INSTITUTE OF MATHEMATICAL ECONOMICS

WORKING PAPERS

No. 252

The Difference Between Common Knowledge of Formulas and Sets: Part I

by

Robert Samuel Simon

March 1996

University of Bielefeld

33501 Bielefeld, Germany

The Difference Between Common Knowledge of Formulas and Sets: Part I

Abstract

This article concerns the interactive modal propositional calculus, using the multi-agent epistemic logic $S5$. Let there be at least two agents and let Ω be the state space of maximally consistent sets of formulas. When is the member of the meet partition on Ω generated by the knowledge partitions of the agents determined by the set of formulas held in common knowledge? In part I, this question is investigated for common knowledge generated by a finite set of formulas.

1 Introduction

My investigations of the multi-agent epistemic logic $S5$ began with an attempt to "prove" the opposite of one of the results in this paper! For at least two agents and a non-empty set of primitives, let Ω be the state space of maximally consistent sets of formulas. I was "convinced" that the partitions of Ω generated by the knowledge of the agents would generate a unique meet partition member for which only the tautologies would be held in common knowledge, and that this special meet partition member would, in some sense, cover virtually all of the state space Ω. My ultimate goal was to show that the formulas held in common knowledge at some point of $\omega \in \Omega$ would determine the member of the agents' meet partition containing ω. In this attempt, I failed! (See Simon [Si].) My initial motivation was to show that, except in very special situations, this meet partition would be far too coarse to convey anything meaningful about common knowledge.

The two main results of Part I are the following:

Theorem 1: A set of formulas held in common knowledge at some point in Ω either determines the member of the agents' meet partition or there is an uncountable set of members of the agents' meet partition for which these and no other formulas are held in common knowledge.

Theorem 2: If a set of formulas held in common knowledge at some point in Ω is generated by a finite set of formulas, then this set of formulas held in common knowledge determines the member of the agents' meet partition if and only if this set of formulas is maximal by inclusion among all the sets of formulas that can be held in common knowledge if and only if the subset of the state space holding this set of formulas in common knowledge is finite.

Fagin, Halpern, and Vardi [Fa-Ha-Fa] had proven that there exists at least one meet partition member for which only the tautologies are held in common knowledge. To do this, they used what they called the "no information extension." Given that the knowledge of the agents has been determined only up to some finite rank, the no information extension prescribes a canonical way to extend this knowledge to all higher finite ranks. Theorem 1 and Theorem 2 are proven without any use of the Fagin, Halpern, and Vardi result. However, to show that the above theorems are not empty statements, that indeed there exists at least one uncountable-to-one correspondence between members of the agents' meet partitions and a set of formulas held in common knowledge, the Fagin, Halpern, and Vardi result is used. I don't know if there is a way to prove this existence without using something essentially equivalent to their construction.

Part II concerns common knowledge of formulas that are not finitely generated and the relationship of the above correspondence to other well studied properties of points in Ω. Part III concerns generalizations of this correspondence for Kripke structures of transfinite rank.

2 Background

Construct the set $\mathcal{L}(X,J)$ of legitimate formulas using an alphabet set X of primitive propositions with a set of agents indexed by 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 four above rules are members of $\mathcal{L}(X,J)$.
If there is no ambiguity with regard to X and J, we will use simply $\mathcal{L}$. We define $g \vee h$ to be $\neg(\neg g \wedge \neg h)$ and $g \Rightarrow h$ to be $\neg g \vee h$. $E(f) = E^1(f)$ is defined to be $\wedge_{j \in J} k_j f$ and for $i \geq 2$, $E^i(f) := E(E^{i-1}(f))$.

These symbols have standard interpretations. "x" means that x is true. "$\neg g$" means that means that g is not true. "$g \wedge h$" means that both g and h are true. "$g \vee h$" means that either g or k is true (not necessarily mutually exclusive.) "$k_j g$" means that agent j knows that g is true. "$\neg k_j \neg g$" means that agent j considers g to be possibly true. "$E(g)$" means that all agents know that g is true.

We will work with a standard indexing of a finite set of n agents, namely $\underline{n} := \{1, 2, \ldots, n\}$.

Throughout this article, the multi-agent epistemic logic $S5$ will be assumed, also refered to as $S5_n$ when n is the number of agents. For a discussion of the $S5$ logic system, see Hughes and Cresswell, An Introduction to Modal Logic [Hu-Cr]; and for the multi-agent variation $S5_n$, see Halpern and Moses [Ha-Mo]. Briefly, the $S5_n$ logic system 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 $f \Rightarrow g$ 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 $1 \leq j \leq n$. The axioms are the following, for every $f, g \in \mathcal{L}(X, \underline{n})$ and $1 \leq j \leq n$:

1) all formulas resulting from theorems of the propositional calculus through substitution,

2) $(k_j f \wedge k_j(f \Rightarrow g)) \Rightarrow k_j g$,

3) $k_j f \Rightarrow f$,

4) $k_j f \Rightarrow k_j(k_j f)$,

5) $\neg k_j f \Rightarrow k_j(\neg k_j f)$.

A set of formulas $\mathcal{A} \subseteq \mathcal{L}(X, \underline{n})$ is called "complete" if for every formula $f \in \mathcal{L}(X, \underline{n})$ 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, (using the $S5_n$ logic system.) Define a formula $f \in \mathcal{L}(X, \underline{n})$ to be a tautology (of the logic $S5_n$) if f is in every complete and consistent set of formulas.

A formula $f \in \mathcal{L}(X, J)$ is common knowledge in a set of formulas $\mathcal{A} \subseteq \mathcal{L}(X, J))$ if $f \in \mathcal{A}$ and for every $m \geq 1$ and every function $a : \{1, \ldots, m\} \to J$ the formula $k_{a(m)} \cdots k_{a(1)} f$ is in $\mathcal{A}$ [Le].

Consider any set S with partitions $(\mathcal{P}^j \mid j \in J)$ of S, sometimes called an Aumann structure [Au1]. For each $j \in J$ define a mapping $K_j : \mathcal{P}(S) \to \mathcal{P}(S)$, from the set of subsets of S to itself, by

$$K_j(A) := \{a \in A \mid a \in B \in \mathcal{P}^j \Rightarrow B \subseteq A\} \quad \text{[Au1]}.$$

(Notice that $K_j(A) = \emptyset$ is possible even when $A \neq \emptyset$.) One can interpret $\mathcal{P}^j$ as the collection of sets representing the finest instrument providing discrete measurements available to the jth agent, that is $A \in \mathcal{P}^j$ is a set such that for every $a \in A$ and $b \in S$ the jth agent can discriminate between a and b if only if $b \notin A$ [Au1].

With a set S and partitions $(\mathcal{P}^j \mid j \in J)$ of S one can define a semantic concept of common knowledge. Consider the meet partition $\vee_{j \in J} \mathcal{P}^j$, which is the finest partition equal to or coarser than $\mathcal{P}^j$ for all $j \in J$. The set $A \subseteq S$ is common knowledge at s if and only if $s \in B \in \vee_{j \in J} \mathcal{P}^j$ implies that $B \subseteq A$ [Au1]. Equivalently, one can define A to be common knowledge at $s \in S$ if and only if for all $1 \leq m < \infty$ and functions $a : \{1, \ldots, m\} \to J$ it follows that $s \in K_{a(m)}(\cdots(K_{a(1)}(A))\cdots)$ [Au1].

Using these partitions, one can define the "adjacency" distance between any two points in S as follows: $\rho(s, s') := \min\{d \mid$ there is a sequence

$s = s_0, \ldots s_d$, a function $a : \{1, \ldots, m\} \to J$ and sequence of sets $D_i \in \mathcal{P}^{a(i)}$ such that s_i and s_{i-1} both belong to $D_i\}$, with $\rho(s,s) = 0$ and $\rho(s,s') = \infty$ if there is no such sequence from s to s'. Then the meet partition $\vee_{j \in J} \mathcal{P}^j$ is determined by the equivalence relation for which s is related to s' if and only if $\rho(s,s') < \infty$. It follows that A is common knowledge at s if and only if $\rho(s,s') < \infty$ implies that $s' \in A$. [Au2].

