SURJECTIVITY AND FINITE FANOUT

ZIV HELLMAN AND ROBERT SAMUEL SIMON

ABSTRACT. There are situations in which players have common knowledge of a fact, yet there is no way for them to state this explicitly. One way to ascertain when this happens is to map semantic knowledge models into purely syntactic models of knowledge in a canonical manner, taking each state into the complete and consistent set of formulae that hold true at that state. If a common knowledge component in the domain is mapped non-surjectively to a common knowledge component in the range, a property called non-surjectivity, then by definition something known in common in the domain has been excluded from explicit expression. Simon (1999) shows that a sufficient condition for surjectivity is for every partition element in the partition profile to contain only a finite number of states, a property termed finite fanout. The converse, however, does not hold. We construct here an example of a knowledge structure that is surjective yet lacks finite fanout.

Keywords: Semantic knowledge structures, Syntactic knowledge structures, Common Knowledge, Surjectivity, Electronic mail game
JEL classification:

1. INTRODUCTION

There are situations in which players have common knowledge of a fact, yet there is no way for them to state this explicitly. An example of this can be seen in the following three-player game, loosely based on the two-player electronic mail game of Rubinstein (1989). We present it first from the perspective of a semantic model, along the lines initiated by Robert Aumann's seminal paper, Aumann (1976).

There are three players. The space of states $\mathcal{K}$ is labelled by the non-negative integers, 0, 1, ... There is one state of nature, p, that is false only at 0 and holds true at all other states. The partitions of the players are $\{\{0\},\{1,2\},\ldots\}$ for player 1, $\{\{0,1\},\{2,3\}\ldots\}$ for player 2, and the trivial partition $\{\{0,1,\ldots\}\}$ for player 3.

Date: February 2014.
Department of Economics, Bar Ilan University, Ramat Gan, Israel, and Department of Mathematics, London School of Economics, London, UK.

The states in $\mathcal{K}$ are obviously distinguished by the labels given to them, but at a deeper level they can be distinguished by the hierarchies of mutual knowledge associated with them. For example, at state 1, the formulae k_1p, $\neg k_2p$ hold, where k_1 and k_2 stand for 'player 1 knows' and 'player 2 knows', respectively. At state 2, k_1p, k_2p, k_2k_1p, but $\neg k_2k_1k_2p$ and $\neg k_1k_2k_1p$. At state 3, k_1p, $k_2p, k_2k_1p, k_1k_2k_1p$, but $\neg k_1k_2k_1k_2p$ and $\neg k_2k_1k_2p$. At state 4, $k_2k_1k_2p$, but $\neg k_2k_1k_2k_1k_2p$. And so on.

One can, alternatively, analyse all this syntactically. From the syntactic perspective, it is the formulae of the form k_2p, $\neg k_1k_2k_1p$, $k_1\neg k_3k_2p$, and so on, that are the primitive givens. Out of these one forms a state space Ω, where each 'state of the world' is a consistent and complete collection of such formulae. A partition element of a player j at a state z is then the set of all formulae that player j knows to hold at z, as expressed by the operator k_j.

One advantage of the syntactic approach is that while a semantic model may contain 'extraneous' structure due to the initial choice of states (which are primitives in the semantic approach), the syntactic model is purely epistemic, with everything built out of 'what is true and what is known', nothing added beyond. Furthermore, a semantic model can be mapped to a corresponding syntactic model using a canonical mapping ϕ, defined by $\phi(s)$, for each state s in the semantic model, being the set of formulae that hold true at s. This provides a tool for studying aspects of the semantic model. For example, two states that are labelled differently in the semantic model but are epistemically indistinguishable will be mapped to the same state in the syntactic model, indicating that they are, in some sense, redundant.

In our example, when we map $\mathcal{K}$, as defined above, to Ω using ϕ the following phenomenon is observed. The elements $0, 1, \ldots$ of $\mathcal{K}$, being epistemically distinguished, are injectively mapped into Ω. However, Ω also contains additional states that are not in $\phi(\mathcal{K})$. One such additional state is what might be termed a 'state at infinity', ω_∞, which contains formulae expressing that the event that p occurs is common knowledge among players 1 and 2, while player 3 still has no knowledge at all of what mutual knowledge players 1 and 2 have.

