Faculty Working Papers

LEARNING TO AGREE

Jonathan A. K. Cave, Lecturer, Department of Economics

#633

College of Commerce and Business Administration
University of Illinois at Urbana-Champaign

FACULTY WORKING PAPERS

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

January 2, 1980

LEARNING TO AGREE

Jonathan A. K. Cave, Lecturer, Department of Economics

#633

Summary:

In this note, we generalize some of the results pertaining to the phenomenon of common knowledge. An event is common knowledge if each agent knows it, each knows that the others know it, etc. We define a class of procedures whereby agents take actions based on private information, and make further inferences from their observations of the actions of other agents. We are able to characterize the information thus revealed and obtain general versions of theorems to the effect that agents cannot agree to disagree, and will not make trades based solely on differences in information.

Learning to Agree

I. Introduction

In this paper we explore some of the implications of the idea of "common knowledge" and define a general model whereby agents can make inferences about the knowledge held by other agents on the basis of actions those agents are publicly observed to take. The intuitive data of common knowledge is that an event is known to all (each agent can tell whether or not it has occurred), and each agent knows that every other agent knows this, each agent knows that every other agent knows that every other agent knows this, etc. etc. The discussion of the phenomenon of common knowledge began with the publication by Aumann (Al) of a result which holds that, if two agents form conditional probability assessments of the likelihood of a given event, and if these assessments are common knowledge, then it is not possible for the assessments to differ. Milgrom (Ml) then gave a formal characterization of common knowledge that expressed Aumann's definition in Axiomatic terms. Geanakopolous and Polemarchakis (G&P) extended the collection of events to which the appelation "common knowledge" could be applied by defining an explicit communications mechanism whereby agents exchange conditional probability assessments about the likelihood of a given event, revise their private beliefs in the light of these disclosures, and continue until a sort of steady state is reached at which there is no further revision. Another sort of communications mechanism was defined by Milgrom and Stokey (M&S) in the context of a market with uncertainty and risk-averse agents; the

---

I am indebted to Francoise Schoumaker and Al Roth for helpful discussions.

information exchanged constituted a set of feasible and conditionally individually rational trades, and the result was the no-trade point; risk-averse agents will not make trades solely on the basis of differences in information.

What we shall do in this paper is to present simple definitions of the idea of common knowledge consistent with a general model of a communications process. With the aid of this model, we shall be able to reproduce the existing results and obtain some extensions of them. In fact, it will turn out that we can characterize any such model in terms of the common-knowledge partition on the states of nature that characterizes equilibrium. Another advantage is that we shall be able to extend the notion that agents cannot "agree to disagree" to a more general context, and relax some of the stringent assumptions necessary for that result.

The plan of the paper is as follows: in Section II we record for posterity in its most general form, the story of the "adulterous couples" which speakers on the topic are fond of reciting, but which has never appeared in print, to our knowledge. In Section III we present the model of information, define the general communication mechanism, and describe its equilibria. In Section IV we give the main results and indicate how they relate to the literature. Section V contains some examples and lists some open problems that are currently being explored.

II. The Adulterous Couples

This is a story whose origins are unclear, but one that clearly demonstrates the principle of common knowledge. It belongs to a class of stories usually couched in terms of people with marks on their foreheads trying to determine something about their own mark, but the present version has slightly more intrinsic interest, if not plausibility.

The story goes as follows: in a certain country, the rule is that any wife who can prove that her husband is committing adultery must brand him on the forehead with the letter "A". Moreover, it is common knowledge that no wife knows whether or not her husband is unfaithful, but that each wife knows how many of the other husbands are. In other words, if there are no couples, and k of the husbands are unfaithful, each wife knows (and is known to know ...) that either k or k-1 of the n-1 other husbands are cheating. However, since no wife knows the truth about her own husband, matters continue in this wise without any branding of husbands. Continue, that is, until a travelling moralist, excessively concerned with other people's business, happens to visit the country. Outraged by the immorality he detects, he calls a meeting which is attended by the wives, at which he announces "there are at least m unfaithful husbands" (where $m \leq k \leq n$). As long as $m < k$, this comes as no news to anyone, so that everyone believes the moralist, and nothing much appears to change. Disgruntled, our latter-day Muggeridge leaves in a huff. For a while nothing happens, and matters go on as usual. However, on the $k-m+1^{st}$ day all the guilty husbands are branded!