If in addition to a set S and partitions $\{\mathcal{P}^j \mid j \in J\}$ of S we have an alphabet X and a truth assignment $\psi : X \to \mathcal{P}(S)$, the quintuple $\mu = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ is called a Kripke structure for the $S5^n$ logic. (For the rest of this article, it will be called just a Kripke structure.) A Kripke structure is one easy way to generate complete and consistent sets of formulas. We can define a mapping $\alpha^\mu : \mathcal{L}(X,J)) \to \mathcal{P}(S)$ inductively on the structure of the formulas in the following way:

Case 1 $f = x \in X$: $\alpha^\mu(x) := \psi(x)$.
Case 2 $f = \neg g$: $\alpha^\mu(f) := S - \alpha^\mu(g)$,
Case 3 $f = g \wedge h$: $\alpha^\mu(f) := \alpha^\mu(g) \cap \alpha^\mu(h)$,
Case 4 $f = k_j(g)$: $\alpha^\mu(f) := K_j(\alpha^\mu(g))$.

For any point $s \in S$ one can consider the set of formulas defined by

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

Such a set of formulas is complete due to Case 2. Consistency results also from Case 2; the implication of f and $\neg f$ from the use of the multi-agent $S5$ logic system would imply the containment of the point s in both $\alpha^\mu(f)$ and $\alpha^\mu(\neg f)$, a contradiction. (See Hughes and Cresswell [Hu-Cr] and also Fagin, Halpern, and Vardi [Fa-Ha-Va].)

For a Kripke structure $\mu = (S; J; (\mathcal{P}^j \mid j \in J); X; \alpha)$, if $s \in \alpha^\mu(f)$, or equivalently $f \in \phi^\mu(s)$, we say that f is true at s in μ, or with respect to μ. We say that f is valid in the Kripke structure μ if f is true at every $s \in S$ with respect to μ.

For a given alphabet X and the index set $\underline{n}$ of n agents, consider the set of all consistent and complete sets of formulas in $\mathcal{L}(X, \underline{n})$, and give this set of sets of formulas the symbol

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

If there is no ambiguity, we will write simply Ω. Can one consider the space Ω as a Kripke structure itself? Yes. For $1 \le j \le n$ consider the partition Q^j

of Ω generated by the inverse images of the function $\beta^j : \Omega \to \mathcal{P}(\mathcal{L}(X, \underline{n}))$ defined by

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

Consider the mapping $\overline{\psi} : X \to \mathcal{P}(\Omega)$ defined by $\overline{\psi}(x) := \{\omega \in \Omega \mid x \in \omega\}$. Now we have a Kripke structure $\Omega = (\Omega; \underline{n}; \mathcal{Q}^1, \ldots, \mathcal{Q}^n; X; \overline{\psi})$.

Theorem: Every formula valid in Ω is a theorem of the $S5_n$ logic, and vice versa. Furthermore, $\alpha^\Omega(f) = \{\omega \mid f \in \omega\}$ for every $f \in \mathcal{L}(X, \underline{n})$ and $\phi^\Omega(\omega) = \omega$ for every $\omega \in \Omega$.

For a proof of this theorem, see Halpern and Moses [Ha-Mo] and Hughes and Cresswell [Hu-Cr]. Also see Aumann [Au2].

For the purposes of this paper, we will call this result the "Completeness Theorem." The most important lemma used to prove the Completeness Theorem states that any consistent set of formulas can be extended by inclusion to a consistent and complete set of formulas. We will call this the "Extension Lemma."

The Completeness Theorem has fascinating consequences for the meet partition $\mathcal{Q}^1 \vee \cdots \vee \mathcal{Q}^n$ of $\Omega = (\Omega; \underline{n}; \mathcal{Q}^1, \ldots, \mathcal{Q}^n; X; \overline{\psi})$. Assume that $\mu = (S; \underline{n}; \mathcal{P}^1, \ldots, \mathcal{P}^n; X; \overline{\psi})$ is a Kripke structure and f is common knowledge in $\phi^\mu(s)$ for some $s \in S$. If $s \in B \in \mathcal{P}^1 \vee \cdots \vee \mathcal{P}^n$ and $s' \in B$ then it is an easy induction proof to show that f is also common knowledge in $\phi^\mu(s')$. On the other hand, assume for some $B \in \mathcal{P}^1 \vee \cdots \vee \mathcal{P}^n$ that $B \subseteq \alpha^\mu(f)$. If there were some $s \in B$ such that f were not common knowledge in $\phi^\mu(s)$, then $s \notin \alpha^\mu(k_{a(j)} \cdots k_{a(1)} f)$ for some j and $a : \{1, \ldots, j\} \to \{1, \ldots, n\}$ would mean that $A \not\subseteq \alpha^\mu(k_{\alpha(j-1)} \cdots k_{m(1)} f)$ for the $A \in \mathcal{P}^{\alpha(j)}$ containing s. Also by induction we would get an $s' \in B$ with $s' \notin \alpha^\mu(f)$, a contradiction. Therefore we have a nice elementary result:

Lemma 0: If $s \in B \in \mathcal{P}^1 \vee \cdots \vee \mathcal{P}^n$ then f is common knowledge in $\phi^\mu(s)$ if and only if $B \subseteq \alpha^\mu(f)$ (Lemma 4.1, [Ha-Mo]).

For the state space Ω, using the Completeness Theorem, Lemma 0 implies:

For any $B \in \mathcal{Q}^1 \vee \cdots \vee \mathcal{Q}^n$, $\{f \mid f$ is common knowledge in ω for some $\omega \in B\} = \{f : f$ is common knowledge in ω for all $\omega \in B\} = \{f : f \in \omega$ for all $\omega \in B\}$.

We call any member of $\mathcal{Q}^1 \vee \cdots \vee \mathcal{Q}^n$ a cell of Ω.

Define a Kripke structure $\mu = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ to be "connected" if the meet partition $\vee_{j\in J}\mathcal{P}^j$ is a singleton (equal to $\{S\}$,) or equivalently that the adjacency distance between any two elements is finite.

As background for this article, we must consider the canonical hierarchical constructions of Ω. There are two equivalent ways to introduce sets Ω_i whose inverse limit is Ω. In one formulation, Ω_i are the complete and consistent sets of formulas of depth less than or equal to i. In another formulation, the elements of Ω_i are considered to be "worlds." There are 2^X different 0-level worlds, one for each truth assignment on the elements of X. For $i > 0$ an i-level world is a $i-1$ level world plus a determination for each agent which of the $i-1$ level worlds she considers to be possibly true (with certain consistency conditions.) For this article, I will adopt the former construction, though I will use results from the latter and I will call a member of Ω_i a world.

