JOURNAL OF ECONOMIC THEORY **70**, 249–256 (1996)
ARTICLE NO. 0084

"Knowing Whether," "Knowing That," and The Cardinality of State Spaces[1]

SERGIU HART

Department of Mathematics, Department of Economics, and the Center for Rational and Interactive Decision Theory, The Hebrew University of Jerusalem, Jerusalem, Israel 91905

AVIAD HEIFETZ

CORE, Université Catholique de Louvain and the School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel 69978

AND

DOV SAMET

Faculty of Management, Tel Aviv University, Tel Aviv, Israel 69978

Received January 14, 1994; revised June 14, 1995

We introduce a new knowledge operator called "*knowing whether*." Knowing *whether* an event occurred means either knowing *that* it occurred or knowing *that* it did not occur. We demonstrate the following difference between "knowing whether" and "knowing that." In a multiple agent model, a sequence of events generated by successively applying "knowing that" operators, or their negations, may be contradictory. But when "knowing whether" operators are used instead, the sequence is never contradictory. Using this property of the "knowing whether" operator we construct a multiple agent model with a continuum of knowledge states. This simplifies such a construction due to Aumann [1]. *Journal of Economic Literature* Classification Number: D80. © 1996 Academic Press, Inc.

1. INTRODUCTION

How many different states of knowledge are there with regard to some given fact, when several agents are involved? This question has some bearing

[1] We thank Robert J. Aumann for introducing to us the problem of the cardinality of the state space and for helpful discussions and comments.

0022-0531/96 $18.00

on the problem of naming states of knowledge. If there are only finitely or denumerably many states then they can have finite names, in a language with a finite or countable alphabet. But if the cardinality of the set of these states is that of the continuum, then it is impossible to name them all, in much the same way that we cannot name each point in our physical space. In such a case, the idea that agents know the "epistemic model" that describes their knowledge should be interpreted very cautiously. If the model is rich enough, then "know" in this context does not mean that agents can talk (in ordinary language) about every state.

The answer to our question is that there is indeed a continuum of different knowledge states, and therefore we cannot name them all. This answer is by no means novel. Aumann [1], in a set of widely circulating but unpublished lecture notes, constructed, in a rather non-trivial way, a continuum of different states of knowledge in an interactive environment. Here we present a very simple construction of such a continuum using an epistemic operator which we call "Knowing whether." The standard knowledge operator, by contrast, is "Knowing that." By saying "Alice knows whether the hatter is mad" we mean that either she knows that the hatter is mad or she knows that the hatter is not mad. Thus, while "Alice knows that the hatter is mad" implies that the hatter is indeed mad, "Alice knows whether the hatter is mad" is just about Alice's state of mind, with no implication regarding the sanity of the hatter.

We are making no philosophical claim about the primacy of either epistemic operator or the "correct" way to model knowledge. On the contrary. Both operators can be used in the same model since each can be expressed in terms of the other. Thus, "Bob knows whether X" is the same as "Bob knows that X or he knows that not X," while "Bob knows that X" is exactly the same as "X, and Bob knows whether X." The question of which operator to use is one of convenience.

What this paper demonstrates is that in some cases "Knowing whether" is much more convenient to use than "Knowing that." When one considers sentences (or events, depending on the formal framework) in which "Knowing that" operators of several agents alternate, as in "Bob knows that Alice does not know that he knows ...," one discovers that the logical relationship between such sentences can be rather complicated. This complexity is the reason why the construction of a continuum of states is cumbersome when the operator "Knowing that" is used as a building block. These complications disappear when "whether" is used instead of "that." The logical relationship between sentences formed by alternating "Knowing whether" operators of different agents is extremely simple; they are logically independent. This simplicity is at the base of the uncomplicated construction we present here. The usefulness of the "Knowing whether" operator in a different context is demonstrated in Heifetz and Samet [3].

The next section describes the formal framework of a state space with partitions, which is the standard model in game theory and economics. In Section 3 we show the difficulties one encounters when dealing with alternating "Knowing that" operators. In Section 4 we introduce the "Knowing whether" operator and construct a continuum of knowledge states.

