



Common Priors: A Reply to Gul

Robert J. Aumann

Econometrica, Vol. 66, No. 4 (Jul., 1998), 929-938.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28199807%2966%3A4%3C929%3ACPARTG%3E2.0.CO%3B2-I>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

COMMON PRIORS: A REPLY TO GUL

BY ROBERT J. AUMANN¹

1. INTRODUCTION

GUL (1998) (henceforth simply Gul) attacks the conceptual foundation of the Common Prior Assumption (CPA), which is essential for the derivation of correlated equilibrium in Aumann (1987) (henceforth [A]). We are grateful for the opportunity to discuss the CPA and the underlying information model in more detail; evidently, our 1987 discussion was too sketchy.

To avoid misunderstanding, we stress that we do not consider the CPA “true;” the concept of truth does not apply here. We do think that it embodies a reasonable and useful approach to interactive decision problems, though by no means the only such approach; and that it does not suffer from the deficiencies attributed to it by Professor Gul.²

Section 2 is devoted to technical preliminaries, and can perhaps be omitted or only briefly scanned at a first reading. In Sections 3 through 6, we flesh out an argument for the CPA that was presented sketchily in [A]. This is then used in Sections 7 and 8 to respond to Gul’s points.

2. PARTITION STRUCTURES AND HIERARCHIES

There are several representations of information in n -person environments. One consists of a space X of parameter values, a finite³ space Ω of “states of the world,” a function $x: \Omega \rightarrow X$, and for each of the n players i , an information partition of Ω and a probability distribution $p(\cdot; i, P)$ on each atom P of his partition; call this a *partition structure*. In such a structure, the meet⁴ of the information partitions is called the *common knowledge (ck) partition*; its atoms are the *ck components* of the structure. A *connected*⁵ structure is one with a single ck component; since at any state ω , the “true”

¹Research support from NSF Grant SES-9209277, and from the Game Theory Program at SUNY-Stony Brook, is gratefully acknowledged. Also, we heartily thank Professor Dov Samet for several important discussions regarding this material. Finally, we wish to express our gratitude, appreciation, and admiration to the co-editor of *Econometrica* in charge of this note, and to two anonymous referees. Over a period of years, they have invested an unusually large amount of time, effort, and energy in the processing of this note, far beyond the call of duty, and so have played a large role in shaping its current form.

²Gul starts by disputing our claim, in the introduction to [A], that “the notion of equilibrium is an unavoidable consequence of (Bayesianism).” We readily agree that (objective) correlated equilibrium is not literally an *unavoidable* consequence of Bayesianism alone; one also needs the CPA. This is made abundantly clear at several places in [A] (Sections 3, 4b, and 5).

Even with the CPA, Bayesianism itself is not enough for correlated equilibrium; one needs *common knowledge* of Bayesianism, which is a different kettle of fish altogether. We are sorry if anyone was misled.

³The finiteness restriction is not without loss of generality, as noted both in [A] and by Gul.

⁴Finest common coarsening.

⁵With a given partition structure \mathbf{P} , one can associate a graph G , whose vertices are the states ω , and whose edges are pairs of states lying in the same information set for some player. Then \mathbf{P} is connected iff G is.

ck component is commonly known by the players, one may for many purposes limit discussion to connected structures.

Information may also be represented by an n -tuple of knowledge-belief hierarchies like⁶ those in Harsanyi (1967–68), Mertens–Zamir (1985), and others; call this a *hierarchy n -tuple*.

These two representations of information (by partitions and by knowledge-belief hierarchies) are closely related; they are two ways of looking at the same thing, two sides of the same coin. Thus, each state ω in a partition structure \mathbf{P} induces a unique hierarchy n -tuple in a natural way; we call it the *hierarchy n -tuple at ω* . Conversely,⁷ given a hierarchy n -tuple h , one may find a connected partition structure \mathbf{P} with a unique state ω at which h is the hierarchy n -tuple. If, moreover, one asks that \mathbf{P} be *reduced*—i.e., that each two different states in \mathbf{P} induce different hierarchy n -tuples—then \mathbf{P} is uniquely determined.⁸

Partition structures look much like the model of [A] and Gul’s “information model,” but there is a noteworthy difference. There each player i starts out with a prior on all of Ω ; his posterior probabilities given his information can then be *derived*, using ordinary conditional probability calculations. Here, we *start* with the posteriors. Mathematically, the two kinds of model are equivalent; each can easily be derived from the other. Here we have chosen to avoid defining the basic model in terms of a prior on the entire state space, in order to emphasize our primary concern with the players’ actual information and probability assessments—not with how the information and assessments were derived, and not with the situation at any previous time.