Furthermore, we will also perceive an Ω_i in two ways: first as a seperate Kripke structures in its own right, and second as a canonical projective image of Ω inducing a partition of Ω through inverse images. We define $\mathcal{L}^i := \{f \in \mathcal{L} \mid \text{depth } (f) \leq i\}$ and we define $\Omega_i := \{S \subseteq \mathcal{L}^i \mid S$ is consistent and for every $f \in \mathcal{L}_i$ either $f \in S$ or $f \notin S\}$. Define $\pi_i : \Omega \to \Omega_i$ to be the canonical projection

$$\pi_i(\omega) := \{f \in \mathcal{L}_i \mid f \in \omega\} = \omega \cap \mathcal{L}_i$$

and $\mathcal{F}_i$ will be the partition of Ω induced by the inverse images of π_i,

$$\mathcal{F}_i := \{\pi_i^{-1}(w) \mid w \in \Omega_i\}.$$

Due to the Extension Lemma, the mappings π_i are surjective. Furthermore, from the definition of the Ω_i, $\wedge_{i=1}^{\infty}\mathcal{F}_i$ is the discrete partition of Ω. For every $0 \leq i < \infty$ we consider the Kripke structure $\Omega_i = (\Omega_i; \underline{n}; \overline{\mathcal{F}}_i^1, \ldots, \overline{\mathcal{F}}_i^n; X; \overline{\psi}_i)$; where $\overline{\psi}_i = \pi_i\overline{\psi}$ and for $i > 0$ $\overline{\mathcal{F}}_i^j$ is induced by the inverse images of the function $\beta_i^j : \Omega_i \to \mathcal{P}(\mathcal{L}_{i-1}(X, \underline{n}))$ defined by

$$\beta_i^j(w) := \{f \in \mathcal{L}_{i-1}(X, \underline{n}) \mid k_j(f) \in w\};$$

and define $\overline{\mathcal{F}}_0^j = \{\Omega_0\}$ for every $1 \leq j \leq n$. (To formulate this equivalence relation in the worlds terminology of Fagin, Halpern, and Vardi, we could use equivalently $\neg k_j \neg(f)$ to define β_i^j.) Let $\mathcal{F}_i^j$ be the partition on Ω, coarser than $\mathcal{F}_i$, defined by $\mathcal{F}_i^j := \{\pi_i^{-1}(B) \mid B \in \overline{\mathcal{F}}_i^j\}$. From the definition of the Ω_i

and the $\mathcal{F}_i^j$ it follows that $\wedge_{i=0}^{\infty}\mathcal{F}_i^j = Q^j$. For the sake of notational simplicity, we define $\mathcal{F}_\infty^j$ to be Q^j.

There are several useful properties of the Ω_i.

First, Ω_i is finite for every $0 \leq i < \infty$ (For a more general statement, see Lismont and Mongin [Li-Mo].)

Second, for every $0 < i < \infty$, if $F \in \mathcal{F}_i^j$ and $B \in \mathcal{F}_{i-1}$ with $F \cap B \neq \emptyset$, then every $F' \in \mathcal{F}_{i+1}^j$ contained in F also has a non-empty intersection with B. (See Axiom K2 of Fagin, Halpern, and Vardi [Fa-Ha-Va].) This we call the "Consistency Property."

Third, for every formula $f \in \mathcal{L}^i$ (depth $(f) \leq i$)) and $l \geq i$

$$\pi_l^{-1}(\alpha^{\Omega_l}(f)) = \alpha^{\Omega}(f).$$

(See Lemma 2.5 of Fagin, Halpern, and Vardi [Fa-Ha-Va] and compare with the Completeness Theorem.) This we call the Stability Lemma.

Fourth, for every $0 \leq i < \infty$ the finiteness of Ω_i allows us for every world $w \in \Omega_i$ to define a formula $f(w)$ of depth i such that $\alpha^{\Omega_i}(f(w)) = w$. $f(w)$ is defined inductively on $0 \leq i < \infty$ in the following way: if $i = 0$ then $f(w) = \wedge_{x \in w} x \wedge_{x \notin w} \neg x$. If $i > 0$ and v is the element of Ω_{i-1} such that $v = \pi_{i-1}(\pi_i^{-1}(w))$ and for every j F^j is the member of $\mathcal{F}_i^j$ containing $\pi_i^{-1}(w)$ then

$$f(w) := f(v) \bigwedge_{j \in \underline{n}} \Big(\bigwedge_{u \in \Omega_{i-1},\ \pi_{i-1}^{-1}(u) \cap F^j \neq \emptyset} \neg k_j \neg f(u) \bigwedge_{u \in \Omega_{i-1},\ \pi_{i-1}^{-1}(u) \cap F^j = \emptyset} k_j \neg f(u) \Big).$$

The definition of $f(w)$ follows directly from the "worlds" formulation of Ω_i. (See Fagin, Halpern and Vardi [Fa-Ha-Va] and also Fagin and Vardi [Fa-Va].) We call this the Formula Determination Property. For any subset $A \subseteq \Omega_i$ define $f(A) := \vee_{w \in A} f(w)$.

Fifth, for every Kripke structure $\mu = (S; \underline{n}; \mathcal{P}^1, \ldots, \mathcal{P}^n; X; \psi')$ define the mapping $\phi_i^\mu : S \rightarrow \Omega_i$ by

$$\phi_i^\mu(s) := \phi^\mu(s) \cap \mathcal{L}_i(X, \underline{n}),$$

so that $\cap_{i=0}^{\infty} \pi_i^{-1}(\phi_i^\mu(s)) = \phi^\mu(s)$. It follows from the Stability Lemma that for all $i < \infty$ and $s \in S$ $\phi_i^\mu(s) = \phi^{\Omega_i}(\phi_i^\mu(s)) \cap \mathcal{L}_i(X, \underline{n})$. We call this the

Universal Mapping Property. (Also see Heifetz and Samet [He-Sa].) For the sake of notational convenience we define $\phi^\mu_\infty : S \to \Omega$ to be ϕ^μ. Notice from the definition of α^μ and the operator K_j that if s and s' belong to the same member of $\mathcal{P}^j$ then $\phi^\mu_\infty(s)$ and $\phi^\mu_\infty(s')$ belong to the same member of $\mathcal{F}^j_\infty$.

Two Kripke structures $\mu = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ and $\mu' = (S'; J'; (\mathcal{P}^{j'} \mid j' \in J'); X'; \psi')$ are isomorphic if there are bijections $\gamma_1 : S \to S'$, $\gamma_2 : J \to J'$, and $\gamma_3 : X \to X'$ such that for every $x \in X$ $\gamma_1(\psi(x)) = \psi'(\gamma_3(x))$ and for every pair $x, x^* \in X$ and every $j \in J$, x and x^* share the same member of $\mathcal{P}^j$ if and only if $\gamma_1(x)$ and $\gamma_1(x^*)$ share the same member of $\mathcal{P}^{\gamma_2(j)}$. We say that $\mu = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ and $\mu' = (S'; J; (\mathcal{P}^j \mid j \in J); X; \psi')$ are isomorphic "with fixed ground set and agents" if μ and μ' are isormorphic using the identity maps for γ_2 and γ_3. If μ and μ' are isormorphic with fixed ground set and agents and $\gamma_1 : S \to S'$ is the bijection used, then it follows from the definition of α and ϕ that $\alpha^{\mu'}(f) = \gamma_1(\alpha^\mu(f))$ for all $f \in \mathcal{L}(X, J)$ and $\phi^{\mu'}_i(\gamma_1(s)) = \phi^\mu_i(s)$ for all $s \in S$ and $i \geq 0$, including $i = \infty$. .

Given a Kripke structure $\mu = (S; \underline{n}; \mathcal{P}^1, \dots \mathcal{P}^n; X; \psi)$, let $s \in C \in \mathcal{P}^1 \vee \dots \vee \mathcal{P}^n$. From the definition of α and the finitely constructive nature of any formula, whether or not s is contained in $\alpha^\mu(f)$ for any formula $f \in \mathcal{L}(X, \underline{n})$ is completely determined by the restriction of the truth assignment ψ and the partitions of the agents within the set C. For any union D of members of $\mathcal{P}^1 \vee \dots \vee \mathcal{P}^n$ we define the Kripke structure $\mathcal{V}^\mu(D) := (D; \underline{n}; \mathcal{P}^1|_D, \dots \mathcal{P}^n|_D; X; \psi|_D)$ where $\mathcal{P}^j|_D := \{F \cap D \mid F \cap D \neq \emptyset \text{ and } F \in \mathcal{P}^k\}$ for all $1 \leq j \leq n$ and $\psi|_D(x) = \psi(x) \cap D$ for all $x \in X$. It follows for such a subset D and $s \in D$ that $\phi^{\mathcal{V}^\mu(D)}_i(s) = \phi^\mu_i(s)$ for all $i \geq 0$, including $i = \infty$.

