On the Unique Extensibility and Surjectivity of Knowledge Structures

Robert Samuel Simon

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

Abstract: With the $S5$ multi-agent epistemic logic we consider the canonical maps from Kripke structures to knowledge structures. We define a cell to be a minimal subset of knowledge structures known in common semantically by the agents. A cell has finite fanout if at every point every agent considers only a finite number of other points to be possible. We define a cell to be surjective if every Kripke structure that maps to it does so surjectively. All cells with finite fanout are surjective, but the converse does not hold. To construct a counter-example we need topological insights concerning the relationship between the logic and its semantic models. The difference between syntactic and semantic common knowledge is central to this construction.

Key words: Kripke Structures, Knowledge Structures, Common Knowledge, Baire Category, Cantor Sets, Belief Revision

1 Introduction

In Fagin (1994), for every ordinal number γ and sets of agents and primitive propositions a hierarchically constructed canonical multi-agent Kripke structure W_γ was defined for the $S5$ epistemic logic. This canonical structure respresented all the possible truth evaluations for the statements of the logic, with the ordinal numbers representing the levels in the construction of these statements. For every Kripke structure for the same set of agents and primitive propositions canonical maps were defined from the Kripke structure to the canonical structures W_γ. These maps were defined in a straightforward way; one reads off the truth evaluations from the Kripke structure in the domain.

If there is an ordinal such that the map to one of the canonical structures is injective, then the Kripke structure is called *non-flabby*, and the first such ordinal is called the *distinguishing* ordinal. There is another minimal ordinal β, possibly larger than the distinguishing ordinal, for which the image of the Kripke structure in W_β can be extended to any W_γ with $\gamma > \beta$ in only one unique way. This latter ordinal is called the *uniqueness* ordinal. Fagin (1994) proved that the uniqueness ordinal is a limit ordinal and never greater than the next limit ordinal above the distinguishing ordinal. Behind this result lies the concept of common knowledge (Lewis, 1969). The uniqueness of extension depends upon the common knowledge of the agents concerning the structure of possibilities. Common knowledge is completed only through a limiting process; the common knowledge of a statement is an unlimited sequence of statements concerning mutual knowledge of this statement, such as the statement "I know that you know that I know that .... that the proposition p is true."

For the rest of this paper, all multi-agent Kripke structures will be models for the $S5$ epistemic logic, and we assume that there are only finitely many agents and primitive propositions.

Our interest in unique extension originates from the difference between semantic and syntactic formulations of common knowledge. Syntactic common knowledge is a sequence of statements of ever increasing mutual knowledge. Semantic common knowledge is the implicitely shared knowledge of what is possible that is inherent in the structure of a Kripke structure. Syntactic common knowledge may provide much less information that semantic common knowledge. This distinction is essential to understanding the main result of this paper, and we return to it later.

We define a *cell* of a Kripke structure to be a minimal subset that is closed under the possibility considerations of all agents. A subset C is a cell if and only if for every distinct points $x, y \in C$ there is a sequence $x = x_0, x_1, \ldots, x_n = y$ such that for every $0 < i \leq n$ there is some agent at x_{i-1} who considers both x_{i-1} and x_i to be possible, and furthermore for every $x \in C$ and every agent all the points considered possible by this agent at the point x already lie in C. The cell is the semantic formulation of common knowledge; it is the minimal set on which the interactive knowledge of the agents can be defined. It can be argued that all points outside of a cell are irrelevant to the interactive knowledge of the agents inside that cell, and therefore the cells are the important units with regard to the applications of interactive epistemology (for example in economics see Aumann 1976). We define a Kripke structure to be connected if it contains only one cell, (namely the whole structure). It is straightforward to show that a connected Kripke structure maps into only one cell of a canonical structure W_γ.