3. THE INTUITION FOR COMMON PRIORS

A *common prior* for a partition structure \mathbf{P} is defined as a probability measure q on Ω such that $p(E; i, P)q(P) = q(E \cap P)$ whenever E is an *event* (subset of Ω), i a player, and P an atom of i ’s information partition. In words, each player’s probability is derived from q by conditioning on the information in his partition. The *common prior assumption* (for \mathbf{P}) says that \mathbf{P} is endowed with a common prior.

The basic intuition for the CPA, prominently cited by Gul, was set forth in [A, pp. 13–14] as follows: “... the CPA expresses the view that probabilities should be based on

⁶But not identical with them. They deal with beliefs—probabilities—only; we are interested also in knowledge, in the sense of absolute certainty (as distinguished from probability 1). Already in the two-player case, a pair of knowledge-belief hierarchies is an awesomely complex object. At the first level of the hierarchy, each player has a knowledge-belief profile about X ; that is, he knows that the true parameter value is in a certain subset B of X , and he has a probability distribution on B . The knowledge of the two players must be *consistent*; one player cannot know something that the other knows to be false. At the second level of the hierarchy, each player has some knowledge and beliefs about pairs consisting of elements of X and the other players’ first-level knowledge and beliefs; these second-level knowledge-belief profiles of the two players must be consistent with each other in a sense like that described for the first level, and each player’s second stage profile must also be consistent with his first level profile (e.g., the first level must be the marginal of the second level when projected onto the first level). And so on, ad infinitum.

⁷See, for example, Mertens and Zamir (1985)—though we stress again that their set-up is a little different, in that they do not deal with absolute certainty (as distinguished from probability 1 belief).

⁸Indeed, not only each hierarchy n -tuple, but even just the hierarchy of one player alone, uniquely determines the associated partition structure. Since a player may always be presumed to know his own knowledge-belief hierarchy, it follows that he also knows precisely the relevant partition structure.

information; that people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so. . . . Under the CPA, differences in probabilities express differences in information *only*.”

A decade later, we must admit that we could have been more explicit here. After all, the players do *not* in general have the same information; why is the argument relevant for them?

To answer this, let us imagine that the players forget all their differential information; formally, this is represented by a partition structure in which each of the partitions is replaced by the trivial partition (consisting of Ω only). In this new situation, the players *do* have precisely the same information. So we *can* apply the above argument from [A], and conclude that then they all *do* have the same probability; denote it q .

Now, suppose the players are reminded of what they forgot; formally, return to the original partitions. Being reminded is a way of acquiring information; so each player's current probabilities may be derived from q by conditioning on his information. Since we are now back where we started, it is reasonable to suppose that also the probabilities are the same as those with which we started; that is, they are all derived from q by conditioning on the private information. But this means that q is a common prior.

This argument is informal. *If* we wish to formalize it, we must expand the framework to a dynamic one, in which a player's information can change. And, we must formulate axioms that embody the underlying principles that

- (1) players with the same information have the same probabilities; and
- (2) when information is acquired, probabilities are updated by the usual rules of conditioning.

Do we wish to formalize it? We think yes. In conceptual discussion, it is easier to come to grips with an argument—to identify the essential points, to separate the difficulties—if it is formal. Nevertheless, the “theorem” below is nothing more than a precise formulation of the above informal argument from [A].

It appears that our main difference with Gul revolves around the dynamic framework. He views the given informational situation as static, and considers it inappropriate to ask, “where did we come from, what might we have known before?” Imbedding the given set-up in a larger one with hypothetical elements bothers Gul, even though such imbeddings are standard in axiomatic analyses (see Section 8a).

Note that the argument in this section depends basically on the substantive notion of “acquiring information” rather than on this or that formalism.

