points of a centered countable cell is a good subset.*

Proof: Let A be the subset of isolated points of the cell C. By Lemma 8 we know that A is not empty. Assume that $F \in \mathcal{Q}^j$ contains some $a \in A$. Let $g \in \mathcal{L}$ be a formula such that $\bar{C} \cap \alpha(g) = \{a\}$. Let $e \in \mathcal{L}$ be any formula such that $\alpha(e) \cap F \neq \emptyset$. Since $\alpha(e) \cap F$ is a countable closed set, it has an isolated point. Let f be a formula such that z is the only point of $\alpha(f) \cap \alpha(e) \cap F$. Consider the formula $h := e \wedge f \wedge (\neg k_j \neg g)$: It follows that z is the only point of $\alpha(h) \cap \bar{C}$ and thus is also a member of A. q.e.d.

Lemma 9. *If A is a good subset of a cell C such that for every possibility set F $A \cap F \neq \emptyset$ implies that $A \cap F$ is closed, then $A = C$.*

Proof: If F is a possibility set and $F \cap A \neq \emptyset$ then $A \cap F = F$ by the hypothesis. The rest follows from the definition of a cell as a meet partition member. □

Corollary 2. *A centered cell of finite fanout contains only isolated points. In particular, an infinite centered cell of finite fanout cannot hold in common knowledge a set of formulas that is maximal in $\mathcal{T}$.*

Proof: It follows directly from Theorem 3b and Lemma 9. □

Proposition 1 shows that the centered condition cannot be dropped from the first half of Corollary 2. The same holds for the second half of Corollary 2; see the next section.

6. Related problems

6.1. Cardinalities

In this subsection, we survey the relations between various properties of cells. We restrict ourselves in this subsection to a finite number of agents and primitive propositions.

A possibility set $F \in \mathcal{Q}^j$ is *expressible* if there exists some formula f such that the formulas held in common knowledge in any point of F together with $k_j f$ are sufficient to imply $k_j g$ for all the formulas g that agent j knows in F; (see Samet 1990 and Maruta 1997). It is not hard to show (given the $S5$ logic) that if C is the cell containing F then F is expressible if and only if it contains an open set of C if and only if F is an open set of C, (Maruta 1997).

Consider the following properties of a cell C:

(a) C is centered,
(b) at least one possibility set of C is expressible,
(b′) all possibility sets of C are expressible,
(c) $F(C)$ is maximal in $\mathcal{T}$,
(d) C has finite fanout,
(e) $F(C) = \underline{Ck}(\{f\})$ for some single formula f
(f) C is finite.

If (e) holds, then (a), (b), (b$'$), (c), and (f) are all equivalent, (Simon 1997). (e) and (d) but none of the others can hold, as witnessed in Proposition 1. (e) can be interpreted to be the common knowledge parallel to expressibility for knowledge.

(b) implies (a), but the converse does not hold, even if one also assumes (c); a counter-example is in Simon (1997). By Lemma 8, such a counter-example must be an uncountable cell topologically equivalent to a Cantor set and by Proposition 2 it has finite adjacency diameter.

(b$'$) implies (d): A point in Ω is determined by the truth assignments for the primitive propositions and the possibility sets of the agents containing the point. As long as the agents and primitive propositions are finite, (b$'$) implies that all points of C are isolated. With the compactness of possibility sets, this implies (d).

(c) and (d) together fail to imply (a); a counter-example is in Simon (1998).

(b$'$) does not imply (c), as is demonstrated by any infinite centered cell of finite fanout and Corollary 2.

The counter-example demonstrating (a) and (c) but not (b) used three agents, (Simon 1997). The counter-example demonstrating (c) and (d) but not (a) used four agents; and this counter-example has the additional property that the cardinality of every possibility set in $\bar{C}$ is no more than 2, (Simon 1998). There is a counter-example demonstrating (c) but not (a) with three agents, such that all the possibility sets in $\bar{C}$ for one special agent are topologically equivalent to Cantor sets, (Simon 1997). We do not know if such counter-examples can be found demonstrating the same principles but with fewer agents.

