

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

**Box 2125, Yale Station
New Haven, Connecticut 06520**

COWLES FOUNDATION DISCUSSION PAPER NO. 639

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

WE CAN'T DISAGREE FOREVER

John D. Geanakoplos and Heraklis M. Polemarchakis

July 29, 1981

We Can't Disagree Forever*

by

John D. Geanakoplos and Heraklis M. Polemarchakis**

July 1978

Final Revision, December 1980

*This research was partially supported by National Science Foundation Grant SOC 77-0600 to the Institute for Mathematical Studies in the Social Sciences at Stanford University.

**The second author's work was partially supported by National Science Foundation Grant SES 78-25910 to Columbia University.

Abstract

Under the assumption of common priors, if the information partitions of two agents are finite, then simply by communicating back and forth and revising their posteriors the two agents will converge to a common equilibrium posterior, eventhough they may base their posteriors on quite different information. Furthermore, given any integer, n , one can construct an example in which the revision process not only takes n steps to converge, but no evident revision occurs -- for $(n-1)$ steps both agents repeat the same conflicting posteriors -- until the last step when the two agents decide to agree. Common knowledge of each other's posterior does not necessarily lead agents to the posterior they would have agreed upon had information been directly exchanged. On the other hand, the examples that are characterized by a discrepancy between the direct and indirect communication equilibrium, are rare: with probability 1, the revision process constructed here leads the two agents in one step to the direct communication equilibrium.

WE CAN'T DISAGREE FOREVER

by

John D. Geanakoplos* and Eeraklis M. Polemarchakis**

Opinions depend on prior beliefs held and additional information received. Whether communication between agents is sufficient to eliminate differences of opinion is a question of primary interest—normative as well as positive—for the theory of decentralized decision making. Aumann [1976] gives an elegant and elementary answer to this question. He demonstrates that if two persons have identical priors, and if their posteriors for a given event A are common knowledge, then these posteriors must be equal, even though the two agents may base their posteriors on quite different information. Aumann goes on to emphasize that the key notion is that of common knowledge, which he explains intuitively as follows: Call the two people α and β . To say that an event is common knowledge means more than just that α and β know it; it must also be the case that α knows that β knows it, β knows that α knows it, α knows that β knows that α knows it, and so on. If one merely assumes that the two persons know each other's posteriors, the result is not true.

In the present paper we want to answer a number of questions left open by Aumann's argument. First, observe that the equality of the posteriors is claimed by Aumann only for those (rare) events whose posteriors happen to be common knowledge. No process is described according to which an

* Cowles Foundation, Yale University.

** Department of Economics, Columbia University.

We are grateful to Robert Aumann for lengthy and enlightening discussions.

agent might revise his own posterior when informed of the posterior of the other agent. We shall demonstrate that for an arbitrary event A , under the assumption of common priors, if the information partitions of both agents are finite, then simply by communicating their posteriors back and forth the agents will be led to make revisions that in finitely many steps converge to a common, equilibrium posterior. This posterior we refer to as the indirect communication equilibrium--indirect, since α and β communicate their posteriors on A but not directly the information they received. Furthermore, we shall show that given any positive integer, n , one can construct an example in which the revision process not only takes n steps to converge, but no evident revision occurs--for $(n-1)$ steps both agents repeat the same conflicting posteriors--until the last step when the two agents decide to agree.

Second, observe that, in principle at least, agents α and β could communicate to each other not only their posteriors, but every shred of knowledge, including everything conceivably related to the event A , they possessed. In such a case, i.e., with pooled information (direct communication equilibrium), equality of the two posteriors is trivial, since agents are assumed to have identical priors. The question follows whether the common posterior arrived at in the indirect communication equilibrium coincides with the common posterior arrived at with pooled information. We shall construct an example to demonstrate that this is not necessarily the case--common knowledge of each other's posterior does not necessarily lead agents to the posterior they would have agreed upon had information been directly exchanged. On the other hand, the examples that are characterized

by a discrepancy between the direct and indirect communication equilibria exhibit a high degree of symmetry, which is clearly atypical. We shall conclude by demonstrating that "with probability 1"—in a sense to be made precise—the revision process we construct leads the two agents in one step to the direct communication equilibrium.

In discussing the probability of some event occurring, for instance the election of a particular presidential candidate, individuals are not likely to recount every experience and piece of information they have ever had. What we show is that if two agents have the same priors and if they know the type of information the other is capable of having obtained (they know each other's "information partitions") then simply the terse announcement of probability assessments back and forth will lead to a common judgment; this judgment, however, may not be the same common judgment individuals would make if they shared all their information.

