

A Comment on Aumann's Bayesian View[†]

Faruk Gul
Princeton University

February 1997

In Aumann [1987], it is asserted that for those who adhere to the "... Bayesian view of the world, the notion of equilibrium is an unavoidable consequence..." I discuss two possible interpretations of the information model and show that neither interpretation supports this assertion. The *hierarchy representation* interpretation renders the prior stage meaningless and hence both the key assumption of Aumann's theory and its conclusion become impossible to interpret. The *prior* interpretation, on the other hand, is distinctively non-Bayesian. Furthermore, both the common prior assumption and the notion of having beliefs over one's own actions are problematic in the latter interpretation.

Key Words: correlated equilibrium, Bayesian view of probability, common prior assumption, Savage's foundations of statistics.

[†] I am indebted to Eddie Dekel-Tabak, David Kreps, and Robert Wilson for their help. Financial support from the Alfred P. Sloan Foundation and the National Science Foundation is gratefully acknowledged.

Aumann [1987] presents a result which he interprets as establishing that, for those who adhere to the “... Bayesian view of the world, the notion of equilibrium is an unavoidable consequence...”¹ The purpose of this comment is to refute this claim.

An information model $I = \{\Omega, (T_i, p_i)_{i=1}^n, \tilde{x}\}$ consists of a set of states of nature Ω , partitions T_i of Ω , prior probability distributions p_i on Ω ,² and a function $\tilde{x} : \Omega \rightarrow X$, where X is the set of relevant parameter values for the underlying problem. Let $t_i(s)$ denote $t_i \in T_i$ such that $s \in t_i$.

Aumann [1987] considers an information model in which $X = A \times Y$, where $A = \prod_{i=1}^n A_i$ is the set of pure strategy profiles for some finite normal form game of complete information, $G = \{(A_i, u_i)_{i=1}^n\}$, and Y is the set of all other relevant parameters. Let \tilde{x}^n denote the first n coordinates of \tilde{x} . Aumann makes the following assumptions:

Assumption 1 (Common Priors): $p = p_i$ for all i .

Assumption 2 (Bayesian Rationality): For all i , $s \in \Omega$, and $z_i \in A_i$,

$$Eu_i(\tilde{x}^n | t_i(s)) \geq Eu_i(z_i, \tilde{x}_{-i}^n | t_i(s)).$$

Theorem (Aumann [1987]): If \tilde{x}_i^n is measurable with respect to T_i for all i (i.e., each player knows his own actions) and Assumptions 1 and 2 are satisfied, then the distribution of \tilde{x}^n is a correlated equilibrium distribution.

Aumann’s defense of his assumptions includes the following claims:

- (1) The statement, “ I is common knowledge” is a tautology. Hence all players know I .³
- (2) The common prior assumption reflects a “reasonable” philosophical position, and nearly all of information economics is predicated on this assumption: “...[P]eople with different information may legitimately entertain different probabilities but there is no rational basis for people who have always been fed the same information to do so.”⁴

¹ Aumann [1987], p. 2.

² For the sake of simplicity, consider only finite Ω ’s. The papers cited below by Mertens and Zamir [1985] and Brandenburger and Dekel [1993] do not guarantee the finiteness of Ω .

³ In the discussion of the prior view below, I challenge even the weaker assertion that players know I (i.e., each other’s beliefs, for example), not just the common knowledge assumption.

⁴ Aumann [1987], pp. 13–14.

There are two possible ways to interpret the information model. In the first of these, we imagine a situation prior to the economic problem to be considered. We assume that the p_i 's described the beliefs of players and were common knowledge at this prior stage. At a subsequent stage, a state of nature is realized, agents receive their information $t_i(s)$, update their priors, and make the appropriate adjustments in their beliefs and behavior (as embodied in the function \tilde{x}). In this interpretation, the prior stage represents a situation which actually occurred at some previous time. The distribution of the random variable \tilde{x} implied by any one of the existing prior distributions p_i is a meaningful vehicle for stating predictions or making welfare comparisons.