2. Stating the Result Formally

We model interactive knowledge in a standard way, (see for example the survey, Geanakoplos [2]) by means of an *information structure* (Ω, Π_1, Π_2), where Ω is a state space and Π_1, Π_2 are partitions of Ω representing the information that agents 1 and 2 have about the states. Subsets of Ω are called *events*. The complementary event to E, $\Omega \backslash E$, is denoted by $\neg E$ (read "not E"). For $\omega \in \Omega$ and $i = 1, 2$ denote by $\Pi_i(\omega)$ the unique element of Π_i containing ω. We define two knowledge operators K_1 and K_2 for agents 1 and 2, $K_i \colon 2^\Omega \to 2^\Omega$, $i = 1, 2$, as follows. For each event E,

$$K_i E := \{\omega \in \Omega \mid \Pi_i(\omega) \subseteq E\}.$$

Fix a nontrivial event X. In order to describe all events that express interactive knowledge concerning X, we construct inductively a sequence $\mathscr{B}_0, \mathscr{B}_1, \mathscr{B}_2 \ldots$ of finite Boolean algebra[2] of events. $\mathscr{B}_0 := \{\Theta, X, \neg X, \Omega\}$, is the Boolean algebra generated by X. For $n \geqslant 1$, $\mathscr{B}_n$ is the Boolean algebra generated by $\mathscr{B}_{n-1} \cup \{K_i E \mid E \in \mathscr{B}_{n-1}, i = 1, 2\}$. Let $\mathscr{B} := \bigcup_{n=1}^{\infty} \mathscr{B}_n$. The events in $\mathscr{B}$ are said to be (*finitely*)[3] *generated by* X. Such events correspond to sentences constructed using the phrases "X," "not," "or," "and," "1 knows," and "2 knows." A typical event generated by X is, for example,

$$(K_1(K_2 X \cup \neg K_1 \neg X)) \cap K_2 K_1 X. \qquad (*)$$

With some abuse of language we call expressions like $(*)$ events generated by X, although strictly speaking such expressions describe events only for a given information structure and event X.

Some facts about X and knowledge concerning X, which may be quite complex, can be true in one state and false in another one. In such a case we say that the two states are *separated by* X. Formally,

[2] A Boolean algebra is a subset of 2^Ω closed under union, intersection, and complementation.

[3] Events generated by X for infinite ordinals are defined in Heifetz and Samet [3].

DEFINITION 2.1. Two states ω and ω' are separated by X if there exists an event E which is generated by X, such that $\omega \in E$ and $\omega' \in \neg E$.

One may wonder how many states can be in an information structure (Ω, Π_1, Π_2) such that an event X separates any two of them. We prove in Section 4:

THEOREM 2.2. *There exists an information structure* (Ω, Π_1, Π_2) *and an event* $X \subseteq \Omega$, *such that all the states in* Ω *are separated by* X, *and* Ω *has the cardinality of the continuum.*

Clearly the cardinality of such a set of states cannot be greater than $2^{\aleph_0}$ since there are only a denumerable number of events generated by X. Theorem 2.2 shows that this upper bound is indeed attained.

3. WHY ISN'T IT STRAIGHTFORWARD?

At first glance, the existence of a continuum of states describing interactive knowledge of X seems rather obvious. Consider the following construction.

Let $(E_1, E_2, \ldots)$ be a list (finite or infinite) of events generated by X. We say that the list is *consistent* if there is an information structure (Ω, Π_1, Π_2), an event $X \subseteq \Omega$, and a state $\omega \in \Omega$ such that all the events in the list are true in ω (that is, if the intersection of all the events in the list is non-empty).

Consider now the two lists,

$$(X, K_1 X), \qquad (X, \neg K_1 X).$$

Clearly, both are consistent, expressing the idea that the truth of X does not imply knowledge of X or lack of it. We can extend each of the two lists in two ways by applying knowledge of agent 2, or lack of it, to the last event in the list. Thus we obtain four lists,

$$(X, K_1 X, K_2 K_1 X), \qquad (X, K_1 X, \neg K_2 K_1 X),$$

from the first one, and

$$(X, \neg K_1 X, K_2 \neg K_1 X), \qquad (X, \neg K_1 X, \neg K_2 \neg K_1 X),$$

from the second. It is easy to see that all four lists are consistent by constructing an appropriate information structure. Applying now K_1 and $\neg K_1$ to the last event in each list we have already eight lists which again are consistent. We call a list of this type a *K-list*. We can go on, ad infinitum, generating infinitely many infinite *K*-lists, and the cardinality of this set of lists is $2^{\aleph_0}$.

The-fault in this line of proof is that not every K-list is consistent, as the following example shows. Consider the K-list $(E_1, ..., E_5)$ defined by

$$(X, K_1 X, \neg K_2 K_1 X, \neg K_1 \neg K_2 K_1 X, K_2 \neg K_1 \neg K_2 K_1 X).$$

This list is inconsistent; that is, for each information structure and event X the intersection of the events in the list is empty. Indeed, the intersection of the third and fifth events is already empty. We demonstrate this by a simple story followed by a rigorous proof.