We shall illustrate the process of deduction for two cases: $n=k=m+1=2$; and $n=k=m+2=3$, from which it should be clear how one proceeds in the general case. (The case n=4, m=1 is worked out as an application of our model in Section V below.) It should also be clear from this exposition that nothing like the entire infinite regress of 'I knows that II knows that I knows that II knows that ...' involved in common knowledge is required to make this work: I am indebted to Al Roth for this observation.

In the case n=k=m+1=2, wife 1 reasons as follows: "I know that husband 2 is an adulteror, since he is involved with me, but I do not know whether my own husband is. However, I know that wife 2 does know this. Therefore, if my husband is faithful, then wife 2 will conclude from the moralist's announcement that it is her own husband who is the cheat. Therefore, if my husband is faithful, husband 2 will be branded tomorrow." Wife 2 reasons similarly, of course, and so neither husband gets branded on the first night. However, each wife then learns the truth from the inaction of the other and on day 2 both husbands come home to find the iron hot and waiting.

The case n=k=m+2=3 is slightly more complicated. We illustrate only wife 1's chain of reasoning. "If my husband is faithful, then wife 2 will know this. This means that wife 2 will observe that one other husband is unfaithful, although she may not know which one. She will then suppose that, if her husband (husband 2) is faithful, then one of either wife 3 or myself will observe no adulteries. That one of us will brand our husband tonight." Each of the other wives reasons similarly, so there is no branding on the first night. On the second day, when wife 2 observed that neither husband 1 nor husband 3 has been branded, then she should go home and brand her own husband, according to wife 1's reasoning. When she does not, and the third day dawns with no branded husbands, all the chains of reasoning based on the supposition by wife 1 that her husband is innocent collapse, and each infers the awful truth.

Obviously, what is important in this class of examples is that each person knows enough about what the other ones know to be able to interpret what they do. This is the importance of common knowledge, and the reason why it is not enough merely for each person merely to know what information about the true state the other people know. In other words, it is necessary that wife 2 knows that wife 3 knows that wife 1 knows whether 0, 1, or 2 of husbands 2 and 3 is unfaithful, so that wife 1 can form an accurate appraisal of wife 2's assessment of wife 3's actions. In the following section, we present a model wherein these notions can be given precise definition and their implications explored.

III. The Model

We begin with a measureable space (Ω,β) where β is a σ-field of subsets of Ω. This space constitutes the states of nature. For precise definitions of these and other measure-theoretic concepts, the reader is referred to any standard textbook on measure theory, such as Halmos (H). We are also given a finite set $N = \{1,\ldots,n\}$ of agents. Each agent i is endowed with a measureable partition P^i on the states of nature, and with a prior μ^i, which is a probability measure on the space (Ω,β). We shall interpret these objects as follows: if the state $\omega\epsilon\Omega$ occurs, agent i is told only that one of the states in that element of P^i that contains ω has occurred, but not which one. We shall denote this collection of states by $P^i(\omega)$ and refer to it as agent i's private information. An event is any member of $B\epsilon\beta$, and each agent possesses a prior belief as to the likelihood of this event given by $\mu^i(B)$. After the true state has been chosen, and the agent has received his private information, he can form a posterior belief as to the likelihood of B's occurrence, given by $q^i(B;\omega) = \frac{\mu^i(B\cap P^i(\omega))}{\mu^i(P^i(\omega))}$; this is easily seen to be a simple conditional probability. Before we describe the model any further, we must define some operations on partitions. In keeping with the above notation, if R is any partition of Ω, and if $\omega\varepsilon\Omega$ is any state, we shall denote by $R(\omega)$ that element of R that contains ω. Now suppose that R and S are two different partitions. We say that R is <u>finer than</u> S (S is <u>coarser than</u> R) iff every set in S is a union of sets in R. We write this relation $R \supset S$. Given any partition R we define the <u>field generated by R</u>, F(R), to be the collection of all sets which are unions of sets in R, together with the empty set. For full generality, we note that we can define F(S) for any collection of sets S by taking the closure of S under complementation and unions and intersections, together with the empty set. Given any two partitions R and S we define two new partitions:

i) their <u>meet</u> $R \wedge S = \{B\varepsilon\beta:\ B\varepsilon F(R) \text{ and } B\varepsilon F(S)\}$ - this is the partition consisting of sets that can be 'detected' using either R or S; the finest common coarsening of R and S; and

ii) their <u>join</u> $R \vee S = \{B\varepsilon\beta:\ B = B_1 \cap B_2, \text{ where } B_1\varepsilon F(R),\ B_2\varepsilon F(S)\}$ - this is the collection of sets that can be detected using both R and S together; the coarsest common refinement of R and S.