Lastly we need to define a topology for Ω. Let $\{\alpha^\Omega(f) \mid f \in \mathcal{L}\}$ be the base of open sets of Ω. (A topology is defined by the fact that $\alpha^\Omega(f) \cap \alpha^\Omega(g) = \alpha^\Omega(f \wedge g)$.) In this topology Ω is compact for the following reason. $\cup_{f \in S} \alpha^\Omega(f) = \Omega$ for some subset $S \subseteq \mathcal{L}$ is equivalent to $\cap_{f \in S} \alpha^\Omega(\neg f) = \emptyset$, which, by the Extension Lemma, is equivalent to the set of formulas $\{\neg f \mid f \in S\}$ being inconsistent. Since a set of formulas is inconsistent if and only if some finite subset is inconsistent, if $\{\neg f \mid f \in S\}$ is inconsistent then $\{\neg f \mid f \in S'\}$ is inconsistent for some finite subset $S' \subseteq S$, and hence $\cup_{f \in S'} \alpha^\Omega(f) = \Omega$. Furthermore, due to the Formula Determination Property, every member of $\mathcal{F}_i$ is an open and closed set for all $i < \infty$. The topology used on a subset A of Ω will be the relative topology for which the open sets of A are $\{A \cap O \mid O \text{ is an open set of } \Omega\}$.

3 The common knowledge correspondence

For every $\omega \in \Omega$ define $Ck(\omega) := \{f \in \mathcal{L} \mid f \text{ is common knowledge in } \omega\}$. For every set of formulas $T \subseteq \mathcal{L}$ define the set

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

and define the set of formulas $\underline{Ck}(T) \subseteq \mathcal{L}$ by

$$\underline{Ck}(T) := \bigcap_{\omega \in \mathbf{Ck}(T)} Ck(\omega).$$

$\mathbf{Ck}(T)$ are the members of Ω with at least T in common knowledge and $\underline{Ck}(T)$ is the set of formulas whose common knowledge is implied logically in some sense by the common knowledge of T.

Proposition 1: $\underline{Ck}(T) = \{f \in \mathcal{L} \mid$ for every $l < \infty$ there exists an $i(l) < \infty$ and a finite set $T' \subseteq T$ with $(\wedge_{t \in T'} E^{i(l)}(t)) \Rightarrow E^l(f)$ a tautology $\}$

Proof: If the set of formulas $\{E^i(t) \mid i < \infty, t \in T\}$ were inconsistent, then there would be nothing to prove since both sets would be all of $\mathcal{L}$. So in what follows we assume that this set is consistent.

It is straightforward to show that the set on the right is contained in the set on the left.

Let us assume that f is not a formula in the set on the right. That means that there is an $l < \infty$ such that for every pair $i < \infty$ and finite $T' \subseteq T$ $(\wedge_{t \in T'} E^i(t)) \Rightarrow E^l(f)$ is not a tautology. This means that $\{E^i(t) \mid i < \infty, t \in T\} \cup \{\neg E^l(f)\}$ is a consistent set of formulas, contained in some $\omega \in \Omega$ by the Extension Lemma. We see that $f \notin Ck(\omega)$ but $\omega \in \mathbf{Ck}(T)$. q.e.d.

We define $\mathcal{CK} = \{\underline{Ck}(T) \mid T \subseteq \mathcal{L}\} - \{\mathcal{L}\}$. Any member of $\mathcal{CK}$ we call a set of "a-priori" common knowledge; and we say that T generates $\underline{Ck}(T)$. Any set $Ck(\omega)$ for some $\omega \in \Omega$ we will call a set of "actual" common knowledge. Since we defined $\mathcal{CK}$ so that $\mathcal{L}$ is not a member, ($\mathcal{L}$ is generated by any contradictory set of formulas,) all maximal sets of $\mathcal{CK}$ are sets of actual common knowledge (but as we will see later the converse is not true!)

For every $S \in \mathcal{CK}$ we define the correspondence $F(S) \subseteq \mathcal{F}^1_\infty \vee \cdots \vee \mathcal{F}^n_\infty$ by $F(S) := \{C \in \mathcal{F}^1 \vee \cdots \vee \mathcal{F}^n \mid S = Ck(\omega) \text{ for all } \omega \in C\}$. By Lemma 0 $F(S)$ is empty if and only if $S \in \mathcal{CK}$ is not of actual common knowledge.

If $F(S)$ is not empty, then every member of $F(S)$ is dense in $\mathbf{Ck}(S)$. To see this, suppose that a cell $C \in F(S)$ has an empty intersection with $\alpha^{\Omega}(f)$ for some formula $f \in \mathcal{L}$ such that $\alpha^{\Omega}(f) \cap \mathbf{Ck}(S)$ is not empty. Then we can conclude by Lemma 0 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.

For every i and pair $\omega, \omega' \in \Omega$ define $\rho_i(\omega, \omega')$ to be the adjacency-distance in Ω between ω and ω' with respect to the partitions $\mathcal{F}_i^1, \ldots, \mathcal{F}_i^n$. Define ρ_∞ to be the adjacency-distance in Ω between ω and ω' with respect to the partitions $\mathcal{F}_\infty^1, \ldots, \mathcal{F}_\infty^n$.

For every pair of numbers $0 \leq i < \infty$ and $0 \leq d < \infty$ and $\omega \in \Omega$ define the closed set $R_i^d(\omega) := \{\omega' \mid \rho_i(\omega, \omega') \leq d\}$ and define $R_\infty^d(\omega) := \{\omega' \mid \rho_\infty(\omega, \omega') \leq d\}$.

Lemma 1: For every $\omega \in \Omega$ and $0 \leq d < \infty$, $R_\infty^d(\omega) = \cap_{i=0}^\infty R_i^d(\omega)$, and therefore this set is closed.

Proof: It suffices to prove that $\cap_{i=0}^\infty R_i^d(\omega) \subseteq R_\infty^d(\omega)$, since the opposite containment is obvious.

We proceed by induction on d. If $d = 0$ then $\cap_{i=1}^\infty R_i^0(\omega) = \{\omega\} = R_\infty^0$. Let us assume that $\cap_{i=0}^\infty R_i^{d-1}(\omega) \subseteq R_\infty^{d-1}(\omega)$ for all $\omega \in \Omega$. Let $\omega' \in \cap_{i=1}^\infty R_i^d(\omega)$ and define $S_i := R_i^1(\omega') \cap R_i^{d-1}(\omega)$. $\omega' \in \cap_{i=0}^\infty R_i^d(\omega)$ implies that the S_i are a nested sequence of non-increasing non-empty closed sets. Therefore by the compactness of Ω the intersection $\cap_{i=0}^\infty S_i$ is not empty; and we assume that ω^* is a member of this intersection. That $\omega^* \in R_i^1(\omega')$ for all i implies that for some $1 \leq j \leq n$ ω shares with ω' the same partition $\mathcal{F}_i^j$ for all $i < \infty$, and therefore the same is true of $\mathcal{F}_\infty^j$. It follows that $\omega^* \in R_\infty^1(\omega')$. Furthermore, w^* is in $R_\infty^{d-1}(\omega)$ by the induction assumption. q.e.d.

For every $T \subseteq \mathcal{L}$, $\mathbf{Ck}(T)$ is a closed set, because its complement is $\cup_{f \in S^*} \alpha^{\Omega} f$ where $S^* = \{\neg E^i(g) \mid 1 \leq i < \infty, g \in T\}$. Along with Lemma 1, this fact leads to the following theorem.

Theorem 1: For every $S \in \mathcal{CK}$, $F(S)$ is either empty, a singleton, or an uncountable set.