Let (Ω, \mathcal{S}, p) be a probability space, P^α and P^β partitions (known to both agents) of Ω whose join $\frac{1/}{P^\alpha \vee P^\beta}$ consists of non-null events. In the interpretation, (Ω, \mathcal{S}) is the space of states of the world, p is the common prior of α and β , and P^i is the information partition of person i , $i = \alpha, \beta$. If the true state of the world is ω , then i is informed of that element $P^i(\omega)$ of P^i that contains ω . Given $\omega \in \Omega$, an event E is called common knowledge at ω if E includes that member of the meet $\frac{2/}{P^\alpha \wedge P^\beta}$ that contains ω .

Let A be an event, and let $q^i(\cdot)$ denote the posterior probability $p(A|P^i)$ of A given i 's information; i.e., if $\omega \in \Omega$, then $q^i(\omega) = p(A \cap P^i(\omega))/p(P^i(\omega))$.

Proposition (Aumann): Let $\omega \in \Omega$, and let q^a and q^b be numbers. If it is common knowledge at ω that $q^a(\omega) = q^a$ and $q^b(\omega) = q^b$, then $q^a = q^b$.

The concept of common knowledge of the two posteriors can be made explicit as follows: Consider the events

$$E^a(\omega) = \{\omega' \in \Omega \mid q^a(\omega') = q^a\} \quad \text{and}$$

$$E^b(\omega) = \{\omega' \in \Omega \mid q^b(\omega') = q^b\}$$

and let $P(\omega)$ be the member of the meet $P^a \wedge P^b$ which contains ω . To assume that "it is common knowledge at ω that $q^a(\omega) = q^a$ and $q^b(\omega) = q^b$ " is equivalent to the assumption that $E^a(\omega) \supset P(\omega)$ and $E^b(\omega) \supset P(\omega)$. Intuitively then, the posteriors of the two agents on an event A can be common knowledge only if it is the case that the probability of A conditional on any element of P^i contained in $P(\omega)$ is the same. But this is a rather unlikely situation—for most events A it will not be the case. As a result, Aumann's proposition demonstrates the identity of posteriors for a very limited class of events. In the argument that follows we do not impose any restrictions on the event A . Of course, if it is the case that the event A satisfies Aumann's conditions, the equilibrium of our revision process is indeed Aumann's common knowledge equilibrium.

Consider the following example:^{3/} Let

$$\Omega = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}, \quad p(k) = \frac{1}{9} \text{ for all } k \in \Omega;$$

$$P^{\alpha} = \{\{1,2,3\}, \{4,5,6\}, \{7,8,9\}\};$$

$$P^{\beta} = \{\{1,2,3,4\}, \{5,6,7,8\}, \{9\}\};$$

$$P^{\alpha} \vee P^{\beta} = \{\{1,2,3\}, \{4\}, \{5,6\}, \{7,8\}, \{9\}\};$$

$$P^{\alpha} \wedge P^{\beta} = \{1,2,3,4,5,6,7,8,9\};$$

$$u = .1; \text{ hence } \underline{P}^{\alpha}(\omega) = \{1,2,3\} \text{ and } \underline{P}^{\beta}(\omega) = \{1,2,3,4\} .$$

The only events Λ which could be common knowledge are by definition those for which $p(\Lambda \cap E)/p(E)$ is the same for all events E in $P^{\alpha} \vee P^{\beta}$. In the example that leaves only $\Lambda = \Omega$ or $\Lambda = \phi$ and it is clear that in the former case $q^{\alpha}(\omega) = 1 = q^{\beta}(\omega)$ and in the latter case $q^{\alpha}(\omega) = 0 = q^{\beta}(\omega)$. Suppose, on the other hand, that $\Lambda = \{3,4\}$. Then $q^{\alpha}(\omega) = p(\Lambda \cap \underline{P}^{\alpha}(\omega))/p(\underline{P}^{\alpha}(\omega)) = 1/3$ and $q^{\beta}(\omega) = p(\Lambda \cap \underline{P}^{\beta}(\omega))/p(\underline{P}^{\beta}(\omega)) = 1/2$, and these differ.

