INFORMATION AND META INFORMATION

Itzhak Gilboa
Department of Managerial Economics
and Decision Sciences
Kellogg Graduate School of Management
Northwestern University
Evanston, Illinois 60208

ABSTRACT

A model of information is presented, in which statements such as "the information sets are common knowledge" may be formally stated and proved. The model can also abe extended to include the statement: "this model is common knowledge" in a well-defined manner, using the fact that when an event A is common knowledge, it is common knowledge that A is common knowledge. Finally, the model may also be used to define a "natural" topology on information.

1. Introduction

In 1976, Aumann published his seminal paper, "Agreeing to Disagree," introducing the concept of common knowledge in game theory. In his model, every player t has an information partition Π_t of the set of all possible states of the world Ω such that when $\omega \in \Omega$ obtains, player t is informed of the member of Π_t which contains ω. This model, which may be derived from more primitive assumptions--as is shown in Aumann ("An Axiomatization of Knowledge") and Bacharach (1985)--seems to be sufficient for describing the structure of information in a game and, indeed, for all practical purposes encountered so far.

However, game theory models quite frequently call for informal assumptions regarding the meta-information, that is, the information about the information. A prevalent assumption of that sort is that the information structure (i.e., the partitions of Ω) is itself common knowledge. Such an assumption is supposed to be expressible in the partition model. Indeed, in Aumann (1976), we find (p. 1237):

> Worthy of note is the implicit assumption that the information partitions P_1 and P_2 are themselves common knowledge. Actually, this constitutes no loss of generality. Included in the full description of a state ω the world is the manner in which information is imparted to the two persons. This implies that the information sets $P_1(\omega)$ and $P_2(\omega)$ are indeed defined unambiguously as functions of ω, and that these functions are known to both players.

It should be stressed that this assumption is implicit in Aumann's model. Moreover, it is not trivial that it may be formalized in a well-defined manner. A straightforward formalization will have to include subsets of Ω (the information sets) in the definition of each member of Ω, thus rendering this definition self-referring.

This implicit self-reference has troubled many game theorists in the last decade. At least four independent trials to cope with this problem were made in the last few years: Tan and Werlang (1985), Gilboa (1986), Kaneko (1987), and Samet (1987). We will now give a (very) brief survey of these papers, while the content of Gilboa (1986) will be given in the following sections.

Tan and Werlang (1985) begin with a basic uncertainty space Ω and construct upon it the spaces of beliefs, beliefs regarding beliefs, and so on in an infinite-recursion model. They redefine common knowledge and show that for a given Aumann model there exists an infinite-recursion model such that the two notions of common knowledge coincide. They note that their results, together with those of Brandenburger and Dekel (1985), show that the two approaches are equivalent. One of the merits of their model is that they can formally state that "the information partitions are common knowledge."

However, the partitions they refer to are those of the "basic" uncertainty space, and not of the "universal" space obtained from the former by constructing the hierarchies of belief upon it. In other words, they do not formalize Aumann's implicit assumption but rather construct a super-model to deal with the meta-information. Needless to say, they do not try to formally construct a model which will itself be common knowledge.

The model presented in Gilboa (1986) is supposed to be an exact formalization of Aumann (1976). Its first and most primitive version shows that the problem we began with can be quite easily solved: one can write down a well-defined model in which every state of the world specifies "the manner in which information is imparted" to the players--in the sense of subsets of the same uncertainty space Ω.

However, this is not all we intended to do: although the basic self-reference problem has been solved, the partitions themselves or, to be precise, the fact that there are such partitions, are not yet common knowledge, let alone the model itself. The solution suggested by Gilboa (1986), though not carried out formally, is the introduction of logic into the model, that is to say, explicitly writing down the assumption that all the axioms of logic <u>and the model we have so far described</u>, are common knowledge. Using the trivial result that if an event A is common knowledge, then it is common knowledge that A is common knowledge--one deduces (rather than assumes) that this last assumption is itself common knowledge, and hence the model itself (including this assumption) is common knowledge.

Gilboa (1986) also suggests extensions of the model and a few other results which are included in the sequel.

Kaneko (1987) also introduces logic into the model (but, as opposed to Gilboa (1976), in a detailed formal way). He formally discusses the notion of "sharing the epistemic world" defined by a proposition, thus providing the tools to deal with a model which is itself common knowledge, although he does not explicitly assume that his model is such.