Now, it turns out that ω_∞ contains formulae in common with all the states in the image of $\mathcal{K}$ in Ω, namely the formulae stating that player 3 does not know whether players 1 and 2 have k-level mutual knowledge of p for all k. Given the way the knowledge partitions of the players are defined in Ω, this means that $\phi(0), \phi(1), \ldots$ and ω_∞ are all elements within the same common knowledge component of Ω. In other words, ϕ maps $0, 1, \ldots$, which comprise a common knowledge component in $\mathcal{K}$, *non-surjectively* into the common knowledge component containing $\phi(\mathcal{K})$. It is in this sense that something that is 'common knowledge' in $\mathcal{K}$ never the less lacks expression within that structure, namely, the fact that player 3 never can know what mutual knowledge players 1 and 2 have.

The only way to attain surjectivity here is to add a point at infinity explicitly into the semantic model, yielding a model $\mathcal{K}'$ with states $\{0, 1, \dots\} \cup \{\infty\}$. The partitions are, for player 1: $\{\{0\}, \{1, 2\}, \dots \{\infty\}\}$, for player 2: $\{\{0, 1\}, \{2, 3\} \dots \{\infty\}\}$, for player 3: $\{\{0, 1, \dots \infty\}\}$. But this amounts to adding a state to the semantic model in which the implicit statement 'no matter how large n is, at every state n player 3 has no idea what level of mutual knowledge players 1 and 2 have' is stated explicitly.

Non-surjectivity of the mappingfrom a common knowledge component, or *cell*, of a semantic model $\mathcal{M}$ to the corresponding cell of the canonical syntactic structure Ω indicates that some common knowledge in $\mathcal{M}$ is 'excluded'. This motivates seeking characterisations of when surjectivity holds.

Based on the above example, one might conjecture that lack of surjectivity arises if and only if at least one partition element of one of the players contains an infinite number of states (as is the case with player 3). Partition profiles in which each element of each player's partition contains a finite number of states are said to have *finite fanout*.

Finite fanout has recently been prominent in counter-examples and results related to the existence of equilibria in Bayesian games with infinitely many states. In a now classical result, Harsányi (1967) showed that every finite Bayesian game (meaning finite players, finite actions, finite states) has a Bayesian equilibrium, essentially by reducing the question to one of the existence of Nash equilibria in an appropriately defined corresponding game. Existence of equilibria when state spaces are of the cardinality of the continuum, however, is a much more complex matter with which to contend.[1]

Restricting attention to Bayesian games with player partition profiles with finite fanout has enabled some progress in the study of such equilibrium existence (or non-existence) questions. Simon (2003) presented an example of a three-player Bayesian game with a continuum of states of the world and finite fanout that admits no measurable Bayesian equilibria. More recently, Hellman (2013) showed an example of a two-player Bayesian game with finite fanout that lacks even approximate equilibria.

[1] Perhaps the most prominent publication containing sufficient conditions for the existence of Bayesian equilibria is Milgrom and Weber (1985), which makes use of conditions guaranteeing that the payoff functions are continuous with respect to the weak* topology. These conditions, however, are rather restrictive; see Simon (2003) for more details.

This was followed up by Hellman and Levy (2013), which presents necessary and sufficient conditions for the existence of measurable equilibria in Bayesian games with finite fanout.[2]

Finding epistemic characterisations of finite fanout is therefore of interest, alongside characterising surjectivity. Simon (1999) showed that all common knowledge components of semantic structures with finite fanout are also surjective components, thus confirming some connection between the two concepts, and that any surjective common knowledge component of Ω must be at most countable.

We show here, however, that failure of finite fanout does not necessarily lead to non-surjectivity. Our main result is an example of a structure with a cell that is surjective but without finite fanout. Surjectivity and finite fanout, it turns out, are not equivalent concepts.

2. Semantic and Syntactic Models

2.1. **The Semantic Approach.**

One approach for rigorously defining knowledge and common knowledge is by way of semantic models, making use of partition profiles with evaluations (see also (Aumann, 1976) for a similar model).

Let X be a set that intuitively represents every primitive, binary statement that can be made about a particular subject, e.g., 'the share price rose today', 'rain is not falling on the plain' and so forth. Hence, if one knows for each $x \in X$ whether 'the value' of x is true or false, one knows all that there is to know that is relevant to a particular subject.