Adam and Eve are waiting for a phone call from Godot. When finally Godot calls, only Adam is at home and he receives the call, an event which we denote by X. Obviously the event $E_2 = K_1 X$ holds in this case when we designate Adam as agent 1. Adam leaves a message on Eve's answering machine, telling her about the call he received. But Eve does not check her answering machine and so clearly the event $E_3 = \neg K_2 K_1 X$ is true, where Eve is agent 2. Adam does not know that Eve did not check her answering machine and thus the event $E_4 = \neg K_1 \neg K_2 K_1 X$ is also true.

Suppose now that Mr. Beckett, who is aware of all these events, brings to Eve's attention that Adam does not know that she does not know that he, Adam, received the call. In other words, Mr. Beckett tells her that E_4 holds true and thus the event $E_5 = K_2 \neg K_1 \neg K_2 K_1 X$ is now true. Is E_3 still correct? That is, knowing E_4, is Eve still ignorant of Adam receiving the call? Certainly not. Had Adam not received the call, he would have known for sure that Eve could not have possibly known that he received it, contrary to E_4. Therefore, knowing E_4, Eve concludes that Adam received the call; that is, she knows E_2 and hence E_3 is false. Thus E_5 contradicts E_3, or formally, $E_5 \cap E_3 = \emptyset$.

Now we state and prove formally:

PROPOSITION 3.1. *In any information structure and for any event X,*

$$(K_2 \neg K_1 \neg K_2 K_1 X) \cap (\neg K_2 K_1 X) = \emptyset.$$

Proof. To prove it we record three properties of the operators K_i that follow readily from the definition. Let K stand for either K_1 or K_2, then for any events E and F: (a) $KE \subseteq E$, (b) $K \neg KE = \neg KE$, (c) if $E \subseteq F$ then $KE \subseteq KF$.

Now by (a), $K_2 K_1 X \subseteq K1X$, and therefore by taking complements $\neg K_1 X \subseteq \neg K_2 K_1 X$. Hence by (c), $K_1 \neg K_1 X \subseteq K_1 \neg K_2 K_1 X$. Applying (b) to the left-hand side of this inclusion we have $\neg K_1 X \subseteq K_1 \neg K_2 K_1 X$. Taking complements again yields $\neg K_1 \neg K_2 K_1 X \subseteq K_1 X$. Finally by (c) we conclude that $K_2 \neg K_1 \neg K_2 K_1 X \subseteq K_2 K_1 X$ which completes the proof. ∎

Thus not all K-lists of length 5 are consistent. The list we examined is the shortest inconsistent one. The next new[4] consistent K-list is of length 9. Does the set of all consistent K-lists have the cardinality of the continuum? We do not know the answer, and it seems that to prove or disprove it is not straightforward. Aumann [1], constructs another continuum of consistent lists, but the construction is rather tricky and does not have the simplicity of the, alas, faulty K-lists. In the next section we construct an even simpler set of lists that does the job.

4. Knowing Whether

Define the operators J_1 and J_2 by

$$J_i E := (K_i E) \cup (K_i \neg E), \quad i = 1, 2$$

for each event E. We read $J_i E$ as "i knows whether E" meaning that agent i knows whether E is true or not, without saying whether E is true or not. Obviously $J_i E$ is the union of all elements of Π_i which are entirely included in either E or $\neg E$. It is easy to see that the operator K_i can be expressed in terms of J_i, namely, $K_i E = E \cap J_i E$. Note that the operators J_i are invariant under complementation; i.e., $J_i E = J_i \neg E$ for each event E.

Consider now lists of events of the same form as the K-lists of Section 3, only with the operators J_i replacing the K_i operators. We call such a list a *J-list*. A typical J-list is, for example,

$$(X, J_1 X, \neg J_2 J_1 X, \neg J_1 \neg J_2 J_1 X, J_2 \neg J_1 \neg J_2 J_1 X, ...).$$

But in light of the complementation invariance of the J operators we can omit all complementation signs which do not precede an event. Thus the previous list is the same as

$$(X, J_1 X, \neg J_2 J_1 X, \neg J_1 J_2 J_1 X, J_2 J_1 J_2 J_1 X, ...).$$

Thus each event in a J-list is generated by applying alternately J_1 and J_2 or J_2 and J_1, starting from X, and then possibly taking the complement. We show now that, as opposed to the K-lists, all $2^{\aleph_0}$ infinite J-lists are consistent. This follows from the proof of Theorem 2.2 below.