6.2. Extending the central result

One can try to go further than Theorem 1, and ask the following question.

Question 2: If one does not assume the continuum hypothesis, can one prove for every $S \in \mathscr{T}$ that if $F^{-1}(S)$ is uncountable then it has the cardinality of the continuum?

Question 3: Does the conclusion of Theorem 1 hold for either arbitrary X or J?

Conjecture 1: The answers to Questions 2 and 3 are *no*.

If the number of agents, but not the number of primitive propositions, is countable, either *yes* or *no* as an answer to Question 3 seems plausible to us.

Question 2 and Question 3 are related: if one must use Baire Category to prove Theorem 1 then its generalization suggested by Question 3 is unlikely if there are uncountably many formulas. Without assuming the continuum hypothesis, a cover of a compact subset of real numbers by meagre closed sets need not have the cardinality of the continuum; see Bartoszynski and Judah (1995).

6.3. *Unique extension and surjectivity*

In the proof of Theorem 3b we showed that if z is a cluster point of a possibility set F of a countable centered cell C that there should be a Kripke structure that maps to $C \backslash \{z\}$. We did not show this in the proof of Theorem 3a.

Question 4: If C is an uncountable cell and z is a cluster point of a possibility set contained in C, what additional conditions are necessary, if any, for there to exist a Kripke structure that maps to $C \backslash \{z\}$?

From Lemma 8, it follows that there can be no disconnected Kripke structure that maps injectively to a countable centered cell.

Question 5: Does there exist a disconnected Kripke structure that maps injectively to a countable un-centered cell?

The following question seems now to be a curiosity, but may prove later to be important.

Question 6: Given a cell C and a subset $A \subseteq C$ is it possible for $\phi^{\mathscr{V}(A)}(A)$ to be contained in either A, C, or the closure of C and dense in the closure of C without $\phi^{\mathscr{V}(A)} = z$ for every $z \in A$?

6.4. *Other canonical models*

For every ordinal α, let U_α be the universal semantic model associated with α, as presented by Heifetz and Samet (1997), and we will assume that the number of agents and primitive propositions are both finite. For pairs of ordinals $\beta < \alpha$ define π^α_β to be the canonical surjective map from U_α to U_β. For every $j \in J$, let $\mathscr{P}^j_\alpha$ be the partition of U_α induced by the knowledge of agent j concerning the subsets of U_β for $\beta < \alpha$, as defined in Heifetz and Samet (1997). There are at least three ways to formulate common knowledge in U_α for limit ordinals α:

1) the meet partition $\mathscr{Q}_\alpha := \bigwedge_{j \in J} \mathscr{P}^j_\alpha$,

2) the equivalence relation $\sim$ on $\mathscr{Q}_\alpha$ defined by $C \sim C'$ if and only if for every $\beta < \alpha$ and every subset $A \subseteq U_\beta$ it follows that $C \subseteq (\pi^\alpha_\beta)^{-1}(A) \Leftrightarrow C' \subseteq (\pi^\alpha_\beta)^{-1}(A)$, the analog to common knowledge of formulas for $\alpha = \omega$,

3) the join partition $\bigvee_{\beta<\alpha} \{(\pi^\alpha_\beta)^{-1}(A) \mid A \in \mathscr{Q}_\beta\}$.

For $a = \omega$, the first infinite ordinal, and at least two agents we have shown that all three of these types of common knowledge are different.

For every ordinal number α greater or equal to ω we define a corresponding cardinal number $f(\alpha)$ in the following way:

$f(\omega) = c$, where c is the cardinality of the continuum,
$f(\alpha + 1) = 2^{f(\alpha)}$,
if α is a limit ordinal then $f(\alpha) := \prod_{\beta<\alpha} f(\beta)$.

Conjecture 2: For every $\alpha > \omega$, the cardinality of U_α is $f(\alpha)$ and for every limit ordinal α there is a coset of the equivalence relation of 2) such that there are $f(\alpha)$ different members of $\mathscr{Q}_\alpha$ in this coset.