We also assume the existence of a non-empty set of players J. Although one may consider the case in which either X or J is infinite, in this paper for simplicity we will assume throughout that both X and J are finite. In addition, there is a set S that is the set of *states of the world*, along with an *evaluation* function $\psi : S \to \{0,1\}^X$ (with ψ^x furthermore standing for the function ψ projected to the x coordinate). If one knows ψ and knows the 'true state of the world' $s \in S$, then one knows the 'truth' about all that there is to know, by reading off $\psi(s) \in \{0,1\}^X$.

To model uncertainty, we suppose that players do not know the one true state of the world; for each possible true state they regard several such states as possible. More formally, each player $i \in J$ has a partition $\mathcal{P}^i$ of the set of states. The tuple $(S, \mathcal{P}, J, X, \psi)$, with $\mathcal{P} = (\mathcal{P}^i \mid i \in J)$, is then called a *partition profile with evaluations*.

[2] In fact, those same conditions also characterise the existence of measurable equilibria in Bayesian games with countable fanout.

Common knowledge of an event E by a set of persons intuitively means that for every string of persons $i_1, i_2, \ldots, i_k$ it is the case that i_k knows that i_{k-1} knows that ... i_1 knows that the event E has occurred (Lewis, 1969). To translate this intuition to our formalism we make use of what are termed cells (sometimes also called common knowledge components) of partition profiles.

Given the partition profile $\mathcal{P}$, the meet $\bigwedge \mathcal{P}$ is defined as the finest common coarsening of $\mathcal{P}$. An element of $\bigwedge \mathcal{P}$ is a *cell*. The intuitive notion of common knowledge is then captured formally in the context of partition profiles with evaluations by defining an event A (i.e., a subset of S) to be *known in common* by all the players in J at the state $s \in A$ if the cell containing s is contained in the set A.

2.2. **The Syntactic Approach.**

We can run through these same basic ideas using a syntactic approach. Doing so in fact has the added value of leading to the construction of a canonical partition profile with evaluations.

As before, let X be a set of primitive propositions and let J be a set of players. Use these to construct a set $\mathcal{L}(X, J)$ of formulae (writing simply $\mathcal{L}$, in place of $\mathcal{L}(X, J)$ if there is no ambiguity) as follows:

(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 formulae constructed through application of the above four rules are members of $\mathcal{L}(X, J)$.

As standard, $\neg f$ stands for the negation of f, $f \wedge g$ stands for f and g. $f \vee g$ stands for either f or g (inclusive) and $f \to g$ stands for $\neg f \wedge g$.

If $\mathcal{K} = (S, \mathcal{P}, J, X, \psi)$ is a partition profile with evaluations then define a map $\alpha^{\mathcal{K}}$ from $\mathcal{L}(X, J)$ to 2^S, the set of subsets of S, inductively as follows, using the structure of the formulae:
Case 1 $f = x \in X$**:** $\alpha^{\mathcal{K}}(x) := \{s \in S \mid \psi^x(s) = 1\}$.
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 $f = k_j(g)$**:** $\alpha^{\mathcal{K}}(f) := \{s \mid s \in P \in \mathcal{P}^j \Rightarrow P \subseteq \alpha^{\mathcal{K}}(g)\}$.

There is a very elementary logic defined on the formulas in $\mathcal{L}$ called $S5$.[3] Briefly, the $S5$ logic is defined by two rules of inference (modus ponens and

[3] For a more extensive discussion of $S5$ logic, see Cresswell and Hughes (1968); for the multi-player variation, see Halpern and Moses (1992) and Bachrach et al. (1997).

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 formulae resulting from theorems of the propositional calculus via 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 formulae $\mathcal{A} \subseteq \mathcal{L}(X, J)$ 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 formulae 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.

Define

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

$\Omega(X, J)$ is itself a partition profile with evaluation. For each player $j \in J$ we define the corresponding partition $\mathcal{Q}^j(X, J)$ to be that generated by the inverse images of the function $\beta^j : \Omega(X, J) \to 2^{\mathcal{L}(X,J)}$, namely,

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

In words, for any state $z \in \Omega(X, J)$, the set $\beta^j(z)$ is the set of formulae that player j knows to hold at z. The partition element $\mathcal{Q}^j(X, J)(z)$ of player j at z is then the set of all states in $\Omega(X, J)$ at which each of the formulae in $\beta^j(z)$ hold true. That is, intuitively, $z' \in \mathcal{Q}^j(X, J)(z)$ means that player j 'knows' at state z that state z' is possible because for each $k_j f \in z$, it is the case that $f \in z'$.