4. FORMAL TREATMENT

Define a *dynamic framework* as a family \mathfrak{F} of partition structures \mathbf{P} , with the same state space Ω and the same player set, that is closed under coarsening; i.e.,

(0) *If \mathcal{P} is the partition profile⁹ in some structure \mathbf{P} in the family \mathfrak{F} , and \mathcal{Q} is obtained by coarsening¹⁰ one or more of the partitions in \mathcal{P} , then there is a structure \mathbf{Q} in \mathfrak{F} with partition profile \mathcal{Q} .*

The structures in a dynamic framework \mathfrak{F} represent different situations that may arise as time progresses and the players acquire more and more information. Each \mathbf{P} in \mathfrak{F} is

⁹ n -tuple of information partitions.

¹⁰ One partition *coarsens* another if each atom of the coarser partition is a union of atoms of the finer one.

viewed as resulting from some history; we ask that \mathfrak{S} contain all possible structures that may have preceded \mathbf{P} in this history. Call \mathfrak{S} *consistent* if for all \mathbf{P} and \mathbf{Q} in \mathfrak{S} , events E , and players i and j , we have

(1) *If the same event P is an atom of the partitions of both i and j in \mathbf{P} , then $p_{\mathbf{P}}(\cdot; i, P) = p_{\mathbf{P}}(\cdot; j, P)$, where $p_{\mathbf{P}}$ denotes the probabilities in \mathbf{P} ; and*

(2) *Suppose that i 's partition in \mathbf{P} refines his partition in \mathbf{Q} , while each other player's partition is the same in \mathbf{P} as in \mathbf{Q} . Let P be an atom of i 's partition in \mathbf{P} , and Q that atom of i 's partition in \mathbf{Q} that includes P . Then $p_{\mathbf{P}}(E; i, P)p_{\mathbf{Q}}(P; i, Q) = p_{\mathbf{Q}}(E \cap P; i, Q)$, and $p_{\mathbf{P}}(\cdot; j, \cdot) = p_{\mathbf{Q}}(\cdot; j, \cdot)$ for $j \neq i$.*

These two axioms embody the two numbered "principles" in Section 3. Note that (1) yields

(1a) *If each player's partition has just one atom—namely, all of Ω —then there is a single probability q on Ω that is the same for all players and represents all their probabilities.*

This means that if the players know "nothing," they must have the same probabilities.

THEOREM: *In a consistent dynamic framework, each partition structure \mathbf{P} has a common prior.*

PROOF: For simplicity, take $n = 2$; the case of general n is similar. Call the players 1 and 2.

Starting from the partition profile $\mathcal{P} = (\mathcal{P}^1, \mathcal{P}^2)$ in \mathbf{P} , construct two auxiliary profiles \mathcal{C} and \mathcal{R} . Define \mathcal{C} as the trivial profile, in which \mathcal{C}^1 and \mathcal{C}^2 have only one atom each, namely all of Ω . Define \mathcal{R} by $\mathcal{R}^1 := \mathcal{P}^1$ and $\mathcal{R}^2 := \mathcal{C}^2$. Thus \mathcal{C} arises from \mathcal{P} when both players "forget" their private information; \mathcal{R} from \mathcal{C} , when 1 is then "reminded" of her information, but 2 is not. One may also think of \mathcal{P} as arising from \mathcal{R} when 2 is then reminded of *his* information.

By (0), there are partition structures \mathbf{Q} and \mathbf{R} in \mathfrak{S} whose partition profiles are \mathcal{C} and \mathcal{R} respectively. By (1a), there is a single probability distribution q on Ω that represents both players' probabilities in \mathbf{Q} . Applying (2) twice, first to go from \mathbf{Q} to \mathbf{R} , then from \mathbf{R} to \mathbf{P} , we find that the probabilities in \mathbf{P} are the conditionals of q given the information of each player in \mathbf{P} . Thus \mathbf{P} has the common prior q , so the proof is complete.

5. DISCUSSION

It is worthwhile to be a little more explicit about the motivation for our treatment.

As we have said, the CPA expresses the idea that differences in probabilities should reflect differences in information *only*. If one sets forth all relevant information in sufficient detail, then in principle, there should be no room for differing probabilities. When we say all relevant information, we mean *all*: the schools the players attended, their childhood experiences, even their genes (which indirectly reflect the experience of previous generations). Of course, the players do not need to know all this information; indeed, usually they know very little. What we are saying is that if people have precisely the *same* information about all these factors—no matter how much or how little it is—then it is not unreasonable to assume that they entertain the same beliefs.