Of special interest are the canonical Kripke structures W_ω associated with the first infinite ordinal ω, whose elements are called also *knowledge structures*; (see Fagin, Halpern, and Vardi, 1991). Fagin established the necessary and sufficient condition for a cell of W_ω to have the first infinite ordinal ω as its uniqueness ordinal. This condition is the *finite fanout* property, and it means that at every point in the cell every agent considers only a finite number of other points to be possible. There is an alternative formulation for hierarchically constructed canonical Kripke structures corresponding to the ordinal numbers (Heifetz and Samet 1998, 1999), but with regard to the first infinite ordinal they are the same as Fagin's. W_ω is the set of all possible truth assignments for the modal logic formulas of finite expression built from the most primitive set of operations that include at least the knowledge of the agents. The knowledge of an agent in this Kripke structure W_ω is nothing other than the collection of statements of knowledge that this agent posesses concerning these formulas of finite expression.

In Simon (1999) a cell of the Kripke structure W_ω was defined to be *surjective* if all connected Kripke structures that map to it do so surjectively. It is easy to show that a cell of W_ω with the finite fanout property is also surjective.

In this paper, we construct an example of a countable and surjective cell of W_ω that does not have finite fanout. We do not need to consider uncountable cells, since all surjective cells are countable (Theorem 3a, Simon, 1999).

Central to understanding the relation between surjectivity and unique extension is point-set topology. For every Kripke structure we define a topology on this structure based on the formulas of the logic. The base of open subsets corresponds to the set of all formulas; for every formula f the subset of the structure where f is true is a member of this base. For any agent j and a point x of a Kripke structure M, let $P^j_M(x)$ be the subset that agent j considers possible at the point x. With this topology, W_ω is compact. Without explicitely mentioning this topology, Fagin (1994) showed that any set $P^j_{W_\omega}(x)$ is a compact subset of W_ω. An extension to $W_{\omega+1}$ of a knowledge structure x in W_ω is defined by dense subsets R^j of $P^j_{W_\omega}(x)$ containing x for all agents j; therefore there is a unique extension of a cell if and only if every $P^j_{W_\omega}(x)$ in the cell has only one dense subset, which is equivalent to the finiteness of these sets.

Is the lack of a unique extension from a cell of W_ω (equivalently the lack of finite fanout) equivalent to the existence of some Kripke structure mapping to this cell such that the agents have common knowledge on some transfinite ordinal level that some knowledge structure of this cell is not possible? The surjective property is exactly the impossibility of such common knowledge.

It is a natural question, and surprisingly the answer rests upon the property called *centeredness*. The centered property has several equivalent definitions; the most straightfoward definition is that a cell of W_ω is centered if and only if no other cell of W_ω shares the same set of formulas held in common knowledge (Simon 1999). The centered property means that there is no distinction between syntactic and semantic common knowledge. (The set of formulas held in common knowledge is a constant throughout any given cell; see Halpern and Moses 1992). An equivalent formulation is that the cell is an open set relative to the closure of itself. The difference between centered and uncentered cells is radical; if a cell is not centered then it does not contain even one subset that is open relative to the cell's closure, and then there are uncountably many other cells sharing the same set of formulas in common knowledge. Furthermore a centered cell of W_ω is surjective if and only if it has finite fanout (Theorem 3b, Simon 1999).

Due to topological formulations of the centered property, to construct a surjective cell without finite fanout requires some topological insight. Central to our construction is an old result concerning homeomorphisms of a Cantor

Set, which we call Proposition 1. We do not prove Proposition 1 because it is Theorem 9 of Chapter 12 of E. Moise, (1977).

Proposition 1: Let X and Y be two totally disconnected, perfect, compact metric spaces (equivalently Cantor sets) and let X' and Y' be countable and dense subsets of X and Y, respectively. There is a homeomorphism between X and Y that is also a bijection between X' and Y'.