Note that 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 $\mathcal{Q}^j$, without the additional specification of X and J, when there is no ambiguity.

If $\mathcal{K} = (S, \mathcal{P}, J, X, \psi)$ we define a map $\phi^{\mathcal{K}} : S \to \Omega(X, J)$ by

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

This is a canonical map, as defined in Fagin, Halpern and Vardi (1991).

3. Finite Fanout

Definition 1. A cell C of a partition profile with evaluations has *finite fanout* if for every player $i \in J$ and partition element $P \in \mathcal{P}^i$ contained in C, the set P has finitely many elements.

Intuitively, finite fanout means that every player at every state believes that the number of alternative states that are possible is finite.

Fagin (1994) defines, for each set of players J, set of primitive propositions X and ordinal number γ, a hierarchically constructed canonical partition profile with evaluations W_γ. The canonical structure W_γ represents all possible truth evaluations, with the ordinal numbers here representing the levels in the construction of these statements. At the bottom, or 0th level, are the elementary truth assignments which can be thought of as a state of nature. The next level describes each player's beliefs about nature, which intuitively corresponds to the partition structure. The $(k+1)$st-order belief of each player is modelled by a set of possibilities, each of which is a description of a state of nature and each player's kth-order belief. Letting ω stands for the first infinite ordinal, Under Fagin's construction, W_ω is exactly the $\Omega(X, J)$ defined above.[4]

Two states s and s' are indistinguishable at level γ if they share the same state of nature and if every player has precisely the same hierarchy of beliefs up to level γ. If there is an α such that α is the first ordinal at which all states are distinguishable from level α and above, then we say that W_α is *non-flabby*[5] and α is called the *distinguishing* ordinal. If there is no such ordinal then the distinguishing ordinal is defined to be the first ordinal α at which all states that are eventually distinguishable are distinguished at W_α. indistinguishable.

There is another minimal ordinal β, possibly larger than the distinguishing ordinal, which is the least ordinal such that W_β can be extended to any W_γ with $\gamma > \beta$ in only one unique way. This ordinal is called the *uniqueness* ordinal.

Fagin (1994) proves that the uniqueness ordinal is a limit ordinal and never greater than the next limit ordinal above the distinguishing ordinal. Fagin (1994) further establishes that the necessary and sufficient condition for a cell of Ω to have the first infinite ordinal ω as its uniqueness ordinal is that the cell has finite fanout.

There is a natural topology that can be associated with every partition profile with evaluations $\mathcal{K} = (S, \mathcal{P}, J, X, \psi)$, as introduced in Samet (1990). This topology, which we call the topology *induced by the formulae* is generated by $\{\alpha^{\mathcal{K}}(f) \mid f \in \mathcal{L}\}$ as its basis of open sets of S. The topology of a

[4] There are alternative canonical constructions corresponding to the ordinal numbers (cf. Heifetz and Samet (1998) and Heifetz and Samet (1999)), but with respect to the first infinite ordinal ω they are identical to Fagin's.

[5] Intuitively, in a flabby structure there are at least two states that are different 'syntactically' (they have different 'names'), but not 'semantically'.

subset A of S is then the relative topology for which the collection of open sets of A is given by $\{A \cap O \mid O$ is an open set of $S\}$.

Even without explicitly mentioning this topology, Fagin (1994) shows that for any player j, any partition element P of the partition $\mathcal{Q}^j$ is a compact subset of Ω. Fagin (1994) then proves that for any state x in Ω, extending the players' beliefs regarding x from W_ω to $W_{\omega+1}$ is defined by the set of dense subsets R_j of the set of $P_j \in \mathcal{Q}^j$ containing x, for all j. It follows that there is a unique extension of a cell of Ω to ordinals higher than ω if and only if for every player $j \in J$ every $P_j \in \mathcal{Q}^j$ in the cell has only one dense subset. This is equivalent to the cell having finite fanout.

4. SURJECTIVITY

Definition 2. A cell C of Ω is *surjective* if all partition profiles with evaluations $\mathcal{K}$ that map to it by $\phi^{\mathcal{K}}$ do so surjectively.

This definition of surjectivity (Definition 2) is taken from Simon (1999).