We should note that, as we have defined them, neither the meet nor the join is really a partition, but is actually the field generated by some partition. We shall refer to these two objects interchangeably, when no confusion will result.

Returning now to the general model, we define two partitions that summarize everyone's information:

$M = \bigwedge_{i \in N} P^i$ is the <u>meet</u>, and $J = \bigvee_{i \in N} P^i$ is the <u>join</u>

Clearly, $M \subset J$, M is the set of events that everyone can detect acting indepedently, while J is the set of events that can be detected if agents pool their information. One further assumption that we shall make is that the join consists of non-null events. In other words, if $B \epsilon J$, and $B \neq \emptyset$, then $\mu^i(B) > 0$, for all i.

We now turn to the definition of common knowledge.

<u>Definition</u>: Let $A \epsilon \beta$, and $\omega \epsilon \Omega$. A is said to be <u>common knowledge at ω</u>, if $A \supset M(\omega)$; in other words, if every agent knows whether or not A has occurred when the true state is ω, and every agent knows that every other agent knows whether A has occurred, ... Implicit in this is the assumption that the partitions themselves are common knowledge. The event A is said to be <u>common knowledge</u> iff it is common knowledge at ω for every $\omega \epsilon A$; in other words, iff $A \epsilon F(M)$.

<u>Remarks</u>: We shall later be concerned with the difference between an event which is common knowledge in some states and an event which is common knowledge. Here we present a simple example. Suppose that there are two agents and four states of the world, and that the partitions are given by: $P^1 = (s_1 s_2)(s_3 s_4)$; $P^2 = (s_1)(s_2 s_3)(s_4)$. Consider the event $(s_3 s_4)$: if the true state is s_4, then both agents will know that $(s_3 s_4)$ has occurred, but if the true state is s_3, only agent 1 will know that $(s_3 s_4)$ has occurred. Another point we should make at this juncture is that we are also assuming that the priors μ^i are common knowledge. In most of the work to date, it has been assumed both that they are common knowledge and that they are the same. We shall see to

what extent the assumption of equal priors is necessary in what follows. Aumann discusses the assumption, and mentions that Harsañyi (H2) has defended it on the grounds that the only reason for agents to form different priors is that they have been given different information at some point in the past, and that the supposedly different priors are in fact different posteriors derived from identical priors. While it is difficult to fault this argument if we accept a sufficiently broad concept of information, such a notion of information puts a great deal of strain on the assumption that the partitions are common knowledge. Therefore, while we shall sometimes use the assumption that agents have identical priors, some of our results are independent of this assumption, and we do not wish to prejudice the issue.

We now define a process by which agents can communicate with each other, albeit in what may be a fairly indirect fashion. Let us add to the structure above another piece of common knowledge. For each agent i, we shall define an <u>action rule</u> f^i: $\beta \to Z^i$, where Z^i is some space of actions. The procedure we have in mind can be loosely described as follows: in the beginning, each agent is given some private information $P^i(\omega)$ about the true state $\omega \in \Omega$, and takes the appropriate action $f^i(P^i(\omega))$, whose value becomes common knowledge, along with the actions of the other agents. On the next day, each person forms a new assessemnt as to the true state of the world based on private information and the information revealed by the other agents. Along with this assessment goes an assessment as to what the other agents beliefs on the second round can be, since each agent goes through the same process of revision. Thus, on day 2 agent i will have some subset of $P^i(\omega)$ which is consistent with private information and public observation. Also, conditional on any state in this set, i will be able to work out what j believes, what j believes k believes, and so on. Fortunately, we shall be able to obtain simple expressions of these statements. On this second day, each agent takes the action appropriate to its current information, and the process continues until an equilibrium is reached.