Why is our main result surprising? As stated above, all possibility sets of all agents in W_ω are compact. Using the above notation, the lack of finite fanout implies the existence a cluster point y of some possibility set $P^j_{W_\omega}(x)$ for an agent j. (A cluster point of a set is the limit of a sequence of infinitely many distinct points in that set.) Let C be the cell of W_ω that contains $P^j_{W_\omega}(x)$ (and therefore also y). The point y is a good candidate for an excluded knowledge structure, namely the existence of a Kripke structure that maps to $C\backslash\{y\}$. Given any proper subset D of C, one could define a new Kripke structure $\mathcal{V}(D)$ as the restriction of the existing structure on C to the subset D, meaning that the new possibility set of an agent at a point $x \in D$ is the old possibility set intersected with D. Then one could establish the necessary and sufficient conditions for the canonical map defined from $\mathcal{V}(D)$ to W_ω to be the identity map, meaning that a point $d \in D$ with respect to the Kripke structure $\mathcal{V}(D)$ goes to d in W_ω. If this occurs with any proper subset D, then our cell would not be surjective. The converse holds too, that the image of a Kripke structure is a subset D such that the map from $\mathcal{V}(D)$ back to D is the identity map, (Lemma 6 and Lemma 7 of Simon 1999).

Now the centered property plays an critical role. If the cell C is centered and countable, then D defined to be the set of all isolated points of C is a non-empty subset such that the map from $\mathcal{V}(D)$ back to W_ω is the identity (Theorem 3b, Simon, 1999). (A point is isolated in a set if there is an open subset which contains only this point. In our context a point is isolated in a cell if it can be distinguished from all other points in that cell by a finite set of formulas.) A point y as described above would be not only a cluster point of some possibility set but also a cluster point of the containing cell C, and therefore our set D of isolated points would exclude y. So to find our counter-example, we must work with uncentered cells. But most countable cells that come to mind naturally are centered, because a countable cell is centered if and only if it has at least one isolated point. The usual ideas for countable cells involve sequences of points converging to one or several cluster points that form a proper subset of the cell. However if only one point remains isolated then the cell is centered and cannot be a counter-example. Proving that uncentered countable cells exist is already complicated (Proposition 1 of Simon 1999), let alone finding a counter-example to our problem. In our counter-example, although all points in our cell C will be cluster points of the cell, there will be only one infinite possibility set P (for some agent) and there will be only one cluster point x of P. There will be no proper subset D of C such that the map from $\mathcal{V}(D)$ back to W_ω is the identity map, including the possibility of excluding only this single point x, namely with the set $D = C \backslash \{x\}$.

Why should one care about our main result, the distinction between finite fanout and surjectivity? First, it is intuitive to suspect that the lack of unique extensibility from a cell (or from any Kripke structure) is due ultimately to the possibility of some form of common knowledge concerning the points in this cell that is additional to that given a-priori by the structure of the cell. Our result refutes this intuition. Second, it highlights the distinction between syntactic and semantic common knowledge, bringing the meaning of this distinction to another area that might appear at first to be unrelated. Third, it demonstrates an non-trivial application of topology to philosophical logic. Fourth, we investigated this problem in the context of a very primitive multi-agent modal logic. Ours may be an example of a more general phenomenon, occuring in more complex epistemic logics.

Lastly, our main result is related to belief revision. The operation of creating a new Kripke structure $\mathcal{V}(D)$ from an old structure through restriction to a proper subset D is that of revising what the players know through introducing the common knowledge of the set D. The property that all d in the new Kripke structure $\mathcal{V}(D)$ get mapped back to the same point d in the old structure W_ω means that concerning the truth value of all formulas the results from the belief revision do not contradict the previous beliefs.

In the next section, we provide some background information necessary to understand our solution. In the third and concluding section, we present our example and prove that it has the claimed properties.

2 Background

Let X be a finite set of primitive propositions, and let J be a finite set of agents. Construct the set $\mathcal{L}(X, J)$ of formulas using the sets X and J in the following way:
1) If $x \in X$ then $x \in \mathcal{L}(X, J)$,
2) If $g \in \mathcal{L}(X, J)$ then $(\neg g) \in \mathcal{L}(X, J)$,
3) If $g, h \in \mathcal{L}(X, J)$ then $(g \wedge h) \in \mathcal{L}(X, J)$,
4) If $g \in \mathcal{L}(X, J)$ then $k_j g \in \mathcal{L}(X, J)$ for every $j \in J$,
5) Only formulas constructed through application of the above four rules are members of $\mathcal{L}(X, J)$.
We write simply $\mathcal{L}$ if there is no ambiguity.