For example, in this interpretation, statements of the form “*ex ante* everyone would be better off if...” or “the marginal prior distribution of agent i 's beliefs on the set of action profiles is a correlated distribution” are statements that are independent of the particular representation of the underlying problem. In this interpretation, all states of nature describe contingencies that at a prior stage were considered as possible resolutions of the uncertainty. I will call this interpretation, the *prior* view. It is evident that much of information economics is predicated on this view.

The prior view outlined above is an inherently dynamic story. It forces us to think of a prior stage. Note that *if* a prior stage exists, it is certainly not a tautology (but often boldly assumed in information economics models) that agents know each other's beliefs and how these beliefs will be revised. When priors include beliefs about one's own action as in Aumann's Theorem, the knowledge and common prior assumptions become even more onerous: it is particularly difficult to imagine a moment in time where players 1 and 2 have the same beliefs about player 1's future actions. Moreover, the assumption of common priors in this context is antithetical to the Savage-established foundations of statistics (i.e., the “Bayesian view”), since it amounts to asserting that at some moment in time everyone must have identical beliefs.⁵

⁵ “The criteria incorporated in the personalistic view do not guarantee agreement on all questions among all honest and freely communicating people, even in principle. That incompleteness, if one will call it such, does not distress me, for I think that at least some of the disagreement we see around us is due neither to dishonesty, to errors in reasoning, nor to friction in communication, though the harmful effects of the latter are almost incapable of exaggeration.” Savage [1954] p. 67. Here, Savage appears to be rejecting not only the idea of “neutral” common probabilities at any stage, but also the idea that “difference in probabilities should reflect differences in information only”. I am grateful to an anonymous referee for this citation. Compare with the citation from Aumann in (2) above. For another analysis of Savage's views and the common prior assumption, see Morris [1995].

The second interpretation of the model (which I will call the *hierarchy representation* interpretation is based on a result due to Mertens and Zamir [1985] and Brandenburger and Dekel [1993]. These two papers show that given any collection of infinite hierarchies of beliefs (one for each player), one can construct an information model $I = \{\Omega, (T_i, p_i)_{i=1}^n, \tilde{x}\}$ such that at the true state of nature $s^* \in \Omega$, the posterior of each player i represents the same beliefs as the infinite hierarchy of beliefs of player i over X . Thus, in this interpretation, an information model is simply a notational device for representing the n -tuple of infinite hierarchies of beliefs over X . To fix ideas, consider the following simple pair of infinite hierarchies of beliefs:

Let $X = \{x, y\}$ denote the set of parameters. In particular assume that x, y are the two strategy choices of player 1. Let $[x], [y]$ denote the “events” x is realized and y is realized respectively. (Note that (1a)-(1d) define player 1’s and (2a)-(2d) define player 2’s hierarchy.)

(1a) Player believes with probability 1 that $[x]$ (henceforth, I will use BFC, believes with certainty, as short-hand for “believes with probability 1”). (1b) 1 BFC that 2 BFC that if $[x]$ then 1 BFC $[x]$ and if $[y]$ then 1 BFC $[y]$. (1c) 1 assigns probability .5 to the events E_1 and E_2 where $E_1 = \{2 \text{ assigns probability } .4 \text{ to } [x] \text{ and } .6 \text{ to } [y]\}$ and $E_2 = \{2 \text{ assigns probability } .5 \text{ to } [x] \text{ and } .5 \text{ to } [y]\}$. (1d) 1 BFC that 2 BFC (1b) and (1c), 1 BFC that 2 BFC that 1 BFC that 2 BFC (1b) and (1c) and so forth.

(2a) 2 assigns probability .4 to $[x]$ and .6 to $[y]$ (2b) 2 BFC that if $[x]$, 1 BFC $[x]$, and if $[y]$, 1 BFC $[y]$. (2c) 2 BFC that 1 assigns probability .5 to E_1 and .5 to E_2 (2d) 2 BFC that 1 BFC (2b) and (2c), 2 BFC that 1 BFC that 2 BFC that 1 BFC (2b) and (2c) and so forth.