We present here, as motivation for studying surjectivity the following variation of an example presented in Fagin, Halpern and Vardi (1991), which is itself a variation of Rubinstein's electronic-mail game with an added third player. The example here is slightly different from the one in the introduction, and it is presented in greater detail and rigour.

There are three players. A fact p either holds or does not hold, and players 1 and 3 are informed which one of these is true, while player 2 is not. If p does not hold, nothing happens. If p holds then players 1 and 2 communicate over an unreliable channel. First player 1 tells player 2 that p. If player 2 received this message, he sends an acknowledgement. If player 2 receives the acknowledgement, he acknowledges the acknowledgement and so on. If at any point a message is not received there is no further communication. There is no communication between player 3 and the other players; player 3 does not know how many messages have been passed between player 1 and 2.

For the partition profile $\mathcal{K}$ with evaluations in this example, let J be a three-player set and let X contain a single primitive proposition, labelled p. The set of states of the world S is composed of $\neg p$, the state where p does not hold, and $\{(p, k)\}_{k \geq 0}$, at each of which p holds. Intuitively, at each (p, k), there were k messages received by players 1 and 2, and a $(k + 1)$st message was sent by the recipient of the kth message (or by player 1 if $k = 0$), but not received.

The knowledge of the players is straightforwardly determined at each state, because the number of messages received is all that is needed to unfold all levels of player knowledge. For example, at $(p, 2)$, p holds, player 2 has received player 1's initial message and 1 has received 2's acknowledgement. Hence player 1 considers states $(p, 2)$ and $(p, 3)$ possible whilst player 2 considers $(p, 2)$ and $(p, 1)$ possible. Player 3 considers all the states $\{(p, k)\}_{k \geq 0}$, possible. Similar reasoning holds at all states.

The mapping of $\mathcal{K}$ into the knowledge hierarchies of Ω is similarly straightforward by induction. At $(p, 2)$, player 1 knows p, knows that player 2 knows p and knows that player 2 knows that player 1 knows p. But he does not know that player 2 knows that player 1 knows that player 2 knows p, (because player 2 might think that his acknowledgement has been lost along the way). As for player 2, he knows p, knows that player 1 knows p, but does not know that player 1 knows that player 2 knows p. Player 3 knows p, but does not know what the other players know, other than that player 1 also knows p.

The cell C of Ω containing $\phi^{\mathcal{K}}((p, 2))$ the image of $(p, 2)$ under the canonical mapping, also contains the image of each and every state (p, k). However, this same cell C also contains a 'point at infinity' ω_∞ that is not the image of any state in $\mathcal{K}$. At ω_∞, the event that p occurs is common knowledge amongst players 1 and 2, meaning that ω_∞ contains all possible permutations $k_{i_\ell} k_{i_{\ell-1}} \dots k_{i_1} p$, for all positive integers ℓ, where each $i_x \in \{1, 2\}$. Intuitively, this corresponds to the situation in which infinitely many messages have been passed between player 1 and player 2.

In addition, ω_∞ contains the formulae $k_3 p$, $k_3 k_1 p$, $k_1 k_3 p$, $\neg k_3 k_2 p$ and all possible permutations $\neg k_3 k_{i_\ell} k_{i_{\ell-1}} \dots k_{i_1} p$, for all positive integers ℓ, where each $i_x \in \{1, 2\}$ (other than $\neg k_3 k_1 p$, which must be left out of this list for the sake of consistency). This intuitively means that player 3 'knows' that players 1 and 2 do not have common knowledge of the event that p is true – even if 'infinitely' many messages pass between players 1 and 2, player 3 still is not informed of this.

All of this is, of course, impossible in $\mathcal{K}$ given the way it is constructed, as $\mathcal{K}$ does not include a 'point at infinity'. Ω, however, does contain such a point because ω_∞ is a complete and consistent set of formulae in $\mathcal{L}(X, J)$.

Furthermore, ω_∞ is a point of C. Perhaps the easiest way to see this is to consider the fact that the formulae tracing the levels of player 3's lack of knowledge regarding the mutual knowledge of players 1 and 2, which are listed above as all the formulae $\{\neg k_3 k_{i_\ell} k_{i_{\ell-1}} \dots k_{i_1} p\}$, hold true in all the states $\{\phi^{\mathcal{K}}((p, k))\}_{k \geq 0}$ and in ω_∞. Hence the formulae $\{\neg k_3 k_{i_\ell} k_{i_{\ell-1}} \dots k_{i_1} p\}$

serve to 'witness' the common knowledge that $\{\phi^{\mathcal{K}}((p,k))\}_{k\geq 0}$ and in ω_∞ have as elements in the same common knowledge cell.