$E(f) = E^1(f)$ is defined to be $\wedge_{j \in J} k_j f$, $E^0(f) := f$, and for $i \geq 1$, $E^i(f) := E(E^{i-1}(f))$. A formula $f \in \mathcal{L}(X, J)$ is common knowledge in a subset of formulas $A \subseteq \mathcal{L}(X, J)$ if $E^n f \in A$ for every $n < \infty$.

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 $k_j f$ is also a theorem for all $j \in J$. The axioms are the following, for every $f, g \in \mathcal{L}(X, J)$ 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 $\mathcal{A} \subseteq \mathcal{L}(X, J)$ is called *complete* if for every formula $f \in \mathcal{L}(X, J)$ either $f \in \mathcal{A}$ or $\neg f \in \mathcal{A}$. 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

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

$\Omega(X, J)$ is equivalent to the W_ω of the introduction.

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

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

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 Ω, $\mathcal{L}$ and $\mathcal{Q}^j$ if there is no ambiguity.

In this paper, a Kripke structure is a quintuple $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ where J is a set of agents, for each $j \in J$ $\mathcal{P}^j$ 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.) Define an *adjacency sequence* to be a sequence of points $z_0, z_1, \ldots, z_n$ such that z_i and z_{i-1} both belong to some member of $\mathcal{P}^j$ for some $j \in N$. We define a map $\alpha^{\mathcal{K}} : \mathcal{L}(X,J) \to 2^S$ inductively on the structure of the formulas in the following way:
Case 1 $f = x \in X$**:** $\alpha^{\mathcal{K}}(x) := \psi(x)$.
Case 2 $f = \neg g$**:** $\alpha^{\mathcal{K}}(f) := S \backslash \alpha^{\mathcal{K}}(g)$,
Case 3 $f = g \wedge h$**:** $\alpha^{\mathcal{K}}(f) := \alpha^{\mathcal{K}}(g) \cap \alpha^{\mathcal{K}}(h)$,
Case 4 $f = k_j(g)$**:** $\alpha^{\mathcal{K}}(f) := \{s \mid s \in P \in \mathcal{P}^j \Rightarrow P \subseteq \alpha^{\mathcal{K}}(g)\}$.
We define a map $\phi^{\mathcal{K}} : S \to \Omega(X,J)$ (see Fagin, Halpern, and Vardi 1991) by

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

The map $\phi^{\mathcal{K}}$ is the map to the canonical structure $\Omega(X,J)$ corresponding to the first infinite ordinal.

Consider the map $\psi : X \to 2^{\Omega}$ defined by $\psi(x) := \{z \in \Omega \mid x \in z\}$. We have a Kripke structure $\Omega = (\Omega; J; \mathcal{Q}^1, \ldots, \mathcal{Q}^n; X; \psi)$. (Due to its canonical nature, we index this Kripke structure with Ω.) A possibility set of Ω is defined to be a member of $\mathcal{Q}^j$ for some $j \in J$ and a cell of Ω is a member of the meet partition $\mathcal{Q} := \wedge_{j \in J} \mathcal{Q}^j$.

For every Kripke structure $\mathcal{K} = (S; J; (\mathcal{P}^j | j \in J); X; \psi)$ we define a topology for S, the same as in Samet (1990). Let $\{\alpha^{\mathcal{K}}(f) \mid f \in \mathcal{L}\}$ be the base of open sets of S. We call this the topology induced by the formulas; with the same topology defined on Ω this topology on S is the coarsest topology such that the canonical map $\phi^{\mathcal{K}} : S \to \Omega$ is continuous. The topology of a subset A of S will be the relative topology for which the open sets of A are $\{A \cap O \mid O \text{ is an open set of } S\}$.

Central to this paper is the first part of **Lemma 5** of Simon (1999), which states that if $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ is a Kripke structure and P is a member of $\mathcal{P}^j$ for some $j \in J$ then $\phi^\mu(P)$ is a dense subset of F for some $F \in \mathcal{Q}^j$. This fact was used implicitely by Fagin (1994).