The following information model represents these two infinite hierarchies:

$I = \{\Omega, (T_i, p_i)_{i=1}^n, \tilde{x}\}$ where $\Omega = \{s_1, s_2, s_3, s_4\}$, $T_1 = \{\{s_1, s_2\}, \{s_3, s_4\}\}$, $T_2 = \{\{s_1, s_3\}, \{s_2, s_4\}\}$, $\tilde{x}(s_1) = \tilde{x}(s_2) = x$, $\tilde{x}(s_3) = \tilde{x}(s_4) = y$, $p_1(s_1) = p_1(s_2) = .5\alpha$, $p_1(s_3) = p_1(s_4) = .5\beta$, $p_2(s_1) = .4\delta$, $p_2(s_3) = .6\delta$, $p_2(s_2) = p_2(s_4) = .5\gamma$, for any $\alpha, \beta, \delta, \gamma > 0$ such that $\alpha + \beta = \delta + \gamma = 1$ and $s^* = s_1$ is the true state of nature.

Aumann’s remarks regarding the common knowledge assumptions being tautologies are valid in this context. That is, any conclusion a player can draw from I and his true

type, he can also draw from his infinite hierarchy of beliefs over X . However, there is no prior stage and hence no possible interpretation of the distribution of \tilde{x} with respect to any one of the priors. There is also no way to talk about one player being more or less informed than another or an outside observer since such comparisons necessitate a dynamic framework in which information is actually acquired.⁶ The priors are artifacts of a notational device to represent the infinite hierarchies of beliefs on X of the players, i.e., their “posteriors” at the true state of nature.⁷ The states of Ω other than s_1 never represented contingencies that were considered possible as realizations of the uncertainty; there was never a moment in time at which either player thought s_2 could have occurred. In fact, the occurrence of the state s_2 or the idea of “giving” player 2 the information $\{s_2\}$ does not correspond to meaningful thought experiment. “Conditioning” p_2 on $\{s_2\}$ can not be viewed as giving the information $\{s_2\}$ to player 2. To see this, note that we can not change player 2’s partition without changing the infinite hierarchy of beliefs for 1 as well⁸.

The fallacy of viewing the “priors” in the above representation as beliefs at some hypothetical prior stage becomes clear, if we remember that the information model above is nothing more than an equivalent representation of the infinite hierarchies of beliefs (i.e.,(1a)-(1c) and (2a)-(2c)). Yet there is no data in either of these hierarchies regarding any “prior” stage. Nor is there any data in the infinite hierarchies regarding what the agents *would have believed* had their information been “less” or “more” than what it in fact is. Hence, no such data can be present in the equivalent representation of these hierarchies of beliefs. The fact that the infinite hierarchies of beliefs contain no data about any prior stage can also be verified by observing that $\alpha = p_1(\{s_1, s_2\})$ is indeterminate (i.e., the infinite hierarchy of beliefs of a player does not pin down his own prior for his own types). Thus, it is clear that insisting on $p_1(\{s_1, s_2\}) = p_2(\{s_1, s_2\})$ has nothing to do with the

⁶ Aumann in his reply, proposes that if, in the model representing the infinite hierarchies of beliefs, 1’s partition element is fully contained in 2’s, then 1 is better informed at any state in this partition element of 1. He states that this test “nicely captures the intuitive idea of being more informed”. However, the hierarchy representation view offers no means of determining if the greater confidence of player 1, implied by his finer information partition represents better information, unjustified over-confidence or a legitimate disagreement between reasonable people. One could resolve the issue by *asserting* that 1’s belief is in fact knowledge (i.e. implies truth). The usefulness of this option in developing a subjectivist Bayesian view of the world is, however, questionable.

⁷ Brandenburger and Dekel [1993] explicitly avoid any reference to the prior in their representation theorem mentioned above.

⁸ Note Axiom 2 in the reply.

view that “differences in probabilities reflect differences in information only” even if we could understand what this view means within the hierarchy representation interpretation.