The main difference between Kaneko (1987) and Gilboa (1986) is, at least to the best judgment of the author, that Kaneko distinguishes between factual and structural common knowledge: factual common knowledge is the notion defined by Aumann and its objects are events--subsets of the basic uncertainty space, Ω. Structural common knowledge is only defined in a super-model, the objects of which are statements (or mathematical propositions), including statements about the super-model itself. Like Tan and Welang (1985), Kaneko (1987) does not provide a formalization of Aumann's intuitive framework but rather suggests that the Aumann model (of states of the world, partitions, etc.) be restricted to information about the game ("the basic uncertainty"), and the meta-information will be dealt with in a super-model. These models do not explicitly solve the problem we began with (can Aumann's intuition be justified in his original and succinct model?). The positive answer to this question is given in Gilboa (1986), where information and meta-information are dealt with in exactly the same way and there is complete equivalence between statements (including meta-informational statements) and events which are subsets of the set of the states of the world.

Finally, Samet (1987) also contains a formalization of information and meta-information. His model is similar to Gilboa (1986) in the sense of retaining the relationship between meta-information and the uncertainty space. He, too, only suggests the introduction of logic into the model. However, the main point in Samet's paper regards the substance rather than the formulation of game theory axioms: he shows that one can dispose of the philosophically controversial axiom that every player knows what he does not know (namely, that if player t does not know whether proposition A is true, t knows that he does not know it), and still obtain the "common posterior" result of Aumann (1976).

2. Description of the Model and the Results

The model presented in Gilboa (1986) deals with meaningless characters (symbols), finite strings of these characters and (infinite) sets of strings. Only the assumptions imposed upon these strings and sets of strings allow us to interpret them as representing states of the world, events, game-theoretic axioms, and so forth.

Since information and meta-information are dealt with in the very same way in our model, we will be able to formally state and prove propositions such as "If an event A is common knowledge, then it is common knowledge that A is common knowledge." Furthermore, "event" may also be the fact that A is common knowledge, so that one may also prove that, if A is common knowledge, then $[\text{it is common knowledge that}]^n A$ for any $n \geq 1$.

We will also prove in our framework that the information sets have to form partitions of Ω. The intuitive argument of the proof is very simple and is identical to the one used in Aumann ("An Axiomatization of Knowledge").

Next we turn to the introduction of logic into the model. Intuitively, what that means is to allow the players of the model to *think*, rather than just *know* facts. For instance, in the basic model the players know what the information sets are, and these form partitions, but the players cannot "understand" this last fact. In the extended model it makes sense to state that it is common knowledge that the information sets form partitions. In fact, in the extended model, the (extended) model itself is common knowledge; that is to say, all our assumptions are common knowledge, and hence everything we may state and prove is also known to every player in the model.

Section 5 also suggests extensions of the model by the introduction of topology and of probabilistic statements. However, we do not present any additional results in this paper, but merely discuss the concepts.

This introduction cannot be concluded without an apology. Unfortunately, all results presented in this paper are trivial. It seems that in this respect too, one may quote Aumann (1976): "We publish this observation with some diffidence, since once one has the appropriate framework, it is mathematically trivial."

On the other hand, in spite (and maybe because) of their triviality, the proofs may be confusing. Therefore they are given in detail, and the reader who understands the main idea is advised to skip them.

3. The Basic Model

Let T denote a nonempty set of players. For each player $t \in T$ let there be given two characters, k_t and $\bar{k}_t$, which will be interpreted as "t knows that. . ." and "t does not know whether. . .," respectively. In addition, we will need another character, which will be 'c', to denote common knowledge. Denote $K = \{k_t | t \in T\} \cup \{\bar{k}_t | t \in T\} \cup \{'c'\}$.

Let Ω be a nonempty set of characters, each of which denotes a state of the world. Similarly, let $\mathcal{E}$ be a set of characters, each of which denotes an event.

A statement is one of the following:

(i) A sequence ω '$\in$' A where $\omega \in \Omega$ and $A \in \mathcal{E}$. (That is, the concatenation of three characters, the first of which belongs to Ω, the second is the constant character '$\in$', and the third is a member of $\mathcal{E}$. In the sequel we will drop the apostrophes when they are more likely to cause confusion than to prevent it.) The meaning of such a statement is that the state of the world, ω, belongs to the event A.