Proof: Let us assume first that $F(S)$ is not empty, and that for some $\omega \in C \in F(S)$ and some $d > 0$ the set $R_\infty^d(\omega)$ is not meagre in $\mathbf{Ck}(S)$ and therefore contains $\alpha^{\Omega}(f) \cap \mathbf{Ck}(S)$ for some formula f with $\alpha^{\Omega}(f) \cap \mathbf{Ck}(S) \neq \emptyset$. Let us assume that ω' is a point in $\mathbf{Ck}(S) - C$. It follows that $\alpha^{\Omega}(f) \cap$

$R^e_\infty(\omega') = \emptyset$ for every $e < \infty$, since otherwise ω' would be in $R^{d+e}_\infty(\omega) \subseteq C$ for some e. Therefore we conclude that $\neg f$ is common knowledge at ω'. However Lemma 0 implies that $\neg f$ cannot be common knowledge at ω since $\emptyset \neq \alpha^\Omega(f) \cap \mathbf{Ck}(S) \subseteq C$. Therefore ω' cannot belong to any member of $F(S)$.

Now we assume that $F(S)$ is not empty and that for every $\omega \in C \in F(S)$ and every d the set $R^d_\infty(\omega)$ is meagre in $\mathbf{Ck}(S)$. For all $\omega \in \mathbf{Ck}(S) - \cup_{C\in F(S)} C$ it follows that $\omega \in \mathbf{Ck}(S \cup \{g\})$ for some $g \notin S$. (If $\mathbf{Ck}(S) = \cup_{C \in F(S)} C$ then the following argument remains valid by the non-existence of such a ω.) By the assumption that $F(S)$ is not empty, and therefore every member of $F(S)$ is dense in $\mathbf{Ck}(S)$, if $g \notin S$ then the closed set $\mathbf{Ck}(S \cup \{g\})$ is meagre in $\mathbf{Ck}(S)$. Since the set of all formulas is countable, $\cup_{g\notin S}\mathbf{Ck}(S \cup \{g\})$ is a countable union of meagre closed sets of $\mathbf{Ck}(S)$. Furthermore, by Lemma 1 and our initial assumption, every $C \in F(S)$ is a countable union of meagre closed sets. Therefore the Baire Category Theorem implies that $F(S)$ is uncountable.
q.e.d.

Corollary 1: If $S \in \mathcal{CK}$ and $F(S)$ is not empty with $C \in F(S)$, then $F(S) = \{C\}$ if and only if there exists some pair $d < \infty$ and $\omega \in C$ such that $R^d_\infty(\omega)$ is not meagre in C.

Proof: It suffices to prove that $R^d_\infty(\omega)$ is not meagre in C if and only if it is not meagre in $\mathbf{Ck}(S)$. Since C is dense in $\mathbf{Ck}(S)$ and $R^d_\infty(\omega)$ is a closed set contained in C it follows for all $f \in \mathcal{L}$ that $R^d_\infty(\omega) \supseteq C \cap \alpha^\Omega(f)$ if and only if $R^d_\infty(\omega) \supseteq \mathbf{Ck}(S) \cap \alpha^\Omega(f)$.
q.e.d.

4 Finitely generated common knowledge

Following Fagin, Halpern, and Vardi [Fa-Ha-Va], for $i > 0$ we define a non-empty set $A \subseteq \Omega_i$ to be "Kripke" closed if and only if for every $j \in N$, every $B \in \mathcal{F}_{i-1}$ and every $w \in A$ if $\pi_i^{-1}(w) \subseteq F \in \mathcal{F}_i^j$ and $F \cap B \neq \emptyset$ then $F \cap B \cap \pi_i^{-1}(A) \neq \emptyset$. If $i = 0$, then any non-empty subset of Ω_0 is allowed to be Kripke closed.

For any Kripke closed set $A \subseteq \Omega_i$ define a Kripke structure $\mathcal{S}(A) = (A; N; \overline{\mathcal{F}}^1_A \ldots \overline{\mathcal{F}}^n_A; X; \overline{\psi}_A)$ in the following way:

$$\overline{\psi}_A(x) = \overline{\psi}(x) \cap A$$

$$\overline{\mathcal{F}}^j_A = \{D \cap A \mid D \in \overline{\mathcal{F}}^j_i,\ D \cap A \neq \emptyset\}.$$

Lemma 2:

(a) If $A \subseteq \Omega_i$ is Kripke closed then for every $w \in A$ $\phi_i^{\mathcal{S}(A)}(w) = w$.

(b) For any Kripke structure $\mu = (S; \underline{n}; \mathcal{P}^1, \ldots, \mathcal{P}^n; X; \psi)$ and $i < \infty$ the image $\phi_i^\mu(S)$ in Ω_i is Kripke closed. (Compare with Proposition 4.20 of Fagin, Halpern, and Vardi [Fa-Ha-Va].) Furthermore, if μ is connected then $\mathcal{S}(\phi_i^\mu(S))$ is connected for all $i < \infty$, $\phi_\infty^\mu(S)$ is contained in a cell of Ω and it is dense in this cell.

(c) If C is a finite cell with its restricted Kripke structure $\mathcal{V}^\Omega(C) = (C; \underline{n}; \mathcal{F}^1_\infty|_C, \ldots, \mathcal{F}^n_\infty|_C; X; \overline{\psi}|_C)$ and i is large enough so that $\pi_{i-1} : C \to \Omega_{i-1}$ is injective, then $\mathcal{S}(\phi_i^\Omega(C))$ is isomorphic with fixed ground set and agents to $\mathcal{V}^\Omega(C)$, with $\pi_i|_C$ providing the corresponding bijection between the set C and the Kripke closed set $\phi_i^\Omega(C) \subseteq \Omega_i$. (Compare with Theorem 4.23 of Fagin, Halpern, and Vardi [Fa-Ha-Va].)

(d) If $A \subseteq \Omega_i$ is Kripke closed and $\mathcal{S}(A)$ is connected, then there is a unique finite cell C of Ω such that $\phi_\infty^{\mathcal{S}(A)}(A) = C$ and furthermore $\mathcal{V}^\Omega(C)$ and $\mathcal{S}(A)$ are isomorphic with fixed ground set and agents by the bijection $\pi_i|_C =: C \to A$ or its inverse $\phi_\infty^{\mathcal{S}(A)} : A \to C$.

Proof:

(a) It suffices by the Universal Mapping Property to prove for every formula $f \in \mathcal{L}_i$ that f is true at $w \in A$ as a member of the Kripke structure $\mathcal{S}(A)$ if and only if f is true at w as a member of the Kripke structure Ω_i. We proceed by induction on the structure of formulas. If the depth of f is zero, the claim follows directly from the definition of $\overline{\psi}_A$. Likewise if the claim is true for f and g then it is true for either $\neg f$ or $f \wedge g$ directly from the definition of $\alpha^{\mathcal{S}(A)}$ and α^{Ω_i}.

Let us assume that $w \in A$, depth $(f) < i$, and $k_j(f)$ is true at the world w as a member of the Kripke structure Ω_i. Let $w \in F \in \overline{\mathcal{F}}^j_A$ and $w \in F' \in \overline{\mathcal{F}}^j_i$, so that $F' \subseteq \alpha^{\Omega_i}(f)$. $F = A \cap F'$ by the definition of $\mathcal{F}^j_A$, so it follows that

$F = A \cap F' \subseteq \alpha^{\Omega_i}(f) \cap A = \alpha^{\mathcal{S}(A)}(f)$, the last equality following by the induction hypothesis.