Given a Kripke structure $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ and a subset $A \subseteq S$, we define the Kripke structure $\mathcal{V}^{\mathcal{K}}(A) := (A; J; (\mathcal{P}^j|_A \mid j \in J)); X; \psi|_A)$ where for all $j \in J$ $\mathcal{P}^j|_A := \{F \cap A \mid F \cap A \neq \emptyset \text{ and } F \in \mathcal{P}^j\}$ and for all $x \in X$ $\psi|_A(x) = \psi(x) \cap A$. We define a subset $A \subseteq \Omega$ to be *good* if for every $j \in J$ and every $F \in \mathcal{Q}^j$ satisfying $F \cap A \neq \emptyset$ it follows that $F \cap A$ is dense in F. By Lemma 6 of Simon (1999) A is good if and only if for every $z \in A$ $\phi^{\mathcal{V}^{\mathcal{K}}(A)}(z) = z$.

The next lemmatta relate directly the good property to our problem.

Lemma 7: (Simon 1999) If $\mathcal{K} = (S; J; (\mathcal{P}^j | j \in J); X; \psi)$ is a Kripke structure then $\phi^{\mathcal{K}}(S)$ is a good subset.

Lemma 9: (Simon 1999) 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 $A = C$.

We need a few more facts about $\Omega(X, J)$ for non-empty X. If $|J| \geq 2$ then $\Omega(X, J)$ is topologically equivalent to a Cantor set, (Fagin, Halpern and Vardi 1991). Second we can perceive a Cantor set as $\{0,1\}^{\mathbf{N}}$, where each finite sequence $a = (a^1, a^2, \ldots, a^n)$ defines a cylinder subset $C(a)$ of $\{0,1\}^{\mathbf{N}}$ by $C(a) := \{x \in \{0,1\}^{\mathbf{N}} \mid x^k = a^k \ \forall k \leq n\}$. Furthermore all cylinder subsets are themselves topologically equivalent to Cantor sets, and the same holds for finite unions of cylinder sets. Third, if $|J| \geq 2$ then there exists an uncentered cell of $\Omega(X, J)$ of finite fanout that is dense in $\Omega(X, J)$; an example is constructed in Simon (1999).

3 The Example

We call a partition $\mathcal{P}$ of a topological space D upper (respectively lower) semi-continuous if the set valued correspondence that maps every $d \in D$ to the partition member of $\mathcal{P}$ containing d is an upper (respectively lower) semi-continuous correspondence. (We follow the definitions of Klein and Thompson, 1984.)

Let Ω equal $\Omega(X, \{1,2\})$ with X any finite non-empty set. Let C be an uncentered cell of finite fanout that is dense in Ω. We assume that $\pi : \Omega \to \{0,1\}^{\mathbf{N}}$ is a homeomorphism. For every $n \in \mathbf{N}$ define $\pi_n : \Omega \to \{0,1\}^n$ by $\pi_n(x)$ equaling the $a = (a^1, a^2, \ldots, a^n) \in \{0,1\}^n$ such that $\pi(x)^i = a^i \; \forall \, i \leq n$. If a is the empty sequence in $\{0,1\}^0$ then define $\pi_0(x) := a$ for all $x \in \Omega$.