Finally, it can be shown that no information model that represents the infinite hierarchy of beliefs for 1 defined by (1a)-(1c) can have a common prior, irrespective of the infinite hierarchy of beliefs of 2.⁹ How can it be that insisting on a common prior reflects the view that “differences in probabilities reflect differences in information only” if the beliefs of player 1 *alone* without any mention of player 2’s beliefs and presumably his information, can preclude a common prior? Since Ω and the priors are not primitives, ultimately, the meaning of any assumptions or conclusions regarding them (i.e., statements such as “the priors are common” or “the prior distribution of actions is a correlated equilibrium distribution”) can be understood only with respect to their implications on the posteriors of the players at the true state of nature (i.e., the infinite hierarchies of beliefs over X). The exact implications of the common prior assumption on the infinite hierarchy of beliefs are not known.¹⁰ Thus, within the hierarchy representation interpretation, we do not know what it is that we would be accepting if we were to accept the common prior assumption. Arguing for common priors in the hierarchy representation interpretation on the grounds that people with the same level of information should have the same beliefs is analogous to arguing for risk aversion on the grounds that the incremental units of money give less and less satisfaction. It amounts to an abuse of the formal similarity of two objects (prior and hierarchy representation interpretations of the information model in one case and the von Neumann-Morgenstern utility function versus some cardinal hedonic measure in the other).

Of the two possible interpretations of the information model, only the hierarchy representation interpretation supports the claim that the information model itself is common knowledge among the players. However, this interpretation renders the prior stage meaningless (i.e., it becomes impossible to associate the prior stage with a sensible thought experiment) and hence both the common prior assumption and the conclusion of Aumann’s theorem, that the prior distribution of action profiles is a correlated equilibrium,

⁹ To see this observe that (1b)-(1d) imply E_1 and $[x]$ must be p_1 independent. Obviously, E_1 and $[x]$ can not be p_2 independent.

¹⁰ I am grateful to an anonymous referee for pointing this out.

become uninterpretable. Moreover, the hierarchy representation interpretation offers no argument for identifying the “priors” of the representation with beliefs at any prior stage. Finally, insisting that priors be common does not correspond to anything we might mean by the phrase “differences in probabilities reflect differences in information only”, but rather constitutes a complex and unintuitive restriction on *each* infinite hierarchy. Even if we were to impose common priors, this would not render the prior stage relevant, nor would it render the prior distribution meaningful, nor would it justify identifying this prior with the beliefs of a heretofore unmentioned outside observer within the hierarchy representation interpretation. If an outside observer existed, we would have to model her beliefs the same way we model the players’ beliefs. It can be shown that if the outside observer were modeled as an additional player, the common prior assumption alone would not enable us to conclude that his beliefs over the behavior of the actual players must constitute a correlated equilibrium. In fact, if the underlying game were a 2×2 game, *any* belief for the outside observer that placed zero probability on non-rationalizable actions would be consistent with all of Aumann’s assumptions.

On the other hand, with the prior view, it is not a tautology that players know each others priors. With this view, the assumption that players have priors over their own actions is problematic. Assuming that these priors are commonly known and identical is all the more so. Literally, with the prior view, Aumann’s theorem amounts to noting that if actions are generated by some randomization device, and if every action in the support of the prior distribution is conditionally optimal (i.e., conditional on that action being played), then the prior distribution of the action profiles is a correlated equilibrium (i.e., simply a restatement of the definition of correlated equilibrium). Neither the theorem nor the analysis provides any new argument for why one would expect actions to be generated in this way nor why beliefs for the players (or an outside observer) should be as if behavior were generated in this manner.

Ultimately, Aumann’s defense of the assumptions of his theorem boils down to a peculiar philosophical position. He states, in his reply that assumptions (or axioms) are to be intuitive but not compelling, that we do not “believe” assumptions, that they are not articles of faith and that in the end they will be judged by where they lead rather than

considerations of innate plausibility. This is a strange defense of an axiomatization that claims to establish that for those who adhere to the "...Bayesian view of the world, the notion of equilibrium is an unavoidable consequence...".

References

- Aumann, R., (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* 55: pp. 1–18.
- Aumann, R. and Brandenburger, A., (1995) "Epistemic Conditions for Nash Equilibrium," *Econometrica* 63: pp. 1161–1180.
- Brandenburger, A. and Dekel, E., (1993) "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory* 59: pp. 189–198.
- Mertens, J.F. and Zamir, S., (1985) "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory* 19: pp. 1–29.
- Morris, S., (1995) "The Common Prior Assumption in Economic Theory," *Economics and Philosophy* 11: pp. 1–27.
- Savage, L., (1954) *The Foundations of Statistics*, New York: Wiley.