Formally, we wish to concern ourselves with the following objects; actions f_t^i taken by agent i at stage t, and <u>sets of states</u> $H_t(f_t,\ldots f_1)$ consistent with a sequence of actions. These objects are defined inductively as follows:

$$f_1(\omega) = [f_1^1(\omega),\ldots,f_1^n(\omega)], \text{ where } f_1^i(\omega) = f^i(P^i(\omega))$$

$$H_1(f_1) = \{\omega'\varepsilon\Omega: \ f_1(\omega') = f_1\}$$

$$f_2(\omega) = [f_2^1(\omega),\ldots,f_2^n(\omega)], \text{ where } f_2^i(\omega) = f_2^i(\omega) = f^i(H_1(f_1(\omega)) \cap P^i(\omega))$$

$$\cdot \qquad \cdot \qquad \cdot$$

$$H_t(f_t,\ldots,f_1) = \{\omega'\varepsilon H_{t-1}(f_{t-1},\ldots,f_1): \ f_t(\omega') = f_t\}$$

$f_t(\omega) = [f_t^1(\omega),\ldots,f_t^n(\omega)]$, where $f_t^i(\omega) = f^i(H_{t-1}(\cdot) \cap P^i(\omega))$, and the argument of H_{t-1} has bee suppressed for brevity.

The reader can convince him/herself that what we have described is just what we said in words above, and that it is indeed the best that can be done. The events in the sets H_t are a matter of public record, and therefore constitute a sort of common knowlege belief about the true state. To find any agent's private belief about the true state at time t it suffices to form $P^i(\omega) \cap H_t$. Agent i can do this for itself with no trouble, and can therefore take the appropriate action at each stage. Agent i can also form $P^j(\omega') \cap H_t$ for all $\omega' \epsilon P^i(\omega) \cap H_t$; this is the set of possible $t^{\underline{th}}$ round beliefs of agent j, according to agent i's best information. In the same way, we can derive agent i's assessment of what j thinks that k thinks, and so on. We now define an equilibrium.

Definition: an inference equilibrium consists of an integer T, a sequence of actions $f_1,\ldots f_T$, and a subset H_T of Ω, with the property that $H_T = H_T(f_T,\ldots f_1) = H_{T-1}(f_{T-1},\ldots,f_1)$.

In other words, an inference equilibrium, or IE, is a situation where no further revision of information takes place. It is immediate that for all $k, H_{T+k} = H_T$, and $f_{T+k} = f_T$; so that this is truly an equilibrium. We should also point out that everything in the model is a deterministic function of the true state ω, so that we can also define the inference equilibrium at ω to be just the common-knowledge belief to which the system converges given that the true state is ω. This belief will be denoted $H(\omega)$.

IV. Results

In this section we present several results pertinent to the model of the previous section.

Theorem I: (existance) for an $\omega \epsilon \Omega$, there exists a unique inference equilibrum.

Proof: it is trivial to observe that, for any $t, H_{t+1} \subset H_t$, so that the map H is a contraction. Since it is also single-valued (as long as the f^i are), it must have a unique fixed point. QED

Theorem II: (convergence) for any $\omega \varepsilon \Omega$, if $M(\omega) \cap P^i$ is finite for each i, the inference equilibrium is achieved in a finite number of steps, bounded above by

$$\sum_{i \varepsilon N} \#(M(\omega) \cap P^i)$$

Proof: obvious; we merely remark that what we mean by $\#(M(\omega) \cap P^i)$ is the number of elements of P^i (not the number of states belonging to these elements) that belong to $M(\omega)$, and that the definition of $M(\omega)$ assures us that any element of P^i that intersects $M(\omega)$ is contained in $M(\omega)$ for each i and ω.

Theorem III: H constitutes a partition of Ω.

Proof: suppose to the contrary that there exist distinct states ω and ω' such that

$$H(\omega) \neq H(\omega) \cap H(\omega') \neq \emptyset$$

Now consider as well the sequences $f = f_1, \ldots, f_T$ and $f' = f'_1, \ldots, f'_T$ of actions that establish these inference equilibria. First, suppose that $f_1 \neq f'_1$. By definition of H_1 and the fact that f^i_1 is the value of f^i on a member of a partition, we can see that $f_1 \neq f'_1$ implies $H_1(f_1) \cap H_1(f'_1) = \emptyset$. Therefore, we must have $f_1 = f'_1$. Now suppose that we have $f_t = f'_t$ for all $1 \leq t < s$, but that $f_s \neq f'_s$. Therefore, we have $H_{s-1}(f_{s-1}, \ldots, f_1) = H_{s-1}(f'_{s-1}, \ldots, f'_1)$. Now each agent's partition P^i restricts to a partition of H_{s-1}, so that $f_s \neq f'_s$ implies $H_s(f_s, f_{s-1}, \ldots, f_1) \cap H_s(f'_s, f_{s-1}, \ldots, f_1) = \emptyset$. It will be noted that we made use of $f_t = f'_t$ in writing this last equation, since it would not be true otherwise. At any rate, this provides the inductive step and proves the theorem. QED