Admittedly, in concrete instances it is far beyond anybody's power to write down an explicit model of this kind. Yet solid, down-to-earth conclusions are routinely drawn from

theoretical models whose complexity—or simply sheer size—defies the imagination. General equilibrium theory requires that economic agents have complete preferences over commodity spaces whose *dimension* may run into the tens of thousands. The first theorem in game theory (Zermelo (1913)) says that chess has a pure strategy saddle point; it cannot even be *stated*—much less proved—without considering strategies that are beyond anyone’s power to describe or even imagine in any concrete sense. Using Zermelo’s theorem, one can show that in the game of Hex, the *first* player has a winning strategy (Gale (1979)); nobody knows what it is, and probably it cannot be described in any practical sense. The infinite hierarchies of beliefs to which Gul refers are objects of astounding complexity, which it is utterly unrealistic to ascribe to human agents. One could adduce more and more examples.

In all these cases, the theoretical models have important, widely recognized implications. We suggest that here too, it is perfectly legitimate to see where we are led when we allow the states to include *all* relevant information, without providing details.

6. A QUESTION ABOUT AXIOM (2)

Though the partitions of players other than i are the same in \mathbf{P} and \mathbf{Q} , nevertheless it does not seem that only i ’s information is changed in going from \mathbf{Q} to \mathbf{P} . Let \mathcal{P}^i and \mathcal{Q}^i be i ’s partitions in \mathbf{P} and \mathbf{Q} respectively. Suppose the true state ω of the world lies in an atom of \mathcal{Q}^i that is nontrivially partitioned by \mathcal{P}^i . Let j be a player other than i . In \mathbf{Q} , Player i does not know the element of \mathcal{P}^i in which ω lies, and j knows that i does not know this. In \mathbf{P} , Player i does know the element of \mathcal{P}^i in which ω lies, and j knows that i knows this. Thus in the passage from \mathbf{Q} to \mathbf{P} , both i and j have learned something; i has learned the element of \mathcal{P}^i , and j has learned that i has learned the element of \mathcal{P}^i . Indeed, in \mathbf{P} it is common knowledge that i knows the element of \mathcal{P}^i , so i has learned more than just the element of \mathcal{P}^i ; she has, for example, learned that j has learned that she has learned the element of \mathcal{P}^i . Under these circumstances, is it still true that in going from \mathbf{Q} to \mathbf{P} , Player j should not change his probabilities, and i should simply use the ordinary rules of conditioning?

To answer this, we must look more closely at the meaning of the partition structures \mathbf{P} and \mathbf{Q} . Imagine that the players are called into a room, and in each other’s presence are told Ω , \mathbf{P} , and \mathbf{Q} , that upon leaving the room, i will be told the true element of \mathcal{Q}^i , each other player will be told the true element of his partition, that i will be told the true element of \mathcal{P}^i at a specified time t in the future, and that nobody will find out anything else.¹¹ Then \mathbf{Q} and \mathbf{P} represent the situations before and after t respectively. Clearly i learns something new at time t , but does any other player j learn anything that he did not already know?

We claim that he does not. It is true that j gets to know at time t that i knows the element of \mathcal{P}^i , but he knew beforehand that he would get to know this. It is as if j is taking a train from Basel to Zürich. As the train pulls out of the Basel station, he knows that he is in Basel, and later he knows that he is in Zürich. But that is not “learning.” He knew already in Basel that he would get to Zürich; that he indeed did so should not make him revise his probabilities. Similarly, i *knew* that j will know that she (i) will know the element of \mathcal{P}^i , so what she really learns at time t is *only* the actual element of \mathcal{P}^i . Therefore, i should indeed apply the ordinary rules of conditioning in passing from \mathbf{Q} to \mathbf{P} .

¹¹The story about the room is for vividness only; it should not be taken too literally.

Differently put: the elements of Ω describe the situation in \mathbf{P} —that after time t . The partitions in \mathbf{Q} describe what the players know before t about what they will know after t . Clearly, then, only i learns anything at time t ; what she learns is *only* the true element of \mathcal{P}^i .