Insofar as this paper pertains only to knowledge at the first infinite level, there may be an interesting relation between centeredness and subspaces of the Mertens-Zamir probability spaces (Mertens and Zamir 1985). The most natural way to define a probability measure on a Kripke structure is to determine for every formula the probability that this formula is true. In this way, we are led naturally to Borel measures on Ω, and we can take comfort in the fact that all cells are F_σ sets, (Samet 1990). However, if a cell C is not centered and $S = F(C)$ is the set of formulas held in common knowledge in C then some Borel probability measures on the set of points holding S in common knowledge and giving positive probability to every formula not in contradiction to S may have little or nothing to do with the probability measures on C. (An analogous comparison may be that between the Lesbesgue measure on $\mathbf{R}$ and measures defined on the set of rational numbers). On the other hand, if a cell C is centered then by Proposition 2 it is open in its closure, and hence any Borel probability measure on the set of points holding S in common knowledge and giving positive probability to every formula not in contradiction to S generates naturally a probability measure on the cell C. However, a full investigation of the Mertens-Zamir spaces along these lines would require an independent investigation of these spaces parallel to as not as an extension of the results presented here.

One can investigate other logics for similar distinctions in the content of common knowledge. For this paper, the multi-agent epistemic logic $S5$ was chosen because this logic involves a topology and a partition concept of knowledge that are relatively easy to work with. It would be interesting to know if similar results can be obtained in other logics, with or without using the same mathematical tools. Of special interest would be logics that allow for explicit statements of common knowledge.

Acknowledgements: The inquiry into the number of cells that are dense in Ω was inspired by communication with Robert Aumann and Dov Samet, and later with Aviad Heifetz. Proposition 1 is based on communication with Robert Aumann. Joseph Halpern informed me of related literature and helped sharpen some of the arguments. Mamoru Kaneko helped with some of the writing. Comments from Bart Lipman, Philippe Mongin and Dov Samet were very helpful.

References

Aumann R (1976) Agreeing to disagree. Annals of Statistics 4:1236–1239

Aumann R (1999) Interactive epistemology I: Knowledge. International Journal of Game Theory 28:263–300

Aumann R (1999) Interactive epistemology II: Probability. International Journal of Game Theory 28:301–314

Bacharach M, Gerard-Varet LA, Mongin P, Shin H (eds.) (1997) Epistemic logic and the theory of games and decisions. Kluwer, Dordrecht

Bartoszynski T, Judah, H (1995) Set theory: On the structure of the real line. AK Peters

Cresswell MJ, Hughes GE (1968) An introduction to modal logic. Routledge

Fagin R (1994) A quantitative analysis of modal logic. Journal of Symbolic Logic 59:209–252

Fagin R, Halpern YJ, Vardi MY (1991) A model-theoretic analysis of knowledge. Journal of the A.C.M 91:382–428

Fagin R, Geanakoplos J, Halpern YJ, Vardi MY (1999) The hierarchical approach to modeling knowledge and common knowledge. International Journal of Game Theory 28:331–365

Feinberg Y (1996) Characterizing common priors in the form of posteriors. manuscript

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

Heifetz A (1999) How canonical is the canonical model? A comment on Aumann's interactive epistemology. International Journal of Game Theory 28:435–442

Heifetz A, Samet D (1997) Hierarchies of knowledge: An unbounded stairway. manuscript, revision of "Universal Partition Spaces," IIBR working paper 26, Tel Aviv University

Heifetz A, Samet D (1998) Knowledge spaces with arbitrarily high rank. Games and Economic Behavior 22:260–273

Maruta T (1997) "Information structure on maximal consistent sets. manuscript

Mertens J-P, Zamir S (1985) Formulation of Bayesian analysis for games with incomplete information. International Journal of Game Theory 14:1–29

Samet D (1990) Ignoring ignorance and agreeing to disagree. Journal of Economic Theory 52:190–207

Simon R (1997) The generation of formulas in common knowledge. Discussion Paper 150, (June 1997) Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem

Simon R (1998) Locally finite knowledge structures. manuscript