Remark: this is the most important theorem so far, and justifies the construction of our general model, since the inference partition H thus obtained is characteristic of the communications rule f, and provides a method by which various structures can be compared as to the degree to which information gets revealed as well as which agents obtain which information. It also considerably simplifies the construction and analysis of examples, which we shall address in the next section. However, before we proceed to the specific examples, we shall prove some other properties that are of interest in certain special cases.

Definition: a collection of action rules $f = f^1,\ldots,f^n$ is symmetric iff $f^i = f$ for all i. An action rule is union-consistent if, for any disjoint sets $B, C \varepsilon \beta$, we have:

$$f(B) = f(C) \text{ implies } f(B) = f(B \cup C) = f(C)$$

Examples of union consistent rules include: conditional probabilities for fixed events or collections of events; conditional expectations of random variables; and actions which maximize conditional expectations of functions of random variables. In short, these examples cover most of the concrete applications of communication procedures that have been proposed to date.

Theorem IV: (the impossibility of agreeing to disagree): If f is a symmetric and union-consistent action rule, and H is the inference partition corresponding to f, then for every $\omega \varepsilon \Omega$, $f^i(P^i(\omega) \cap H(\omega)) = f^j(P^j(\omega) \cap H(\omega))$. In other words, at an inference equilibrium, all agents will take exactly the same actions.

Proof: By definition of an IE, for each $i \epsilon N$, and each $\omega' \epsilon H(\omega)$, we must have $f(P^i(\omega) \cap H(\omega)) = f(P^i(\omega') \cap H(\omega))$. If this were not true, then further inference would be possible. However, by the union condition it follows that:

$$f(H(\omega)) = f(P^i(\omega) \cap H(\omega))$$

since $H(\omega)$ is just the intersection of sets of the form $P^i(\omega') \cap H(\omega)$, over $\omega' \epsilon H(\omega)$, all of which lead to the same action. This condition is independent of i, so the theorem is proven. QED

Another question we might ask is what sort of events are common knowledge at an inference equilibrium? Now, it is certainly going to be the case that the inference partition is (weakly) finer than M and (weakly) coarser than J, but the exact ranking of H will depend on f. However, from the definition of common knowledge, we can provide a superficial answer to the above question. At the inference equilibrium, the information possessed by each agent is represented by $P^i(\omega) \cap H(\omega)$. In other words, an agent in a model such as ours can look forward to having the partition $P^i \vee H$. The events which are common knowledge at IE are those belonging to the field generated by the meet of these "final" partitions:

<u>Proposition V</u>: $A \epsilon \beta$ is common knowledge at the inference equilibrium iff $A \epsilon F(M \vee H)$.

Proof: by definition, A is common knowledge iff $A \epsilon F(\bigwedge_{i \epsilon N} (P^i \vee H)) =$

$F((\bigwedge_{i \epsilon N} P^i) \vee H) = F(M \vee H)$. QED

Another question we can ask is whether there is a simple expression for H in terms of P^i and f. For example, if f^i = constant all i, we have H = M, while if $f^i(B) = B$ (or some sufficient statistic), we obtain H = J. While we have no simple expression, we can observe the following: $H_1 = \bigvee_{i\varepsilon N} (f^i)^{-1}(P^i)$, $H_2 = \bigvee_{i\varepsilon N} (F^i)^{-1}(P^i \vee H_1)$, etc. From this expression, we can obtain characterizations of H in various special cases and also derive some conditions under which we get full revelation: H = J. However, those are topics for a subsequent paper.