More succinctly, there is common knowledge at $\{\phi^{\mathcal{K}}((p,k))\}_{k\geq 0}$ and ω_∞ that player 3 never knows how many messages have been passed between players 1 and 2. This common knowledge also holds within $\mathcal{K}$ but there is no way of expressing this within $\mathcal{K}$. This impossibility of explicitly stating 'what is commonly known' is what is captured by the fact that the canonical mapping of $\mathcal{K}$ to Ω is not surjective.

Note that the example just presented, of a non-surjectivity, lacks finite fanout. This naturally leads to the question: is finite fanout the key factor determining whether or surjectivity holds? Indeed, Simon (1999) proves that any surjective cell of Ω must be countable and that any cell of Ω with finite fanout is surjective. Many examples of non-surjectivity involve lack of finite fanout. This might seem to support the conjecture that surjectivity implies finite fanout. This, however, turns out to be false. In Section 6, we construct an example of a countable and surjective cell that does not have finite fanout.

5. Centredness and Goodness

One might expect at first glance that, contrary to the result here, failure of finite fanout would always imply failure of surjectivity, using the following reasoning. Firstly, taking a look at the example in Section 4, the state ω_∞ is a cluster point of the partition element of player 3, using the topology induced by the formulae, as can be seen by noting that $\alpha^{\mathcal{K}(f)}$ intersects player 3's partition element for all formulae f of ω_∞.

Now, let C be a cell of some syntactic structure that lacks finite fanout. Then by definition there is at least one player j who has an information partition $P \in \mathcal{Q}^j$ that is contained in C such that P contains an infinite number of elements. It follows that there exists a state $z \in P$ that is a cluster point of P. It might then seem that one could construct a partition profile with evaluations composed of one cell, whose states are explicitly $C \setminus z$, that would canonically map non-surjectively into C (i.e., missing z) in a manner analogous to the electronic mail game variation of the previous section. One would then, in effect, be able to show that the basic element that enabled us to show non-surjectivity in that example always holds in all structures lacking finite fanout. It turns out, surprisingly, that this is not the case.

Equivalence of finite fanout and surjectivity does hold when a cell satisfies a property called *centredness*, as shown in Theorem 3b of Simon

(1999). Centredness has several equivalent definitions; the most straightforward definition is that a cell of Ω is *centred* if and only if no other cell of Ω shares the same set of formulae held in common knowledge Simon (1999). (The set of formulas held in common knowledge is always constant throughout any given cell; see Halpern and Moses (1992)). An equivalent formulation of centredness is that the cell is an open set relative to the closure of itself.

The difference between centred and uncentred cells is radical; if a cell is not centred then there are uncountably many other cells sharing the same set of formulae in common knowledge (see Simon (1999)) and as stated above, centredness implies the equivalence of finite fanout and surjectivity.

The results of this paper also make crucial use of another result of Simon (1999), namely the first part of Lemma 5 in that paper, which states that if $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ is a partition profile with evaluations and P is a member of $\mathcal{P}^j$ for some $j \in J$ then $\phi^{\mathcal{K}}(P)$ is a dense subset of F for some $F \in \mathcal{Q}^j$. This fact was used implicitly by Fagin (1994).

Given a partition profile with evaluations $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ and a subset $A \subseteq S$, we define a new partition profile with evaluations $\mathcal{V}^{\mathcal{K}}(A) := (A; J; (\mathcal{P}^j|_A \mid j \in J); X; \psi|_A)$, where for all $j \in J$,

$$\mathcal{P}^j|_A := \{F \cap A \mid F \cap A \neq \emptyset \text{ and } F \in \mathcal{P}^j\}.$$

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 directly relate the goodness property to the main question studied here.

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

Lemma 9 of Simon (1999): If A is a good subset of a cell C and if $A \cap F$ is closed for every $F \in \mathcal{P}^j$ with $A \cap F \neq \emptyset$, then $A = C$.

We need a few more facts about $\Omega(X, J)$, where X and J are non-empty. If $|J| \geq 2$ then $\Omega(X, J)$ is topologically equivalent to a Cantor set (see Fagin, Halpern and Vardi (1991)), which with its usual topology is a metric space.