Let z be any member of C and for every $i = 1, 2, \ldots$ let z_i be a member of C such that $\pi_{2i-2}(z_i) = \pi_{2i-2}(z)$ but $\pi_{2i}(z_i) \neq \pi_{2i}(z)$. For every i define non-empty and mutually disjoint sets $A_{i,1}, A_{i,2}, \ldots A_{i,i}$ in the following way. Let $A_{1,1}$ equal $\Omega \backslash \left(\pi_2^{-1}(\pi_2(z_1)) \cup \pi_2^{-1}(\pi_2(z))\right)$. For $1 \leq k < i$ let $A_{i,k} := \pi_{2i-2}^{-1}(\pi_{2i-2}(z_k)) \backslash \, \pi_{2i}^{-1}(\pi_{2i}(z_k))$ and let $A_{i,i} := \pi_{2i-2}^{-1}(\pi_{2i-2}(z)) \backslash \left(\pi_{2i}^{-1}(\pi_{2i}(z_i)) \cup \pi_{2i}^{-1}(\pi_{2i}(z))\right)$. Because for every $a \in \{0,1\}^{2i}$ there are four members b of $\{0,1\}^{2i+2}$ such that $a = \pi_{2i}(\pi_{2i+2}^{-1}(b))$, all the sets $A_{i,j}$ are non-empty and homeomorphic to Cantor sets. By Proposition 1, for every $i \geq 1$ and $1 \leq k \leq i$ there is a homeomorphism $f_k : A_{i,1} \to A_{i,k}$ such that f_k maps $C \cap A_{i,1}$ bijectively to $C \cap A_{i,k}$ This implies for every $i \geq 1$ that there exists an upper and lower semi-continuous partition $\mathcal{P}_i$ of $C \cap (\cup_{k=1}^i A_{i,k})$ such that every partition member of $\mathcal{P}_i$ has i members, one member in $A_{i,k}$ for every $1 \leq k \leq i$. Notice that all the $A_{i,k}$ are mutually disjoint, meaning that $A_{i,k} = A_{i^*,j^*}$ if and only if $i = i^*$ and $k = k^*$. Furthermore the disjoint union $\cup_{i \geq 1} \cup_{1 \leq k \leq i} A_{i,k}$ is equal to $\Omega \backslash \{z, z_1, z_2, \ldots\}$. Let $\mathcal{P}$ be $\cup_{i=1}^{\infty} \mathcal{P}_i \cup \{z, z_1, z_2, \ldots\}$, a partition of C. It is straightforward to check that $\mathcal{P}$ is upper and lower semi-continuous. We define $\mathcal{A}$ be the Kripke structure $(C; \{1,2,3\}; \mathcal{Q}^1|_C, \mathcal{Q}^2|_C, \mathcal{P}; X, \psi|_C)$.

Lemma 1: If $\mathcal{K} := (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ is a Kripke structure with a topology (not necessarily that induced by the formulas) such that
1) for every $x \in X$ $\psi(x)$ is a clopen set and
2) for every $j \in N$ the partition $\mathcal{P}^j$ is lower and upper semi-continuous,
then the map $\phi^{\mathcal{K}} : S \to \Omega(X, J)$ is continuous.

Proof: It suffices to show that $\alpha^{\mathcal{K}}(f)$ is a clopen set for every $f \in \mathcal{L}(X, J)$. We proceed by induction on the structure of formulas. The claim is true for all $x \in X$ by the hypothesis and it is likewise true for $\neg f$ and $f \wedge g$ if it is true for f and g, due to the clopen property being closed under complementation and finite intersection. For some $f \in \mathcal{L}(X, J)$ we assume that $\alpha^{\mathcal{K}}(f)$ is a clopen set. $\alpha^{\mathcal{K}}(k_j f)$ is an open set by the upper semi-continuity of $\mathcal{P}^j$ and the openness of $\alpha^{\mathcal{K}}(f)$. $S \backslash \alpha^{\mathcal{K}}(k_j f) = \alpha^{\mathcal{K}}(\neg k_j f)$ is an open set by the openness of $S \backslash \alpha^{\mathcal{K}}(f)$ and the lower semi-continuity of $\mathcal{P}^j$.
$\square$

The upper and lower semi-continuity of $\mathcal{Q}^j$ in Ω follow directly from the extension conditions of Fagin, Halpern, and Vardi (1991) that define knowledge structures. The upper and lower semi-continuity of $\mathcal{P}$ relative to the topology on $\Omega(X, \{1,2\})$ induced by the formulas in $\mathcal{L}(X, \{1,2\})$ is essential for our argument, since a-priori we would not know that $\phi^{\mathcal{A}}$ maps $\{z, z_1, \ldots\}$ surjectively to a possibility set of $\Omega(X, \{1,2,3\})$.