(ii) A character $A \in \mathcal{E}$. The meaning of the statement "A" is that the event A occurs.

(iii) A sequence kS where $k \in K$ and S is a statement. The interpretation of such a statement depends on k, as explained earlier.

(iv) A sequence ωS where $\omega \in \Omega$ and S is a statement. Such a statement should be read as "S is true if ω obtains."

The set of all statements will be denoted by Σ. Subsets of Σ will be called languages.

We now define a function $E: \mathcal{E} \times 2^{\Sigma} \to 2^{\Omega}$ as follows:

$$E(A,L) = \{\omega \in \Omega | \text{"}\omega \in A\text{"} \in L\} \text{ for } A \in \mathcal{E} \text{ and } L \subset \Sigma.$$

That is, $E(A,L)$ is the set of states of the world which are, according to L, members of the event A.

We define $F = K \cup \Omega$, and denote by F^* the set of all finite-length strings of characters in F (including the empty string).

For a language L and $\underline{f} \in F^*$ we define a language $\underline{f}(L)$ by

$$\underline{f}(L) = \{S \in \Sigma | \underline{f}S \in L\}$$

where $\underline{f}S$ is the statement generated by the concatenation of $\underline{f}$ and S. For instance, $k_1(L)$ will be the language describing what player 1 knows, according to L. Similarly, $\omega k_1(L)$ will denote the language describing what player 1 would know, if the state of the world ω obtained.

We now define a function $M: 2^{\Sigma} \to 2^{\Omega}$ by

$$M(L) = \{\omega \in \Omega | L \subset \omega(L)\}$$

for any $L \subset \Sigma$. That is, $M(L)$ is the set of states of the world which are compatible with L. Note that the definitions of the states of the world are also dependent on L.

We will now list several requirements on languages, designed to support our interpretation of the various symbols.

A language L is said to satisfy the first-order consistency requirements if it is true that:

(1) For any $B \subset \Omega$ there exists a unique $A \in \mathcal{E}$ such that $E(A,L) = B$. This character A will be denoted $E^{-1}(B,L)$. (Thus, E^{-1}: $2^{\Omega} \times 2^{\Sigma} \to \mathcal{E}$.)

(2) For any $A \in \mathcal{E}$, "A" $\in L$ iff

$$E(A,L) \supset M(L).$$

(3) $M(L)$ is not empty.

(4) For any $\omega \in \Omega$ it is true that:

(a) For any $B \subset \Omega$, either "$\omega\ E^{-1}(B,L)$" $\in L$ or "$\omega\ E^{-1}(B^c,L)$" $\in L$ or "$\omega\ E^{-1}(B^c,L)$" $\in L$;

(b) For any $B \subset \Omega$ and any $t \in T$, exactly one of the following is true:

"$\omega\ k_t\ E^{-1}(B,L)$" $\in L$

"$\omega\ k_t\ E^{-1}(B^c,L)$" $\in L$

or

"$\omega\ \bar{k}_t\ E^{-1}(B,L)$" $\in L$

(5) For any $t \in T$ and $S \in \Sigma$,

(a) If $k_tS \in L$, then $S \in L$;

(b) If $k_tS \in L$, then $k_tk_tS \in L$;

(c) If $\bar{k}_tS \in L$, then $k_t\bar{k}_tS \in L$.

(6) For any $s \in \Sigma$, $cS \in L$ if and only if $\underline{k}S \in L$ for any finite-length string

$$\underline{k} \in \cup_{n\geq 1}\ (k_t | t \in T)^n.$$

Requirement (1) is a grammatical requirement. It assures us that any set of states of the world will have a letter $\mathcal{E}$ to denote it, and that this letter will be unique. The second requirement means that those events which are supposed to occur at L (i.e., those A's in $\mathcal{E}$ for which "A" $\in$ L) are exactly those which contain M(L). Recall that M(L) is the minimal event known to occur at L, since it consists of the states of the world which are compatible with L. The third requirement assures us that there are such states of the world. The fourth requirement is designed to capture Savage's (1954) notion of "a state of the world resolving all uncertainty." It first part states that at each ω, for each event B, either B or its complement should occur. Its second part states that, at each ω and for any player t, there are three possibilities, exactly one of which has to realize: (a) t knows B; (b) t knows B^c; (c) t does not know whether B or B^c. Requirement (5) says the following: (a) if somebody (t) knows a fact, then this fact is true; (b) if t knows a fact, then t knows he knows it; (c) if t does not know whether a certain fact is true or not, then t knows that he does not know it. The last requirement (6) is the definition of common knowledge.