V. Examples

In this section, we list some of the examples that have been used in the literature. Perhaps the most important, at least historically, is the action rule specified by Aumann and G&P:

$$f^i(B) = \frac{p(P^i(\omega) \cap B)}{p(P^i(\omega))}$$

where p is the common prior of all the agents. This action rule is symmetric and union-consistent, so the theorems of those two papers are special cases of our theorems I, II, and IV above, for the case where there are only two agents. Another example is that used by Milgrom and Stokey, where the action rule is specified for the whole economy. They do not actually specify an action rule in terms of what an agent knows, but one can infer an action rule from their notion of what is common knowledge at an equilibrium. In essence, they have their agents submit a vector of net trades that is chosen from a set of feasible n-tuples of net trades that are individually-rational conditional on each agent's private information. One can also imagine somewhat more involved formulations of action rules for economic situations; examples have been described by Aumann (A2), Jordan (J), Cave (C) and Radner (R). A common thread in all of these models is that agents make some trades that maximize conditional expected utility based on their private information. The results of these trades (which may be in the form of clearing prices, clearing prices and final allocations, etc.) are then made common knowledge, information is revised and the process repeats. The endpoint of these adjustment procedures is a situation where no further information is conveyed; this means that the final allocation is both <u>ex ante</u> and <u>ex post</u> efficient; it constitutes a <u>rational expectation equilibrium</u> relative to the information structure which gives each agent the partition $P^i \vee H$. Much of the literature on rational expectations, trading with differential information, etc. can be understood in this context. One final class of examples is found in the literature on repeated games of incomplete information, where explicit learning processes are a result of equilibrium behavior.

We conclude the discussion of examples by working out the tale of the adulterous couples for the case n=4, m=1, all values of k. In this case, we can represent the state of nature as $\omega = (h_1,h_2,h_3,h_4)$, where

$$h_i = \begin{cases} 0 \text{ if husband i is innocent} \\ 1 \text{ if husband i is an adulteror} \end{cases}$$

It will be convenient to introduce a more condensed notation for the states and we shall represent them as the numbers of which the original states were binary representations, viz;

$$\omega(h_1,h_2,h_3,h_4) = 8h_1 + 4h_2 + 2h_3 + h_4$$

The action rule in terms of the original state space is:

$$f^i(S) = \begin{cases} \text{B(brand) iff } h_i = 1 \text{ for all } \omega\epsilon S \\ \text{N (not brand) otherwise} \end{cases}$$

In what follows, we shall mark the states where the indicated player is to take action B with an asterisk. The information structure is as follows: we indicate the $k^{\underline{th}}$ element of agent i's partition by P^i_k

i	P^i_1	P^i_2	P^i_3	P^i_4
1	(0,8*)	(1,2,4,9*,10*,12*)	(3,5,6,11*,13*,14*)	(7,15*)
2	(0,4*)	(1,2,5*,6*,8,12*)	(3,7*,9,10,13*,14*)	(11,15*)
3	(0,2*)	(1,3*,4,6*,8,10*)	(5,7*,9,11*,12,14*)	(13,15*)
4	(0,1*)	(2,3*,4,5*,8,9*)	(6,7*,10,11*,12,13*)	(14,15*)

At the present time, $M = \Omega$ (the coarse partition, and $J = [\{0\},\ldots,\{15\}]$ - the fine partition, $H = M$ and we are in inference equilibrium. Now suppose that we add the following piece of common knowledge: the true state is <u>not</u> (0,0,0,0). The information partitions remain the same, except that $P^1_1 = (8)$, $P^2_1 = (4)$, $P^3_1 = (2)$, and $P^4_1 = (1)$. It is therefore obvious that if the true state is a member of $\{1,2,4,8\}$, $H(\omega) = \{\omega\}$, and convergence is immediate (takes one 'day'). In fact, we have that, for example, $\omega = 1$ implies $f_1(\omega) = N,N,N,B$; and $H(\omega) = H_1(N,N,N,B) = \{1\}$. For any other states, we have $f_1(\omega) = N,N,N,N$, so that $H_1(N,N,N,N) = (3,5,6,7,9,10,11,12,13,14,15)$. We can show what happens in each of these states in the following table, where we have shown the true state, the second-round beliefs of each agent in each state, and the second-round actions of the agents.