A Cantor set can also be conceived as $\{0, 1\}^{\mathbb{N}}$, where each finite sequence $a = (a^1, a^2, \ldots, a^n)$ defines a cylinder subset $C(a)$ of $\{0, 1\}^{\mathbb{N}}$ by

$$C(a) := \{x \in \{0, 1\}^{\mathbb{N}} \mid x^k = a^k \; \forall k \leq n\}.$$

Furthermore, all cylinder subsets are themselves topologically equivalent to Cantor sets, and the same holds for finite unions of cylinder sets.

Thirdly, if $|J| \geq 2$ then there exists an uncentred cell of $\Omega(X, J)$ of finite fanout that is dense in $\Omega(X, J)$ (Simon (1999)).

Due to topological formulations of the centered property, to demonstrate that there is a surjective cell without finite fanout requires some topological insight. Central to our results here is the following proposition, which appears as Theorem 9 of Chapter 12 of 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. Then there is a homeomorphism between X and Y that is also a bijection between X' and Y'.*

Following Klein and Thompson (1984), we call a partition $\mathcal{P}$ of a metric space D upper (respectively lower) hemi-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) hemi-continuous correspondence.

Lemma 1. *If $\mathcal{K} = (S; J; (\mathcal{P}^j \mid j \in J); X; \psi)$ is a partition profile with evaluations that is endowed with a topology (not necessarily that induced by the formulas) such that*

(1) *for every $z \in \{0,1\}^X$ the set $\psi^{-1}(z)$ is clopen (closed and open); and*
(2) *for every $j \in \mathbb{N}$ the partition $\mathcal{P}^j$ is lower and upper hemi-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 hypothesis. 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.

Next, assume that, for some $f \in \mathcal{L}(X, J)$, the set $\alpha^{\mathcal{K}}(f)$ is clopen. $\alpha^{\mathcal{K}}(k_j f)$ 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 $\mathcal{P}^j$. □

Lemma 2. *Given finite X and J, for every $j \in J$ the corresponding partition $\mathcal{Q}^j(X, J)$ of $\Omega(X, J)$ is upper and lower hemi-continuous with respect to the topology induced by the formulas.*

Proof: Let $x_1, x_2, \ldots$ be a sequence of points converging to some $x \in P \in \mathcal{Q}^j$, with $x_i \in P_i \in \mathcal{Q}^j$ for every $i = 1, 2, \ldots$.

To prove that $\mathcal{Q}^j$ is upper hemi-continuous, it suffices to show that if $y_1, y_2, \ldots$ is a sequence of points converging to y, with $y_1 \in P_1, y_2 \in$

$P_2, \ldots$ then y is also in P. Let f be any formula such that $k_j f \in y$. Since the y_i converge to y, there is an N such that for every $i \geq N$ it must be the case that $k_j f$ is in both y_i and in x_i. But this means that $k_j f$ is also in x. Exactly the same argument holds for the formula $\neg k_j f$.

To prove that $\mathcal{Q}^j$ is lower hemi-continuous it suffices to show that if $y \in P$ then there is a sequence $y_1 \in P_1, y_2 \in P_2, \ldots$ that converges to y. Since there are only countably many formulas and one can create a new sequence from the diagonal of sequences, getting closer and closer to y, if the claim were not true then there would be some formula f in y and a sufficiently large integer N such that f is not in any member of P_i for all $i \geq N$. This would also imply that $k_j(\neg f)$ is in x_i for all $i \geq N$ and hence that $k_j(\neg f)$ is in x. But this in turn would contradict the assumption that $f \in y \in P$. $\square$.

6. The Example

Let Ω denote $\Omega(X, \{1,2\})$, where X is an arbitrary finite non-empty set. Let C be an uncentred cell of finite fanout that is dense in Ω. We assume that $\pi : \Omega \to \{0,1\}^{\mathbb{N}}$ is a homeomorphism. For every $n \in \mathbb{N}$ define $\pi_n : \Omega \to \{0,1\}^n$ by setting $\pi_n(x)$ by definition to be the unique $a = (a^1, a^2, \ldots, a^n) \in \{0,1\}^n$ such that $\pi(x) = (a_1, \ldots, a_n, \ldots)$. This means that $\pi_n^{-1} \circ \pi_n(x)$ equals $C(\pi_n(x))$, the corresponding cylinder set. For the special case of a equalling the empty sequence in $\{0,1\}^0$, define $\pi_0(x) := a$ and $\pi_0^{-1} \circ \pi_0(x) = \Omega$ for all $x \in \Omega$.