We now proceed to define the second order consistency requirements on a language L:

(7) For any string $\underline{f} \in F^*$, $\underline{f}(L)$ satisfies the first-order consistency requirements.

(8) For any string $\underline{f} \in F^*$, any $\omega \in \Omega$ and any $A \in \mathcal{E}$,

$$\omega(\underline{f}(L)) = \omega(L)$$

and

$$E(A,\underline{f}(L)) = E(A,L).$$

Requirement (7) means that all languages generated by L will also satisfy (1)-(6). (Hence, they will also satisfy (7).) Requirement (8) is again a grammatical one, and ensures that the names of all states of the world and all events are the same for all languages generated by L.

If L satisfies (7)-(8), it is called an information set. Note that if L is an information set, $\underline{f}(L)$ is also an information for any string $\underline{f} \in F^*$.

We will now turn to some trivial results which will demonstrate the way in which meta-information is formalized in this model. Our first result says that, if a certain fact is common knowledge, then it is common knowledge that it is common knowledge.

1. Proposition: Suppose L is an information set, and let $cS \in L$ for some $S \in \Sigma$. Then $ccS \in L$.

Proof: Suppose $\underline{k}_1, \underline{k}_2 \in \cup_{n\geq 1} \{k_t | t \in T\}^n$. Since $cS \in L$, we have $\underline{k}_1\underline{k}_2 S \in L$. This being true for any $\underline{k}_2$, it is true that $\underline{k}_1 cS \in L$ for all $\underline{k}_1$. But this is equivalent to $ccS \in L$.//

Before continuing, we have to simplify our notations. In view of requirement (8), whenever the information set L under discussion is fixed, it may be omitted from the function E and E^{-1}. Furthermore, these functions may be dispensed with altogether: since they translate elements of $\mathcal{E}$ into elements of 2^{Ω} and vice versa, no ambiguity may result from their omission. We will, therefore, use elements of 2^{Ω} to denote elements of $\mathcal{E}$ as well. For instance, if $B \subset \Omega$ then "B" should be read as "$E^{-1}(B,L)$".

Our next result is well-known. It states that each player has a partition of Ω, such that if $\omega \in \Omega$ obtained, the player would know the elements of the partition which contains ω. The intuitive explanation of this result is that if a player knows what he knows and knows what he does not know, at each ω the set of states of the world he considers as possible is exactly those states at which he would know what he indeed knows. We remind the reader that this result is obtained in Aumann ("An Axiomatization of Knowledge"), though in a different framework.

In the sequel, L will be assumed to be an information set unless otherwise stated. We begin with a few lemmas:

2. Lemma: For each $\omega \in \Omega$, $M(\omega(L)) = \{\omega\}$.

Proof: By the definition of M and requirement (8), $\omega \in M(\omega(L))$. However, if for some $\omega' \in \Omega$, $\omega' \in M(\omega(L))$, then $\omega(L) \subset \omega'(L)$. Take $B = \{\omega\}$. If "B^c" $\in \omega(L)$, requirement (2) is contradicted. Hence, by requirement (4), "B" $\in \omega(L)$, whence "B" $\in \omega'(L)$. But this implies $\omega' \in M(\omega'(L)) \subset B$, that is, $\omega' = \omega$.//

For the next few lemmas we will fix a player $t \in T$, and write "k", "$\bar{k}$" instead of "k_t", "$\bar{k}_t$", respectively.