Now suppose, still taking $\Lambda = \{3,4\}$, that the two agents are allowed to announce and revise their posteriors. Then α might begin by announcing $q_1^{\alpha}(\omega) = 1/3$, which informs β that $\underline{P}^{\alpha}(\omega) = \{1,2,3\}$ or $\{4,5,6\}$. β already knew this, however, and so will still announce $q_1^{\beta}(\omega) = 1/2$. This tells α that β has observed $\underline{P}^{\beta} = \{1,2,3,4\}$ yet does not change α 's opinion (since he already has better information); hence, he will announce again $q_2^{\alpha}(\omega) = 1/3$. At this point β can conclude that $\underline{P}^{\alpha}(\omega) = \{1,2,3\}$ for if $\underline{P}^{\alpha}(\omega)$ were $\{4,5,6\}$ then α would have announced $q_2^{\alpha}(\omega) = 1$. Thus β has gained information and

will himself announce $q_2^\beta(\omega) = 1/3$ and agreement is reached. Note carefully

that for $\lambda = \{3,4\}$ the original $q_1^\alpha(\omega) \neq q_1^\beta(\omega)$ even though β knows $q_1^\alpha(\omega)$ before α announces it (i.e., β can deduce from his own information what $q_1^\alpha(\omega)$ must be) and α knows $q_1^\beta(\omega)$ without an announcement; β does not know, however, that α knows $q_1^\beta(\omega)$; hence there is no common knowledge. Finally, we leave it to the reader to see that if $\lambda = \{1,5,9\}$, then we would get the process

$q_1^\alpha(\omega) = 1/3, q_1^\beta(\omega) = 1/4, q_2^\alpha(\omega) = 1/3, q_2^\beta(\omega) = 1/4, q_3^\alpha(\omega) = 1/3, q_3^\beta(\omega) = 1/3$, in which the agents repeated exactly the same opinions, yet still managed to communicate information. In Proposition 2 we show how to extend the repetitions indefinitely.

More generally, let the two information partitions P^α and P^β be finite; $P^\alpha = \{P_1^\alpha, \dots, P_k^\alpha\}$ and $P^\beta = \{P_1^\beta, \dots, P_l^\beta\}$. Consider the process, τ , of revision of their posteriors by α and β defined stepwise as follows:

Step 1:

$$(1) \quad \alpha \text{ announces } q_1^\alpha(\omega) = \frac{P(P_k^\alpha(\omega) \cap A)}{P(P_k^\alpha(\omega))}$$

$$\text{let } a_1 \text{ be the set } a_1 = \{k \mid \frac{P(P_k^\alpha \cap A)}{P(P_k^\alpha)} = q_1^\alpha(\omega)\}$$

$$(2) \quad \beta \text{ announces } q_1^\beta(\omega) = \frac{P((P_{a_1}^\beta(\omega) \cap \bigcup_{k \in a_1} P_k^\alpha) \cap A)}{P(P_{a_1}^\beta(\omega) \cap \bigcup_{k \in a_1} P_k^\alpha)}$$

$$\text{let } b_1 \text{ be the set } b_1 = \{l \mid \frac{P((P_l^\beta \cap \bigcup_{k \in a_1} P_k^\alpha) \cap A)}{P(P_l^\beta \cap \bigcup_{k \in a_1} P_k^\alpha)} = q_1^\beta(\omega)\}$$

Step t:

$$(3) \quad a \text{ announces } q_t^a(\omega) = \frac{P((P^a(\omega) \cap \bigcup_{i \in \mathcal{I}_{t-1}} P_i^b) \cap A)}{P(P^a(\omega) \cap \bigcup_{i \in \mathcal{I}_{t-1}} P_i^b)}$$

$$\text{let } a_t \text{ be the set } a_t = \{ \omega \mid \frac{P((P^a(\omega) \cap \bigcup_{i \in \mathcal{I}_{t-1}} P_i^b) \cap A)}{P(P^a(\omega) \cap \bigcup_{i \in \mathcal{I}_{t-1}} P_i^b)} = q_t^a(\omega) \}$$

$$(4) \quad b \text{ announces } q_t^b(\omega) = \frac{P((P^b(\omega) \cap \bigcup_{k \in \mathcal{K}_t} P_k^a) \cap A)}{P(P^b(\omega) \cap \bigcup_{k \in \mathcal{K}_t} P_k^a)}$$

$$\text{let } b_t \text{ be the set } b_t = \{ \omega \mid \frac{P((P^b(\omega) \cap \bigcup_{k \in \mathcal{K}_t} P_k^a) \cap A)}{P(P^b(\omega) \cap \bigcup_{k \in \mathcal{K}_t} P_k^a)} = q_t^b(\omega) \}.$$