Let z be any member of C. 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 $A_{1,1} := \Omega \backslash \left(\pi_2^{-1} \circ \pi_2(z_1) \cup \pi_2^{-1} \circ \pi_2(z)\right)$. For $1 \leq k < i$ let

$$A_{i,k} := \pi_{2i-2}^{-1} \circ \pi_{2i-2}(z_k) \backslash \, \pi_{2i}^{-1} \circ \pi_{2i}(z_k)$$

and let

$$A_{i,i} := \pi_{2i-2}^{-1} \circ \pi_{2i-2}(z) \backslash \left(\pi_{2i}^{-1} \circ \pi_{2i}(z_i) \cup \pi_{2i}^{-1} \circ \pi_{2i}(z)\right).$$

Since, for every $a \in \{0,1\}^{2i}$ there are four members b of $\{0,1\}^{2i+2}$ such that $a = \pi_{2i} \circ \pi_{2i+2}^{-1}(b)$, 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 that for every $i \geq 1$ 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 $1 \leq k \leq i$.

Note 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 $\Omega\setminus \{z, z_1, z_2, \dots\}$. Let $\mathcal{P} := (\cup_{i=1}^{\infty} \mathcal{P}^i) \cup \{z, z_1, z_2, \dots\}$, which forms a partition of C. It is straightforward to check that $\mathcal{P}$ is upper and lower semi-continuous.

Finally, let $\mathcal{A}$ denote the three-player partition profile with evaluations

$$(C; \{1, 2, 3\}; \mathcal{Q}^1|_C, \mathcal{Q}^2|_C, \mathcal{P}; X, \psi|_C),$$

with the partition $\mathcal{P}$ being the partition of the third player.

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

Proof: By Lemma 1, $\phi^{\mathcal{A}} : C \to \Omega(X, \{1, 2, 3\})$ is continuous. Since every member of $\mathcal{Q}^1|_C$, $\mathcal{Q}^2|_C$, and $\mathcal{P}$ is compact, their images in $\Omega(X, \{1, 2, 3\})$ under the mapping 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 of which are finite partition elements of $\Omega(X, \{1, 2, 3\})$. It follows that 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 $f \in \mathcal{L}(X, \{1, 2\})$, the set $\alpha^{\Omega(X,\{1,2\})}(f)$ gets mapped to $\alpha^{\Omega(X,\{1,2,3\})}(f)$, the mapping $\phi^{\mathcal{A}}$ is an injective and an open mapping (meaning that open sets are mapped to open sets). It follows that the mapping $\phi^{\mathcal{A}}$ is also a homeomorphism of C to $\phi^{\mathcal{A}}(C)$.

Finally, all this implies that the image of the one infinite set in $\mathcal{P}$ is also an infinite set in the cell $\phi^{\mathcal{A}}(C)$, which implies that this cell of $\Omega(X, \{1, 2, 3\})$ does not have finite fanout. □

7. ACKNOWLEDGEMENTS

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

REFERENCES

Aumann, R. J. (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.

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

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

Harsányi, J. C. (1967), Games with Incomplete Information Played by Bayesian Players, *Management Science*, 14, 159–182.

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* 38, pp. 157- 170.

Hellman, Z. (2013), A Game with No Approximate Bayesian Equilibria, *working paper*.

Hellman, Z. and Levy, Y. (2013), Bayesian Games with Continuum State Spaces, *working paper*.

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

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

Milgrom, P. and R. Weber (1985), Distributional Strategies for Games with Incomplete Information, *Mathematics of Operations Research*, 11, 627–631.

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

Rubinstein, A. (1989) The Electronic Mail Game: A Game with Almost Common Knowledge, *American Economic Review* 79, pp. 385-391.

Samet, D. (1990), Ignoring Ignorance and Agreeing to Disagree, *Journal of Economic Theory* 52(1), pp. 190–207.

Simon, R. (2003), Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria, *Israel Journal of Mathematics* 138, pp. 73-92.

Simon, R. (1999), The Difference between Common Knowledge of Formulas and Sets, *International Journal of Game Theory* 28 (3), pp. 367–384.