3. Lemma: For each $\omega \in \Omega$, $\omega \in M(k(L))$ iff $k(L) \subset \omega(L)$.

Proof: By definition of M, $\omega \in M(k(L))$ iff $k(L) \subset \omega(k(L))$. But requirement (8) implies $\omega(k(L)) = \omega(L)$.//

4. Lemma: "$kM(k(L))$" $\in$ L.

Proof: By requirement (2), "M(L)" $\in$ L for any information L and, in particular, "$M(k(L)) \in k(L)$. this is equivalent to "$kM(k(L))$" $\in L$.//

5. Lemma: For each $\omega \in \Omega$, $\omega \in M(k(L))$ iff "$kM(k(L))$" $\in \omega(L)$.

Proof: First assume $\omega \in M(k(L))$. By Lemma 4, "$kM(k(L))$" $\in L$, whence, using requirement (5), "$kkm(k(L))$" $\in L$, or "$kM(k(L))$" $\in k(L)$. As Lemma 3 implies, $k(L) \subset \omega(L)$, the first part is proved. Now assume that "$kM(k(L))$" $\in \omega(L)$. By requirement (5), "$M(k(L))$" $\in \omega(L)$, whence, by requirement (2), $M(k(L)) \supset M(\omega(L))$. However, by Lemma 2, $M(\omega(L)) = \{\omega\}$, and $\omega \in M(k(L))$ has been proved.//

6. Proposition: For any ω, $\omega' \in \Omega$, $\omega' \in M(k(\omega(L)))$ iff $M(k(\omega(L))) = M(k(\omega'(L)))$.

Proof: First assume that $\omega' \in M(k(\omega(L)))$. By the definition of M, $k(\omega(L)) \subset \omega'(L)$. However, by Lemma 4, "$kM(k(\omega(L)))$" $\in k(\omega(L))$, whence "$kM(k(\omega(L)))$" $\in \omega(L)$ or "$M(k(\omega(L)))$" $\in k(\omega'(L))$. By requirement (2) one obtains $M(k(\omega(L))) \supset M(k(\omega'(L)))$. To see that the converse inclusion also has to hold, assume the converse, i.e., $M(k(\omega'(L))) \subsetneq M(k(\omega(L)))$. By requirement (4), exactly one of the following three possibilities is true: (i) "$kM(k(\omega'(L)))$" $\in \omega(L)$; (ii) "$kM(k(\omega'(L)))^c$" $\in \omega(L)$; (iii) "$\bar{k}M(k(\omega'(L)))$" $\in \omega(L)$. If (i) were to hold, then "$M(k(\omega'(L)))$" $\in k(\omega(L))$, which implies $M(k(\omega'(L))) \supset M(k(\omega(L)))$, contrary to our assumption. On the other hand, (ii) would imply "$M(k(\omega'(L)))^c$" $\in k(\omega(L))$, yielding $M(k(\omega'(L)))^c \supset M(k(\omega(L)))$, whence $M(k(\omega'(L)))^c \supset M(k(\omega'(L)))$, contradicting the nonemptiness of $M(k(\omega'(L)))$. Hence we are left with (iii). By requirement (5) we obtain "$k\bar{k}M(k(\omega'(L)))$" $\in \omega(L)$, or "$\bar{k}M(k(\omega'(L)))$" $\in k(\omega(L))$. However, Lemma 3 implies that $\omega' \in M(k(\omega(L)))$ is equivalent to $k(\omega(L)) \subset \omega'(L)$. The former being true, we obtain "$\bar{k}M(k(\omega'(L)))$" $\in \omega'(L)$. On the other hand, Lemma 4 yields "$kM(k(\omega'(L)))$" $\in \omega'(L)$, a contradiction to requirement (4). This completes the first part of the proof. Now assume $M(k(\omega'(L)))= M(k(\omega(L)))$. We note that requirement (5) implies $k(\omega'(L)) \subset \omega'(L)$. By the definition of M and requirement (8), this is equivalent to $\omega' \in M(k(\omega'(L)))$. Hence, $\omega' \in M(k(\omega(L)))$ has been proved, and that concludes the second part of the proposition.//

We note that proposition 6 is equivalent to the following: "The relation $R_t \subset \Omega \times \Omega$ defined by: "$\omega R_t \omega'$ if, when ω obtains, the player t considers ω' as possible"--is an equivalence relation." The equivalence classes of R_t form the information partition for player t.