Proof of Theorem 2.2. Let Ω be the set of all pairs of infinite sequences of 0's and 1's with the same starting digit; that is,

$$\Omega = \{(a_0 a_1 a_2 ..., b_0 b_1 b_2 ...): a_k, b_k \in \{0, 1\} \ \forall k \geq 0, a_0 = b_0\}.$$

Clearly, Ω has the cardinality of the continuum.

[4] That is, a K-list which does not include as a sublist an inconsistent list of length 5.

Let the event X be the set of all states in which $a_0 = b_0 = 1$. Define an equivalence relation $\sim_1$ on states by

$$(a_0 a_1 a_2 \ldots, b_0 b_1 b_2 \ldots) \sim_1 (a'_0 a'_1 a'_2 \ldots, b'_0 b'_1 b'_2 \ldots)$$

if and only if

$$\text{for all } k \geqslant 1\text{: } a_k = a'_k \text{ and } a_k = 1 \Rightarrow b_{k-1} = b'_{k-1}.$$

Let Π_1 be the partition of Ω into equivalence classes of $\sim_1$. Define analogously $\sim_2$ and Π_2, interchanging the roles of the a's and the b's.

For all $k \geqslant 0$ let the event $[a_k = 1]$ be the set of states whose a_kth coordinate is 1. Define similarly the events $[a_k = 0]$, $[b_k = 1]$, and $[b_k = 0]$. Note that $[a_k = 0] = \neg[a_k = 1]$ and $[b_k = 0] = \neg[b_k = 1]$.

LEMMA 4.1. *For all* $k \geqslant 0$

$$[a_k = 1] = \underbrace{J_1 J_2 J_1 \cdots}_{k \text{ operators}} X, \qquad [a_k = 0] = \neg \underbrace{J_1 J_2 J_1 \cdots}_{k \text{ operators}} X,$$

$$[b_k = 1] = \underbrace{J_2 J_1 J_2 \cdots}_{k \text{ operators}} X, \qquad [b_k = 0] = \neg \underbrace{J_2 J_1 J_2 \cdots}_{k \text{ operators}} X.$$

With this lemma the theorem is proved. Any two different states in Ω differ by either the a_kth or the b_kth coordinate for some $k \geqslant 0$. If $k = 0$, the two states are separated by the event X itself (here we use the equality $a_0 = b_0$ in the definition of a state). If $k \geqslant 1$, then by the lemma the two states are separated by some event generated by X. ▮

Lemma 4.1 also shows that every J-list is consistent. Indeed all the events in such a list hold true in all states with coordinates given by the events in the list according to the lemma. Proving the lemma is now straightforward.

Proof of Lemma 4.1. We prove by induction on k. For $k = 0$ the equality in the lemma holds by the definition of X. It suffices now to show that for all $k \geqslant 1$,

$$[a_k = 1] = J_1[b_{k-1} = 1],$$

$$[b_k = 1] = J_2[a_{k-1} = 1].$$

We prove the equality for $[a_k = 1]$. By the definition of J_1 we have to show that

$$[a_k = 1] = (K_1[b_{k-1} = 1]) \cup (K_1[b_{k-1} = 0]).$$

Suppose that $\omega \in [a_k = 1]$ and denote by β the b_{k-1}th coordinate of ω. Then by the definition of $\sim_1$, any state which is $\sim_1$ equivalent to ω must have the same b_{k-1}th coordinate. Thus $\Pi_1(\omega) \subseteq [b_{k-1} = \beta]$ and therefore $\omega \in K_1[b_{k-1} = \beta]$.

Conversely, assume that $\omega \notin [a_k = 1]$. Let ω' be the state which is identical to ω in all coordinates except b_{k-1}. Then $\omega' \sim_1 \omega$, which implies that $\omega \notin K_1[b_{k-1} = \beta]$ for both $\beta = 0$ and $\beta = 1$. ∎

References

1. R. J. Aumann, Notes on interactive epistemology, Cowles Foundation for Research in Economics working paper, 1989.
2. J. Geanakoplos, Common knowledge, *in* "Handbook of Game Theory with Economic Applications" (R. J. Aumann and S. Hart, Eds.), Vol. II, Chap. 40. Elsevier/North-Holland, Amsterdam, 1995.
3. A. Heifetz and D. Samet, Universal partition structures, Faculty of Management, Tel Aviv University, IIBR Working Paper No. 26/93, 1993.