On the other hand, assume that $w \in A$, depth $(f) < i$, and $\neg k_j(f)$ is true at the world w as a member of Ω_i. That means that agent j considers $\neg f$ possible at w as a member of Ω_i, or that $F' \cap \alpha^{\Omega_i}(\neg f) \neq \emptyset$ given that $w \in F' \in \overline{\mathcal{F}}^j_i$. Let $B \in \mathcal{F}_{i-1}$ satisfy $\neg f$ true in B and $\pi_i(B) \cap F' \neq \emptyset$. (The existence of such a B is guaranteed by the Stability Lemma.) By the Kripke closed property of A there is some world $w' \in A$ with $w' \in \pi_i(B) \cap F'$. Let $A \cap F' = F \in \overline{\mathcal{F}}^j_A$ and we have that $w' \in \pi_i(B) \cap F$. Since $\pi_i(B) \subseteq \alpha^{\Omega_i}(\neg f)$, the induction hypothesis implies that $A \cap \pi_i(B) \subseteq \alpha^{\mathcal{S}(A)}(\neg f)$. Hence $F \cap \alpha^{\mathcal{S}(A)}(\neg f) \neq \emptyset$ and $\neg k_j \neg(\neg f)$ and $\neg k_j(f)$ are true at w as a member of the Kripke structure $\mathcal{S}(A)$.

(b) Let $w = \phi^\mu_i(s) \in \Omega_i$. Let $\pi_i^{-1}(w) \in F \in \mathcal{F}^j_i$ and assume that $B \cap F \neq \emptyset$ for some $B \in \mathcal{F}_{i-1}$. Consider the formula $\neg k_j \neg f(\pi_{i-1}(B))$. Since depth $(\neg k_j \neg f(\pi_{i-1}(B))) = i$ and $\neg k_j \neg f(\pi_{i-1}(B))$ is true at $\phi^\mu_i(s)$ with respect to Ω_i, it follows by the Universal Mapping Property that $\neg k_j \neg f(\pi_{i-1}(B))$ is true at $s \in S$. Therefore there exists an $s' \in S$ such that $\phi^\mu_\infty(s') \in B$ and s' belongs to the same member of $\mathcal{P}^j$ as s. It follows that $\phi^\mu_i(s') \in \pi_i(F) \cap \pi_i(B)$ since $\phi^\mu_\infty(s)$ and $\phi^\mu_\infty(s')$ must belong to the same member of $\mathcal{F}^j_\infty$ and $\mathcal{F}^j_i$ is coarser than $\mathcal{F}^j_\infty$.

If there is a sequence $s = s_0, \ldots s_m = s'$ and a function $a : \{1, \ldots, m\} \to J$ such that for every $k = 1, \ldots, m$ s_k and s_{k-1} both belong to the same member of $\mathcal{P}^{a(k)}$ then each pair $\phi^\mu_\infty(s_k)$ and $\phi^\mu_\infty(s_{k-1})$ both belong to the same member of $\mathcal{F}^{a(k)}_\infty$. Therefore the adjacency distance between any two points in $\phi^\mu_\infty(S)$ with respect to the Kripke structure Ω is also finite, and $\phi^\mu_\infty(S)$ is contained in a cell. Since $\mathcal{F}^{a(k)}_i$ is coarser than $\mathcal{F}^{a(k)}_\infty$ for all $i < \infty$, we also conclude that $\phi^\mu_i(s_k)$ and $\phi^\mu_i(s_{k-1})$ belong to the same member of $\overline{\mathcal{F}}^{a(k)}_i$. But from the definition of $\overline{\mathcal{F}}^j_{\phi^\mu_i(S)}$ we must also conclude that $\phi^\mu_i(s_m)$ and $\phi^\mu_i(s_{m-1})$ belong to the same member of $\overline{\mathcal{F}}^{a(k)}_{\phi^\mu_i(S)}$, and likewise that $\mathcal{S}(\phi^\mu_i(S))$ is connected.

Assume that there exists a formula f with $\emptyset \neq C \cap \alpha^\Omega(f) \subseteq C - \phi^\mu_\infty(S)$. This means that $\neg f$ is valid in μ, and hence by Lemma 0 that $\neg f$ is common knowledge in $\phi^\mu_\infty(s)$ for every $s \in S$. But we assumed that f is true at some members of C, a contradiction to Lemma 0. (Warning: $\phi^\mu_\infty(S)$ may be a proper subset of a cell!)

(c) For all $l < \infty$ and $c \in C$, $\pi_l(c) = \phi_l^{\Omega}(c)$ holds by the Completeness Theorem. Let $A = \pi_i(C) = \phi_i^{\Omega}(C)$. Because of the definition of $\overline{\psi}_A$ it suffices to show that for all pairs $c, c' \in C$ the worlds $\pi_i(c)$ and $\pi_i(c')$ belong to the same member of $\overline{\mathcal{F}}^j_A$ if and only if c and c' belong to the same member of $\mathcal{F}^j_\infty$. By the definition of $\overline{\mathcal{F}}^j_A$, if $\pi_i(c)$ and $\pi_i(c')$ don't belong to the same member of $\overline{\mathcal{F}}^j_A$ they won't belong to the same member of $\overline{\mathcal{F}}^j_i$ either, and hence c and c' also won't belong to the same member of $\mathcal{F}^j_\infty$. For $m \geq i$ let $A(m)$ be $\pi_m(C) = \phi_m^{\Omega}(C)$, so that $A(i) = A$. We presume, for the sake of contradiction, that $\pi_i(c)$ and $\pi_i(c')$ belong to the same member of $\mathcal{F}^j_A$, meaning also that c and c' belong to the same member of $\mathcal{F}^j_i$, but c and c' don't belong to the same member of $\mathcal{F}^j_\infty$. Since $\mathcal{F}^j_\infty = \wedge_{i=0}^{\infty}\mathcal{F}^j_i$, we must assume that there is a maximal l with the property that $\pi_l(c)$ and $\pi_l(c')$ do belong to the same member of $\overline{\mathcal{F}}^j_{A(l)}$. Consider the member F of $\overline{\mathcal{F}}^j_{A(l)}$ containing both $\pi_l(c)$ and $\pi_l(c')$ and the member F' of $\overline{\mathcal{F}^j}_{A(l+1)}$ containing $\pi_{l+1}(c)$ but not $\pi_{l+1}(c')$. Consider the member G of $\mathcal{F}_{l-1}$ containing c'. It follows that $\pi_l(G)$ and F have a non-empty intersection. By the Consistency Property $\pi_{l+1}(G)$ and $F^* \in \mathcal{F}^j_{l+1}$ have a non-empty intersection, where $F' = F^* \cap A(l+1)$. Because $A(l+1)$ is a Kripke closed set by (b), there must be a $c^* \in C$ with $c^* \in G \in \mathcal{F}_{l-1}$ and $\pi_{l+1}(c^*) \in F'$. But since $l \geq i$ and $\pi_{i-1} : C \rightarrow \Omega_{i-1}$ is an injection, c^* must be c', a contradiction.

(d) By (b) $\phi_\infty^{\mathcal{S}(A)}(A)$ is a dense subset of a cell C. But $\phi_\infty^{\mathcal{S}(A)}(A)$ is finite, hence closed, and therefore equals C.

As before, for $l \geq i$ define $A(l) := \phi_l^{\mathcal{S}(A)}(A) = \pi_l(\phi_\infty^{\mathcal{S}(A)}(A)) = \pi_l(C)$, the middle equality following by the definition of ϕ. By (c), for $l > i$ we have $\mathcal{S}(A(l))$ isomorphic with fixed ground set and agents to $\mathcal{V}^{\Omega}(C)$ using the mapping $\pi_l|_C$. By (a), $\phi_{i+1}^{\mathcal{S}(A)} : A \rightarrow A(i+1)$ is bijective.. It suffices to show an isomorphism between $\mathcal{S}(A)$ and $\mathcal{S}(A(i+1))$ using the mapping $\phi_{i+1}^{\mathcal{S}(A)} : A \rightarrow A(i+1)$. Notice by (a) that $(\pi_i\pi_{i+1}^{-1})|_{A(i+1)}$ and $\phi_{i+1}^{\mathcal{S}(A)}$ are inverses.