We now turn to a few remarks regarding the cardinality of the set of all information sets. We first note that for a given set Ω, $\mathcal{E}$ is essentially unique. Therefore, the set of all information sets is well defined by the set of player T and the set of states of the world Ω, and it will be denoted by $I(T,\Omega)$. In general it will be true that $I(T,\Omega)$ is larger than Ω. For instance, it is easy to see that

7. Remark: If $|T| \geq 3$ and $|\Omega| \geq 2$, then

$$|I(T,\Omega)| \geq \aleph.$$

Proof: We will show that for each subset M of the natural numbers there exists a distinct information set $L_M \in I(T,\Omega)$. Let there be given M. Assume $\{1,2,3\} \subset T$ and $A = \{\omega_0\} \subset \Omega$. For each $m \geq 1$ let S^m denote the statement "$k_1(k_2k_3)^m k_1 A$" (where "$(k_2k_3)^m$" denotes the m-fold concatenation of "k_2k_3" with itself). Now let L^M be an information set for which

$$\{m \geq 1 | S^m \in L^M\} = M.$$

As $L^M \neq L^{M'}$ for $M \neq M'$, our remark is proved.//

Note: If $|T| = 2$, the previous remark is no longer true. In that case $I(T,\Omega)$ is denumerable for any finite Ω.

Moreover, in general it is true that:

8. Remark: $|I(T,\Omega)| > |\Omega|$.

Proof: Trivial since for any $A \subset \Omega$ there exists an information set L in which, for a given $t \in T$, $M(k_t(L)) = A$.//

This remark may be interpreted as an impossibility result: the set of all information sets $I(T,\Omega)$ is, in a way, the set of all conceivable states of the world. (Technically, not every information set has to "resolve all uncertainty." Hence, not every information set may be a description of a state of the world. It is easy to see, however, that the previous remark is valid even when one restricts one's attention to those information sets which, indeed, "resolve all uncertainty.") The meaning of this result is, therefore, that the states of the world in a model of information (of the type discussed here) cannot exhaust all conceivable states of the world.

4. The Introduction of Logic

The model presented in Section 3 is rich enough to formalize phrases such as "it is common knowledge that A is common knowledge" or to prove that the information is structured as in the classical partitions model. However, it is not rich enough to formalize phrases such as "it is common knowledge that the information is structured as in the classical partitions model." In other words, the fact that Proposition 6 proven above is common knowledge cannot be expressed in the model. This is quite obvious since our model does not contain any logical symbols that may be used to describe mathematical statements. The point we would like to stress here is that these symbols are (almost) all that is needed to extend the model to include mathematical statements.

Since this paper already has a low results/definitions ratio, we do not intend to define the extended model formally. We will, therefore, only describe the changes that have to be made in the model's structure and assumptions:

1) The alphabet of the model has to include the standard logical symbols ($\neg$, $\vee$, $\wedge$, (,) $\forall$, $\exists$, $=>$, ...).
2) Any information set has to satisfy additional requirements,

which will guarantee that these symbols have their common meaning. For instance, if $S \in L$, then it is false that $\neg S \in L$, etc.).

3) Requirements (1)-(8), defining information sets, have to be translated into formal mathematical language, and then it should be assumed that for each information set these requirements are common knowledge.

Assuming such a model, we may write:

Informal Proposition: Propositions 1 and 6 proven above are common knowledge.

Informal Proof: All we have to do is repeat the proofs of Section 2, stating "it is common knowledge that. . ." before each step.

In fact, this informal proposition may be interpreted as stating that under this model's assumptions, the model itself is common knowledge, and so are all the results, including the one you are now reading. What makes this seemingly meaningless statement well defined is Proposition 1: it suffices to assume that the assumptions of the model are common knowledge; this very fact (i.e., the last assumption) will also have to be common knowledge by Proposition 1. Thus it may actually refer to itself without doing so formally.

Remark 1: In this framework, the Savagean notion of state-of-the-world "resolving all uncertainty" will mean that for each $\omega \in \Omega$, $\omega(L)$ is complete and consistent. Since a consistent language may be extended to a complete and consistent one by the axiom of choice (or even without it, in case Ω and T are finite), the existence of information depends, basically, upon that of consistent languages satisfying all other requirements. Of course, we cannot prove that such consistent languages exist, and such a proof is impossible by Goedel's theorems. However, we may use (again) the recursive structure of the model: a reader who has not doubted the consistency of the model presented in Section 2 should not doubt the existence of information sets.

Remark 2: In the model described above, the players are allowed to know any fact we know, including the propositions we have proved. However, they are not allowed to prove these propositions themselves. If one really likes to let the player "think," as we have promised to do in the beginning, one should assume formal logic to be common knowledge, and then one can show that all proofs may be carried out within the model.