Note that at each step only probabilities are communicated, yet the fund of common knowledge is steadily increasing:

$$a_1 \supset a_2 \supset \dots \supset a_t \supset a_{t+1} \supset \dots \text{ and}$$

$$b_1 \supset b_2 \supset \dots \supset b_t \supset b_{t+1} \supset \dots$$

We shall now show that if, for some n , $a_{n+1} = a_n$ and $b_{n+1} = b_n$ then it must be the case that $q_{n+1}^{\alpha(\omega)} = q_{n+1}^{\beta(\omega)}$. By definition, for all $k \in a_{n+1}$,

$$(5) \quad p((P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta}) \cap A) = q_{n+1}^{\alpha(\omega)} p(P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta}) .$$

Similarly, for all $i \in b_{n+1}$,

$$(6) \quad p((P_i^{\beta} \cap \bigcup_{k \in a_{n+1}} P_k^{\alpha}) \cap A) = q_{n+1}^{\beta(\omega)} p(P_i^{\beta} \cap \bigcup_{k \in a_{n+1}} P_k^{\alpha}) .$$

Since $a_n = a_{n+1}$ equation (5) holds for all $k \in a_n$, and hence

$$\sum_{k \in a_n} p((P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta}) \cap A) = q_{n+1}^{\alpha(\omega)} \sum_{k \in a_n} p(P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta}) .$$

Furthermore, since

the events in a partition are disjoint,

$$(7) \quad p(\bigcup_{k \in a_n} P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta} \cap A) = q_{n+1}^{\alpha(\omega)} p(\bigcup_{k \in a_n} P_k^{\alpha} \cap \bigcup_{i \in b_n} P_i^{\beta}) .$$

Similarly, equation (6) implies that

$$p(\bigcup_{i \in b_n} P_i^{\beta} \cap \bigcup_{k \in a_n} P_k^{\alpha} \cap A) = q_{n+1}^{\beta(\omega)} p(\bigcup_{i \in b_n} P_i^{\beta} \cap \bigcup_{k \in a_n} P_k^{\alpha}) .$$

But from (7) and (8),

$$q_{n+1}^{\alpha(\omega)} = q_{n+1}^{\beta(\omega)}$$

since the term $p(\bigcup_{i \in b_n} P_i^{\beta} \cap \bigcup_{k \in a_n} P_k^{\alpha})$ is non-zero by the assumption that

$P^{\alpha} \cap P^{\beta}$ consists of non-null events. ^{4/} The argument for the convergence

of π in finitely many steps can now easily be completed. By the preceding argument at step $t + 1$ either $a_{t+1} \subsetneq a_t$ or $b_{t+1} \subsetneq b_t$, unless the equilibrium has been reached with $q_{t+1}^a(\omega) = q_{t+1}^b(\omega)$. Consequently, in at most $K + L$ steps, π must converge.

Proposition 1: Under the assumption that the two information partitions are finite, given any event A , the revision process π converges in finitely many steps. Furthermore, the sum of the cardinalities of the two partitions is an upper bound on the number of steps required for the process to converge.

It is clear from the argument for Proposition 1 that at every step of the revision process the information of at least one of the two persons changes— $a_t \subsetneq a_{t+1}$ or $b_t \subsetneq b_{t+1}$. This does not imply, however, that any of the posteriors must change—it may well be that $q_{t+1}^a(\omega) = q_t^a(\omega)$ and $q_{t+1}^b(\omega) = q_t^b(\omega)$. We say that evident revision occurs if $q_{t+1}^a(\omega) \neq q_t^a(\omega)$ or $q_{t+1}^b(\omega) \neq q_t^b(\omega)$.

Consider the following example:^{5/}

$$\Omega = \{k \in \mathbb{N} \mid 1 \leq k \leq n^2\} ; p(k) = \left(\frac{1}{n^2}\right) \text{ for all } k \in \Omega ;$$

$$\mathcal{P}^a = \{(1, \dots, n), (n+1, \dots, 2n), (2n+1, \dots, 3n), \dots, (n(n-1)+1, \dots, n^2)\};$$

$$\mathcal{P}^b = \{(1, \dots, n+1), (n+2, \dots, 2n+2), (2n+3, \dots, 3n+3), \dots, \\ \{(n-2)(n+1)+1, \dots, (n^2-1)\}, \{n^2\}\};$$