Theorem: $\phi^{\mathcal{A}}$ maps C bijectively to a cell of $\Omega(\{1,2,3\})$ that is surjective but without finite fanout.

Proof: We have by Lemma 1 that $\phi^{\mathcal{A}} : C \to \Omega(X, \{1,2,3\})$ is continuous. Since every member of $\mathcal{Q}^1|_C$, $\mathcal{Q}^2|_C$, or $\mathcal{P}$ is compact, their images in $\Omega(X, \{1,2,3\})$ are also compact. By Lemma 9 of Simon (1999) $\phi^{\mathcal{A}}$ maps C surjectively to a cell $\phi^{\mathcal{A}}(C)$ of $\Omega(X, \{1,2,3\})$. Between any two points of $\phi^{\mathcal{A}}(C)$ there is an adjacency path using images of members of $\mathcal{Q}^1|_C$ and $\mathcal{Q}^2|_C$, all finite possibility sets of $\Omega(X, \{1,2,3\})$ – therefore there can be no proper good subset of $\phi^{\mathcal{A}}(C)$. By Lemma 7 of Simon (1999) this implies that $\phi^{\mathcal{A}}(C)$ is a surjective cell. Since for every $f \in \mathcal{L}(X, \{1,2\})$ $\alpha^{\Omega(X,\{1,2\})}(f)$ gets mapped to $\alpha^{\Omega(X,\{1,2,3\})}(f)$, $\phi^{\mathcal{A}}$ is an injective and an open map, and therefore the map $\phi^{\mathcal{A}}$ is also a homeomorphism of C to $\phi^{\mathcal{A}}(C)$. Therefore the image of the one infinite possibility set in C is also an infinite possibility set in the cell $\phi^{\mathcal{A}}(C)$, which implies that this cell of $\Omega(X, \{1,2,3\})$ does not have finite fanout. q.e.d.

4 Acknowledgements

Many helped in finding the best citation for Proposition 1; the lemma has many interesting variations, including the same conclusion for open intervals proven by G. Cantor (1895).

5 References

Aumann, R. (1976), "Agreeing to Disagree," Annals of Statistics 4, pp. 1236-1239.

Bacharach, M., Gerard-Varet, L.A., Mongin, P, and Shin, H., eds. (1997), *Epistemic Logic and the Theory of Games and Decisions*, Dordrecht, Kluwer.

Cantor, G. (1895), "Beitraege zur Begruendung der transfiniten Mengenlehre," Mathematische Annalen, Band 46, pp. 481-512.

Fagin, R. (1994), "A Quantitative Analysis of Modal Logic," Journal of Symbolic Logic 59, pp. 209-252.

Fagin, R., Halpern, Y.J. and Vardi, M.Y. (1991), "A Model-Theoretic Analysis of Knowledge," Journal of the A.C.M. 91 (2), pp. 382-428.

Halpern, J., Moses, Y. (1992), "A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief," Artificial Intelligence 54, pp. 319-379.

Heifetz, A. and Samet, D. (1998), "Knowledge Spaces with Arbitrarily High Rank," Games and Economic Behavior 22, No.2, pp. 260-273.

Heifetz, A. and Samet, D. (1999), "Hierarchies of Knowledge: An Unbounded Stairway", Mathematical Social Sciences, Vol. 38, pp. 157-170.

Hughes, G.E., Cresswell, M.J. (1968), *An Introduction to Modal Logic,* Routledge.

Klein, E., Thompson, A. (1984), *Theory of Correspondences,* Wiley.

Lewis, D. (1969), *Convention: A Philosophical Study,* Harvard University Press.

Moise, E. (1977), *Geometric Topology in Dimensions Two and Three,* Graduate Texts in Mathematics 47, Springer.

Samet, D. (1990) "Ignoring Ignorance and Agreeing to Disagree," Journal of Economic Theory, Vol. 52, No. 1, pp. 190-207.

Simon, R. (1999) "The Difference between Common Knowledge of Formulas and Sets", International Journal of Game Theory, Vol. 28, No. 3, pp. 367- 384.