5. Other Possible Extensions

5.1 Topology: There is little doubt that information plays a major role in reality. Information about the information does not seem to be as important as the information itself, but it cannot be dismissed as irrelevant. It is, however, quite rare that the nature of the game played in reality will depend heavily upon what player 1 knows about what player 2 knows about what player 1 knows, etc. It seems that even the wittiest spy will get confused after a relatively small number of iterative concatenations of $\{k_t | t \in T\}$. Therefore, it is preferred that any solution concept applied to the game will not be too sensitive with respect to the highly complicated statements in the information of the game.

This "bounded rationality" type of argument leads us to define a topology on the set of information sets, with respect to which one may require a solution concept to be continuous. The topology we suggest is defined as follows: for each $n \geq 1$, let Σ^n denote the statements in Σ of length n. (The exact definition of "length" is immaterial. For the sake of concreteness we will suppose that it is the number of characters contained in the statement.) Consider the topology generated by the sub-base

$$\{\{L \in I(T,\Omega) | L \cap \Sigma^n = L_0 \cap \Sigma^n\} | L_0 \in I(T,\Omega),\ n \geq 1\}.$$

In this topology, $L_\alpha \to L_0$ iff for each $n \geq 1$ there exists β such that for all $\alpha \geq \beta$, L_α and L_0 are identical as far as n-long statements are concerned. Note that the space of all information sets is a metric space.

In case Ω and T are finite, Σ is denumerable, and for any enumeration of it, $E: N \to \Sigma$, we may map 2^{Σ} into [0,1] by the function:

$$x(E,L) = \Sigma_{i=1}^{\infty} 2 \cdot 1_{\{E(i) \in L\}} \cdot 3^{-i}.$$

We conclude this sub-section with the following trivial remark which may further motivate our choice of topology on $I(T,\Omega)$:

Remark: If $|T|$, $|\Omega| < \infty$, L_n ($n \geq 1$), $L \in 2^{\Sigma}$, then the following are equivalent:

(i) $L_n \to L$;

(ii) $\cap_{n\geq 1} \cup_{k\geq n} L_k = \cup_{n\geq 1} \cap_{k\geq n} L_k = L$;

(iii) For some enumeration $E, x(E,L_n) \to x(E,L)$;

(iv) For any enumeration $E, x(E,L_n) \to x(E,L)$.

5.2 Probability. The last extension of the model to be discussed here is the introduction of probabilistic statements. "Knowledge" may be considered as "belief with probability 1," while the model may include beliefs with other degrees of probability as well. Such an extension may serve as a formalization of Harsanyi's (1967-68) model of beliefs, but we will not expatiate on it since it does not seem to formalize any notion not included in the model of Mertens-Zamir (1985).

Acknowledgments: I wish to thank Professor David Schmeidler for his guidance. I am also grateful to Professors Robert Aumann, Eddie Dekel, Mamuro Kaneko, Dov Samet, and Tommy Tan for comments and references.

References

Aumann, R. J., "An Axiomatization of Knowledge," unpublished paper.

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

Bacharach, M., "Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge," Journal of Economic Theory, 37 (1985), 167-190.

Brandenburger, A. and E. Dekel, "Hierarchies of Belief and Common Knowledge," research paper, Graduate School of Business, Stanford University (1985).

Gilboa, I, "Information and Meta-Information," Working Paper 30-86, Foerder Institute for Economic Research, Tel Aviv University (1986).

Harsanyi, J., "Games of Incomplete Information Played by Bayesian Players," Parts I-III, Management Science, 14 (1967-68), 159=182, 320-336, 486-502.

Kaneko, M., "Structural Common Knowledge and Factual Common Knowledge," RUEF Working Paper No. 87-27, Hitotsubashi University (1987).

Mertens, J. F. and S. Zamir, "Formalization of Harsanyi's Notion of 'Type' and 'Consistency' in Games with Incomplete Information," International Journal of Game Theory, 14 (1), 1985, 1-29.

Samet, D., "Ignoring Ignorance and Agreeing to Disagree," mimeo (1987).

Savage, L. J., The Foundations of Statistics, New York: John Wiley and Sons (1954).

Tan, T. C. C. and S. R. C. Werlang, "On Aumann's Notion of Common Knowledge--An Alternative Approach," Working Paper No. 85-26, Working Paper Series in Economics and Econometrics, University of Chicago, 1985 (revised 1986).