$$A = \{1, n+2, 2n+3, \dots, (n-2)(n+1)+1, n^2\};$$

$$n = 1 .$$

Then $P^{\alpha}(\omega) = \{1, \dots, n\}$ and $P^{\beta}(\omega) = \{1, \dots, n+1\}$. Agent α announces $q_1^{\alpha}(\omega) = 1/n$ while agent β announces $q_1^{\beta}(\omega) = 1/(n+1)$. Furthermore $a_1 = P^{\alpha}$ while $b_1 = P^{\beta} / \{\{n^2\}\}$. At every subsequent step up to the n -th the sets a_t and b_t are reduced by one element but the posteriors $q_t^{\alpha}(\omega)$ and $q_t^{\beta}(\omega)$ remain equal to $q_1^{\alpha}(\omega)$ and $q_1^{\beta}(\omega)$ respectively. At the n -th step, however, agent β "learns" that α has observed $\{1, \dots, n\}$ and, consequently he announces $q_n^{\beta}(\omega) = 1/n$.

Proposition 2: Given any positive integer n , one can construct an example with finite information partitions such that

- (1) The process π takes n steps to converge, and
- (2) No evident revision occurs until the final step. ^{5/}

We now turn to the question whether the indirect communication equilibrium coincides with the direct communication one. Consider the following example:

$$\Omega = \{1, 2, 3, 4\} \quad p(i) = \frac{1}{4} \quad \text{for all } i \in \Omega;$$

$$P^{\alpha} = \{\{1, 2\}, \{3, 4\}\}; \quad P^{\beta} = \{\{1, 3\}, \{2, 4\}\}.$$

Let the true state of the world, ω , be 1, and let A be the event $\{1, 4\}$. Then $P^{\alpha}(\omega) = \{1, 2\}$, $P^{\beta}(\omega) = \{1, 3\}$, $q_1^{\alpha}(\omega) = 1/2$, $q_1^{\beta}(\omega) = 1/2$. Common knowledge of the two posteriors leads to no revision (the process π converges in one step) and the indirect communication equilibrium is attained at $(1/2)$. Suppose, alternatively, that the two agents had exchanged information. Then each of them would be informed of $P(\omega) = P^{\alpha}(\omega) \cap P^{\beta}(\omega) = \{1\}$

and the direct communicational equilibrium would be attained at (1).

Proposition 3: An indirect communication equilibrium is not necessarily a direct communication equilibrium.

Observe, however, that it is the high degree of symmetry that characterizes the example which leads to the divergence between the two types of equilibria. No matter what set in his information partition person i had been informed of, he would have assigned a posterior of $1/2$ to the event A . As a consequence, common knowledge of each other's posterior is of no information value to the two agents. We shall now make precise the intuitive notion that such a situation is atypical: If the event A is chosen randomly, then, with probability 1, the process π converges, for any $\omega \in \Omega$, in one step to the direct communication equilibrium.

Consider the general structure $((\Omega, \mathcal{B}, p))$, $P^\alpha = \{P_1^\alpha, \dots, P_K^\alpha\}$ and $P^\beta = \{P_1^\beta, \dots, P_L^\beta\}$. We assume that p is non-atomic and puts positive weight on every element of the join $P^\alpha \vee P^\beta = \{Q_1, \dots, Q_M\}$; of course, these assumptions still allow for the finite examples we have given simply by taking partitions at uniform intervals. To each event $A \in \mathcal{B}$ we can assign a vector $\sigma(A) = (\sigma_1(A), \dots, \sigma_M(A)) \in I^M = \prod_{j=1}^M [0, 1]$, where

$$\sigma_j(A) = p(A \cap Q_j) / p(Q_j), \quad j = 1, \dots, M.$$

Observe now that a straightforward

application of Liapunov's theorem, since p is non-atomic, implies that the function $\sigma(\cdot): \mathcal{B} \rightarrow I^M$ is surjective. Consequently, if we consider I^M as a measure space endowed with the Lebesgue measure, λ , on the Borel σ -field, there is induced on \mathcal{B} a measurable structure and a measure, μ , as follows: a subset $K \subset \mathcal{B}$ is measurable if and

only if $\sigma^{-1}(G) = K$, where G is a Borel subset of I^M ; $\mu(K) = \lambda(G)$. From the surjectivity of the mapping $\sigma(\cdot)$, the induced measure μ satisfies $\mu(\sigma^{-1}(G)) > 0$ for any open set $G \subset I^M$. Observe now that, the process π will, given λ , converge in one step to the direct communication equilibrium if

- (1) the numbers $p(P_i^a \cap A)/p(P_i^a)$ are distinct for $i = 1, \dots, K$ and
- (2) the numbers $p(Q_j \cap A)/p(Q_j)$ are distinct for $j = 1, \dots, M$.