state 1	$P^1(\omega)$ H_1	$P^2(\omega)$ H_1	$P^3(\omega)$ H_1	$P^4(\omega)$ H_1	f
3	(3,5,6,11,13,14)	(3,7,9,10,13,14)	(3,6,10)*	(3,5,9)*	NNB
5	(3,5,6,11,13,14)	(5,6,12)*	(5,7,9,11,12,14)	(3,5,9)*	NBN
6	(3,5,6,11,13,14)	(5,6,12)*	(3,6,10)*	(6,7,10,11,12,13)	NBB
7	(7,15)	(3,7,9,10,13,14)	(5,7,9,11,12,14)	(6,7,10,11,12,13)	NNN
9	(9,10,12)*	(3,7,9,10,13,14)	(5,7,9,11,12,14)	(3,5,9)*	BNN
10	(9,10,12)*	(3,7,9,10,13,14)	(3,6,10)*	(6,7,10,11,12,13)	BNB
11	(3,5,6,11,13,14)	(11,15)	(5,7,9,11,12,14)	(6,7,10,11,12,13)	NNN
12	(9,10,12)*	(5,6,12)*	(5,7,9,11,12,14)	(6,7,10,11,12,13)	BBN
13	(3,5,6,11,13,14)	(3,7,9,10,13,14)	(13,15)	(6,7,10,11,12,13)	NNN
14	(3,5,6,11,13,14)	(3,7,9,10,13,14)	(5,7,9,11,12,14)	(14,15)	NNN
15	(7,15)	(11,15)	(13,15)	(14,15)	NNN

Therefore, the states that get discovered after the second round are (3,5,6,9,10,12), so for those states ω, we have $H(\omega) = \{\omega\}$. For the others, where F_2 = NNNN, the new common knowledge belief is given by H_2((NNNN),(NNNN)) = (7,11,13,14,15). We can draw the same kind of table to represent the outcomes on the third round for each of these states:

state	$P^1(\omega)$ H_2	$P^2(\omega)$ H_2	$P^3(\omega)$ H_3	$P^4(\omega)$ H_4	f
7	(7,15)	(7,13,14)*	(7,11,14)*	(7,11,13)*	NBBB
11	(11,13,14)*	(11,15)	(7,11,14)*	(7,11,13)*	BNBB
13	(11,13,14)*	(7,13,14)*	(13,15)	(7,11,13)*	BBNB
14	(11,13,14)*	(7,13,14)*	(7,11,14)*	(14,15)	BBBN
15	(7,15)	(11,15)	(13,15)	(14,15)	NNNN

Thus all the information is revealed after the third round, although it may take until the fourth day for the parties to take appropriate action in case the true state were 15. It is therefore obvious that H = J = the fine partition, and that this way of dealing with the adulterous couples story is much clearer and easier to generalize than the verbal approach adopted in Section II above.

There remain several open questions in regard to these mechanisms, in addition to those already raised. For example, if we relax the condition of identical priors, the Aumann/G&P action rule is no longer symmetric; though it is indeed union-consistent. In that case, one can easily modify the results of Theorem IV, but much of its intuitive appeal is lost; it becomes possible for people to agree to disagree. It would be interesting to investigate the effect of different priors on the other union-consistent rules mentioned after the definition. Another line of inquiry is to investigate what the interdependence is between event A about which the Aumann/G&P agents communicate and the degree of revelation. In particular, if we allow agents to communicate about some collection of events, it seems likely that we could improve the performance of the mechanism. An obvious example is having agents communicate their conditional probabilities on all the events in J: this leads to convergence to H = J in one round. Another question left unanswered is whether certain classes of events remain in the middle range between common knowledge at ω, and common knowledge when we move to IE. In particular, does the G&P rule imply that we converge to siutations where the target event becomes common knowledge? The answer is probably no, but the pursuit may turn up some interesting results. So would the analysis of strategic behavior in these models.

References

A1: Aumann, R. J. "Agreeing to Disagree," Ann. Stat. 4/6/1976.

A2: Aumann, R. J. "A General Equilibrium Model with Information," unpublished.

C: Cave, J. "Games of Incomplete Information with Strategic Communication," unpublished.

G&P: Geanakopolos, J. and Polemarchakis, H. "We Can't Disagree Forever," Stanford IMSSS Technical Report 277.

H1: Halmos, P. Measure Theory, New York, Van Nostrand, 1950.

H2: Harsanyi, J. "Games of Incomplete Information Played by Bayesian Players, I-III," Management Sci., Vol. 14, 1967-8.

J: Jordan, J. S. "Expectations Equilibrium with Expectations Conditioned on Past Data," J.E.T.

M: Milgrom, P. "A Formal Characterization of Common Knowledge," Northwestern University, W.P.393, August 1979.

M&S: Milgrom, P. and Stokey, N. "Information, Trade and Common Knowledge," Northwestern University, WP 377, April 1979.

R: Radner, R. "Existance of Equilibrium Plans, Prices and Price Expectations in a Sequence of Markets," Econometrica 40, 1972.

M/D/260