The truth assignment $\overline{\psi}_{A(i+1)}$ satisfies $\overline{\psi}_{A(i+1)} = \phi_{i+1}^{\mathcal{S}(A)}\overline{\psi}_A$. It suffices to show for every $1 \leq j \leq n$ that

(1) $w \in A$ and $w' \in A$ belong to the same member of $\overline{\mathcal{F}}^j_A$

if and only if

(2) $\phi_i^{\mathcal{S}(A)}(w)$ and $\phi_{i+1}^{\mathcal{S}(A)}(w')$ belong to the same member of $\overline{\mathcal{F}}^j_{A(i+1)}$.

(1) implies that $\phi_\infty^{\mathcal{S}(A)}(w)$ and $\phi_\infty^{\mathcal{S}(A)}(w')$ belong to the same member of $\mathcal{F}^j_\infty$,

which in turn implies that $\phi_i^{S(A)}(w)$ and $\phi_i^{S(A)}(w')$ belong to the same member of $\overline{\mathcal{F}}_i^j$ and also (2) from the definition of $\overline{\mathcal{F}}^j_{A(i+1)}$. On the other hand, (2) implies that $\phi_{i+1}^{S(A)}(w)$ and $\phi_{i+1}^{S(A)}(w')$ belong to the same member of $\overline{\mathcal{F}}^j_{i+1}$. But $\mathcal{F}^j_{i+1}$ is a refinement of $\mathcal{F}^j_i$ as a partition of Ω, which implies that $w \in A$ and $w' \in A$ belong to the same member of $\overline{\mathcal{F}}^j_i$. Again by the definition of $\overline{\mathcal{F}}^j_A$ this implies (1).

Lastly, if $c \in \phi_\infty^{S(A)}(A) = C$ then $\phi_\infty^\Omega(c) = \phi_\infty^{\nu^\Omega(C)}(c) = \phi_\infty^{S(A(l))}(\pi_l(c))$, the last equality by the isomorphism just proved. The Completeness Theorem implies that $c = \phi_\infty^\Omega(c)$, meaning that $c = \phi_\infty^{S(A(l))}(\pi_l(c))$.

Theorem 2: Let C be any cell and let S and T be sets of formulas such that $S = \underline{Ck}(T), C \in F(S)$, and $|T| < \infty$. Then the following are equivalent:

1) C is a finite set,

2) $C = \mathbf{Ck}(T)$,

3) S is maximal in $\mathcal{CK}$,

4) S is maximal among all the members of $\mathcal{CK}$ generated by finite sets of formulas,

5) $F(S)$ is a singleton, namely $\{C\}$.

Proof: Since finite conjuctions of formulas are formulas, we will assume that $T = \{g\}$ for a single formula g. Let d be the depth of g.

(1) $\Rightarrow$ (2): Being finite, C is a closed set. Since C is dense in $\mathbf{Ck}(\{g\})$ it follows that $C = \mathbf{Ck}(\{g\})$.

(2) $\Rightarrow$ (3), (3) $\Rightarrow$ (4), (2) $\Rightarrow$ (5): All three implications are obvious.

For all $i < \infty$ $\mathcal{S}(\pi_i(\mathbf{Ck}(\{g\}))$ is connected for every $i < \infty$, since by Lemma 2b $\mathcal{S}(\pi_i(C))$ is connected and by the density of C in $\mathbf{Ck}(\{g\})$ we have $\pi_i(C) = \pi_i(\mathbf{Ck}(\{g\})$.

(4) $\Rightarrow$ (1) and (2): Consider $A := \pi_d(\mathbf{Ck}(\{g\}))$. Since depth $(g) = d$ the Stability Lemma implies that $\pi_d^{-1}(A) \subseteq \alpha^\Omega(g)$. For $l \geq d$ define

$A(l) := \pi_l(\phi_\infty^{\mathcal{S}(A)}(A)) = \phi_l^{\mathcal{S}(A)}(A)$, (the latter equality holding by the definition of ϕ.) Consider the formulas $f(A(l))$. By Lemma 2a we have that $\pi_l^{-1}(A(l)) \subseteq \pi_d^{-1}(A) \subseteq \alpha^\Omega(g)$ for all $l \geq d$, and therefore $f(A(l)) \Rightarrow f(A) \Rightarrow g$ is a tautology for all $l \geq d$ since $\pi_l^{-1}(A(l)) = \alpha^\Omega(f(A(l)))$ by the Stability Lemma. Also by Lemma 2a and the Stability Lemma we have for all $l \geq d$ that $f(A(l))$ is valid in $\mathcal{S}(A(l))$, so that $\phi_\infty^{\mathcal{S}(A)}(A) = \phi_\infty^{\mathcal{S}(A(l))}(A(l)) \subseteq \mathbf{Ck}(\{f(A(l))\}) \subseteq \mathbf{Ck}(S)$, the first equality by Lemma 2d. It follows from (4) that $\mathbf{Ck}(S) = \mathbf{Ck}(\{f(A(l))\})$ for all $l \geq d$. By Lemma 2d there is a finite cell C' of Ω such that $C' = \phi_\infty^{\mathcal{S}(A)}(A)$. But then it follows that $C' \subseteq \mathbf{Ck}(S) = \cap_{l=d}^\infty \mathbf{Ck}(\{f(A(l))\}) \subseteq \cap_{l=d}^\infty \alpha^\Omega(f(A(l)) = \cap_{l=d}^\infty \pi_l^{-1}(A(l)) = \cap_{l=d}^\infty \pi_l^{-1}(\phi_l^{\mathcal{S}(A)}(A)) = \phi_\infty^{\mathcal{S}(A)}(A) = C'$, and we must conclude that $C' = C$.

(5) $\Rightarrow$ (1) By the proof of Theorem 1, C contains a non-empty open set of $\mathbf{Ck}(\{g\})$. Without loss of generality, we assume that the non-empty open set of $\mathbf{Ck}(\{g\})$ is $W^* \cap \mathbf{Ck}(\{g\})$ for some $W^* \in \mathcal{F}_i$ and $i \geq d$. Consider the set $A := \pi_i(\mathbf{Ck}(\{g\}))$. As in the proof of (4) $\Rightarrow$ (1), we have that $f(A)$ is valid in $\mathcal{S}(A)$ and we can conclude from $i \geq d$ that $f(A) \Rightarrow g$ is a tautology and $\phi_\infty^{\mathcal{S}(A)}(A) \subseteq \mathbf{Ck}(\{g\})$. Since $W^* \cap \pi_i^{-1}(A) \supseteq W^* \cap \mathbf{Ck}(\{g\}) \neq \emptyset$ and $W^* \in \mathcal{F}_i$, we conclude by Lemma 2a that $\phi_\infty^{\mathcal{S}(A)}(A)$ has a non-empty intersection with W^*. But since $\phi_\infty^{\mathcal{S}(A)}(A) \subseteq \mathbf{Ck}(\{g\})$ and C contains the non-empty set $\mathbf{Ck}(\{g\}) \cap W^*$, we must conclude that $C \cap \phi_\infty^{\mathcal{S}(A)}(A) \neq \emptyset$. But by Lemma 2d $\phi_\infty^{\mathcal{S}(A)}(A)$ is a cell, hence $C = \phi_\infty^{\mathcal{S}(A)}(A)$, a finite set. q.e.d.

In fact, we can say more about the structure of finitely generated common knowledge.

Using what they called the "least information extension," Fagin, Halpern, and Vardi proved the following (Theorem 4.22).[Fa-Ha-Va]:

If X is not empty, $n \geq 2$ and $A \subseteq \Omega_i$ is a Kripke closed set then $\mathcal{S}(A)$ is connected if and only if $F(\underline{Ck}(f(A)))$ is not empty.

This leads directly to the following result:

Corollary 2: If $n \geq 2$ then there exists a uncountable number of dense cells of Ω. (Also see Simon [Si].)