Conditions (1) and (2) guarantee that each agent reveals all his information with his first announcement. Furthermore, since $p(P_i^a \cap A)/p(P_i^a) = \sum_{j=1}^M p_i^a \sigma_j(A) p(Q_j)/p(P_i^a)$, condition (1) fails for some $A \in \mathcal{B}$ only if at least one of a finite number of linear equations in the variables $\sigma_j(A)$ is satisfied; the same holds for condition (2). But the subset G^* of I^M of points such that none of these linear equations hold has Lebesgue measure 1. It follows then from the definition of the induced measure μ on \mathcal{B} that, if $K^* = \{A \in \mathcal{B}: \pi \text{ converges in one step to the direct communication equilibrium}\}$, $K^* \subset \sigma^{-1}(G^*)$ and hence $\mu(K^*) = 1$ since $1 \geq \mu(K^*) \geq \mu(\sigma^{-1}(G^*)) = \lambda(G^*) = 1$.

Proposition 4: Let the structure $\{(\Omega, \mathcal{B}, p), P^a, P^b\}$ and the function $\sigma(\cdot): \mathcal{B} \rightarrow I^M$ be as above. There exists a σ -algebra of subsets of \mathcal{B} and a probability measure μ on \mathcal{B} such that (1) for any open set $G \subset I^M$, the set $\{A \in \mathcal{B}: \sigma(A) \in G\}$ is μ -measurable and has positive μ -measure; (2) if an event $A \in \mathcal{B}$ is chosen randomly according to μ , then, for any $\omega \in \Omega$, the process π converges in one step to the direct communication equilibrium.

A number of recent works (see Radner [1978] for an overview) have dealt with the problem of the information revealed through the price system in a situation where agents trade under differential information. The analysis has been restricted, however, to situations where the information of one agent (or one group of agents) is "better" than the information of the other; the question asked has then been whether the less informed agent will be able to infer from the equilibrium prices the information received by the other. The situation we have considered is more interesting; the two agents are only assumed to have "different" information. On the other hand, our agents exchange their posteriors—they do not just communicate through the general equilibrium price mechanism. It is a natural question to ask whether—or under what conditions—the result of the existence of a revision process that leads agents to a common posterior is preserved under the additional complication that it is the price mechanism and not the posteriors themselves which is common knowledge.

Footnotes

- 1/ The coarsest common refinement of P^α and P^β .
- 2/ The finest partition refined by both P^α and P^β .
- 3/ Whenever Ω is finite, \mathcal{B} is taken to be the power set of Ω .
- 4/ Observe that it is not the case that if $q_{n+1}^\alpha(\omega) = q_{n+1}^\beta(\omega)$ then $a_{n+1} = a_n$
and $b_{n+1} = b_n$.
- 5/ This example was provided by Robert Aumann in exchange for a Kosher dinner.
- 6/ The archetypal example of this phenomenon goes something like this. There are n students in a classroom, all wearing red hats. Each can see what the others are wearing but does not know the color of his own hat. There is also a rule that anyone can go home if he knows his hat is red, but communication is not allowed nor can anyone look at his own hat. Clearly nobody would ever leave, yet if one day the teacher announced that there was at least one person in the room wearing a red hat, a fact already known to every student, nothing would happen for $n - 1$ more days but on the n th day everyone would get up and go home. We give the explanation for $n = 3$. On the first day after the teacher's announcement, α , β , and γ sit still since each see others with red hats. On the second day α might still be unsure about the color of his hat. But he knows that if it were black then β on day 1, β must have been looking at γ 's red hat and so stayed put and γ must have been looking at β 's red hat and so stayed put. But then on day 2 α figures that β would leave if α 's hat were not red. For if α did not have a red hat then β would know that γ could not have been looking at α 's hat the first day but at β 's. Hence on the second day β would have left. Since β did not leave on the second day α knows that his own hat is red and he leaves on the third day. In summary, it is not until after the second day that α knows that β knows what γ knows and hence not until the third day everyone goes home.

References

- Aumann, R.J. [1976], "Agreeing to Disagree," The Annals of Statistics, Vol. 4, No. 6, pp. 1236-1239.
- Radner, R. [1978], "Rational Expectations Equilibrium with Differential Information," Technical Report No. OW-12, Center for Research in Management Science, University of California, Berkeley.