Proof: Letting $T = \emptyset$ it follows that $\pi_0(\mathbf{Ck}(\emptyset)) = \Omega_0$. Not only is all of Ω_0 trivially Kripke closed, it is also trivially connected, since $\overline{\mathcal{F}}_0^j = \{\Omega_0\}$ for every $1 \leq j \leq n$! By Theorem 2 and the above mentioned Fagin, Halpern, and Vardi result, it suffices to notice that there exists some non-tautological formula that can be held in common knowledge! q.e.d.

After the following lemma, we include two additional results that round out our knowledge of the finitely generated members of $\mathcal{CK}$. Both are corollaries of results of Fagin, Halpern, and Vardi.

Lemma 3:
(a) If $D \in \overline{\mathcal{F}}_A^1 \vee \cdots \vee \overline{\mathcal{F}}_A^n$ then D as a subset of Ω_i is also Kripke closed and $\mathcal{S}(D)$ is connected.

(b) If $A \subseteq \Omega_i$ is Kripke closed, then the members of $\overline{\mathcal{F}}_A^1 \vee \cdots \vee \overline{\mathcal{F}}_A^n$ are the maximal Kripke closed sets both contained in A and whose corresponding Kripke structures are connected.

Proof:
(a) Lemma 2a implies that $\phi_i^{\mathcal{V}^{\mathcal{S}(A)}(D)}(D) = \phi_i^{\mathcal{S}(A)}(D) = D$; therefore it follows by Lemma 2b that D is Kripke closed. Since $\mathcal{V}^{\mathcal{S}(A)}(D)$ is connected, again by Lemma 2b $\mathcal{S}(D)$ is connected.

(b) Assume that $D \subseteq A$ is Kripke closed and $\mathcal{S}(D)$ is connected. From the definition of $\mathcal{S}(D)$ and $\mathcal{S}(A)$, for any $1 \leq j \leq n$ a member of $\overline{\mathcal{F}}_D^j$ is contained in a member of $\overline{\mathcal{F}}_A^j$. Therefore if $\mathcal{S}(D)$ is connected the adjacency distance in $\mathcal{S}(A)$ between any two members of D is also finite. This implies that D is contained in some member of $\overline{\mathcal{F}}_A^1 \vee \cdots \vee \overline{\mathcal{F}}_A^n$. q.e.d.

Proposition 2: For every finite set of formulas T with $\max_{f \in T}$ depth $(f) = d$, and $\mathbf{Ck}(T) \neq \emptyset$ there is a finite set of formulas $T' = \{f_1, \ldots, f_k\}$, all of depth no greater than d, such that $\cup_{i=1}^k \mathbf{Ck}(\{f_i\}) = \mathbf{Ck}(T)$, this union is disjoint, and the $\underline{CK}(\{f_i\})$ are the minimal members of $\{S \in \mathcal{CK} \mid S \supseteq T, F(S) \neq \emptyset\}$.

Proof: Define $A := \pi_d(\mathbf{Ck}(T))$. Let $\{A_1, \ldots, A_k\} = \overline{\mathcal{F}_A^1} \vee \cdots \vee \overline{\mathcal{F}_A^n}$, and for every $1 \leq m \leq k$ let $f_m := f(A_m)$, so that depth $(f_m) = d$.

By Lemma 3a and the above mentioned theorem of Fagin, Halpern, and Vardi, $F(\{f_m\})$ is not empty for every $1 \leq m \leq k$. Since the $\alpha^{\Omega_d}(f_1) = A_1, \ldots, \alpha^{\Omega_d}(f_k) = A_k$ are disjoint in Ω_d, by the Stability Lemma the $\mathbf{Ck}(\{f_1\}), \ldots, \mathbf{Ck}(\{f_k\})$ are disjoint in Ω.

Now assume that $S \in \mathcal{CK}$, $S \supseteq T$, and $C \in F(S)$. Consider $\pi_d(C) \subseteq \pi_d(\mathbf{Ck}(T)) = A \subseteq \Omega_i$. $\pi_d(C)$ is equal to $\phi_d^\Omega(C)$ by the Completeness Theorem. By Lemma 2b $\pi_d(C)$ is Kripke closed and $\mathcal{S}(\pi_d(C))$ is connected, and hence by Lemma 3b $\pi_d(C) \subseteq A_m$ for some $1 \leq m \leq k$. It follows by the Stability Lemma that f_m is valid in $\mathcal{V}^\Omega(C)$, so that $C \subseteq \mathbf{Ck}(f(A(m))$. q.e.d.

Proposition 3: If $i < \infty$, $A \subseteq \Omega_i$ is Kripke closed, $\mathcal{S}(A)$ is connected, and $A(i+1) = \phi_{i+1}^{\mathcal{S}(A)}(A)$ then $\phi_\infty^{\mathcal{S}(A)}(A) = \mathbf{Ck}(\{f(A(i+1))\})$.

Proof: By Lemma 2d $\phi_\infty^{\mathcal{S}(A)}(A)$ is a finite cell C such that $\pi_i|_C : C \rightarrow A \subseteq \Omega_i$ is a bijection. For such a case, in the proof of Theorem 4.23, Fagin, Halpern, and Vardi [Fa-Ha-Va] showed that $\{\phi_\infty^\Omega(c)\} = \pi_{i+1}^{-1}\pi_{i+1}(c) \cap \mathbf{Ck}(f(\phi_{i+1}^\Omega(C)))$ for every $c \in C$. Since $\phi_{i+1}^\Omega(C) = A(i+1)$, by the Stability Lemma we have $\pi_{i+1}^{-1}\pi_{i+1}(C) = \alpha^\Omega(f(A(i+1)))$. It follows that $\mathbf{Ck}(\{f(A(i+1))\}) = \alpha^\Omega(f(A(i+1))) \cap \mathbf{Ck}(\{f(A(i+1))\}) = \pi_{i+1}^{-1}\pi_{i+1}(C) \cap \mathbf{Ck}(\{f(\phi_{i+1}^\Omega(C))\}) = \phi_\infty^\Omega(C)$. The rest follows by the isomorphism of Lemma 2d.

q.e.d.

Does the structure of $\mathcal{CK}$ and the correspondence F behave in general so nicely as they do for finitely generated members of $\mathcal{CK}$? That is the subject of Part II.

5 Acknowledgements

The inquiry into the cardinality of dense minimal common knowledge worlds was inspired by communication with Robert Aumann and Dov Samet, and later with Aviad Heifetz. Joseph Halpern has not only informed me of existing related literature and discoveries, he has also helped me to explore the open problems related to the results of Part I, some of whose solutions will appear in Part II.

6 References

[Au1] Aumann, R. (1976), Agreeing to Disagree, Annals of Statistics 4, pp. 1236-1239.

[Au2] Aumann, R. (1989), Notes on Interactive Epistemology, Cowles Foundation for Research in Economics working paper.

[Fa-Va] Fagin, R., Vardi, M.Y. (1985), Proceedings of the 17th Annual A. C. M. Symposium on Theory of Computing, Providence, R. I., May 6-8, 1985, pp. 305-315.

[Fa-Ha-Va] 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.

[He-Sa] Heifetz, A., Samet, D., Universal Partition Spaces, in preparation.

[Hu-Cr] Hughes, G.E., Cresswell, M.J. (1968), An Introduction to Modal Logic, Routledge.

[Le] Lewis, D. (1969), Convention: A Philosophical Study, Harvard University Press.

[Li-Mo] Lismont, L. and Mongin, P. (1993), Belief Closure: A Semantics of Common Knowledge for Modal Propositional Logic, Mathematical Social Science, to appear.

[Si] Simon, R. (1995), Alienated Extensions and Common Knowledge Worlds, Working Paper 242, (April 1995), Institute of Mathematical Economics, University of Bielefeld, Germany.