7. GUL'S "PRIOR INTERPRETATION"

This interpretation of Gul's "information model" (essentially our "partition structure") involves two stages: "prior" and "current." The given partition structure represents the current stage. At the prior stage—which is meant to model "a situation that actually occurred at some previous time"—there is no differential information; all players consider all elements of Ω possible. As Gul points out, this implies that the agents know each others' beliefs. He says that "much of information economics is predicated on this view" and it "is often boldly assumed," but is "certainly not a tautology." It appears from this that he considers his "prior interpretation" as dubious, even without common priors.

We agree. Though it *may* happen that there is an "actual" prior stage at which beliefs are commonly known, this is the exception rather than the rule. If there is, then the analysis in the foregoing sections applies, and we conclude that at that stage the beliefs must be the same—i.e., we have a common prior. If, as seems more likely, the beliefs at an "actual" prior stage are different and *not* commonly known, then there must be differential information already at that stage; and then again, the foregoing analysis leads us to a common prior at a hypothetical stage that precedes the "prior stage," and so also at the "current" stage. The only case that *would* be inconsistent with the CPA is if the beliefs at the prior stage were commonly known and different; but Gul has adduced no evidence or even argument that such a situation is tenable.¹²

¹²The passage from Savage (1954) cited by Gul in his footnote 6 must be read in context. As his book's title indicates, Savage's primary concern was not decision theory but statistics; specifically, statistics as a tool for scientific induction. Here Savage sensed a difficulty. His theory was personal. It depended on the individual; in principle, *based on the same evidence*, different individuals could legitimately reach different conclusions on, say, the temperature of the sun. This personal approach does not seem to jibe with generally accepted views of science. Many of us view science as an attempt to get at the truth, and that seems unrelated to our personal preferences for tea or coffee.

The whole of Savage's Section 4.6 (from which the passage is quoted) is devoted to grappling with this problem, to trying to bridge this gap. On the one hand, Savage does *not*, in fact, include anything like the CPA or our Axiom (1)—his formal apparatus really is purely personal. On the other, he did consider this a problem, and was looking for ways around it. See footnote 13 in [A], which describes some of Savage's approaches to this problem (and concludes, quite mildly, that "it's just possible that he would have welcomed the CPA").

So there is some ambivalence in this section. In the cited passage, Savage says that the "incompleteness" (the personalistic nature of the probabilities) does not "distress" him. It *does*, however, bother him enough to address; and he *is* willing to call it an "incompleteness." One can read this as saying that though inconsistent beliefs are an embarrassment, they are something one can live with, are not unacceptable, not a reason to reject the theory. But that doesn't mean that he would not have welcomed a more "complete" theory, as [A] suggests.

Note the end of the quote: "...though the harmful effects of (friction in communication) are almost incapable of exaggeration." This is typical of his ambivalence; he seems to be backing away from what he just said—to be saying, "well, maybe disagreements are, after all, the result of 'friction in communication.'"

8. GUL'S "HIERARCHY INTERPRETATION"

This is basically what is presented in Section 2 above. Let us start by noting that Gul does not challenge the formal correctness of [A]; the issues he raises are all interpretive. An axiomatic development, like that in Sections 2–6 above, is ideally suited to deal with such matters: By separating sharply between assumptions and the formal process of deducing conclusions from them, it enables us to focus and clarify the interpretive discussion.

We agree with Gul that the above development of common priors is essentially dynamic, that it depends on the idea of acquiring information and using it to update one's probabilities. And though we consider this perfectly legitimate and intuitive, we also agree that it *would* be desirable to characterize common priors and/or the CPA directly in terms of the "current," posterior probabilities, without any reference, either implicit or explicit, to any prior stage.¹³

We now list Gul's arguments, together with our responses:

(a) *A hierarchy n -tuple does not refer to any prior stage. So the prior stage is irrelevant, and any arguments using a hypothetical, artificially constructed prior stage are inappropriate.*

What Gul fails to realize is that in *any* axiomatic system, the arguments depend crucially on hypothetical, artificial situations that never existed. The essence of the axiomatic approach is that it works with an entire system, and with the relations between the objects in it. In particular, it relates the given, "real," situation to a whole lot of other situations, all of them hypothetical. This happens constantly—in Arrow's (1951) social welfare theory, in the Shapley value (1953), in Nash's (1950) bargaining solution, indeed even in Savage's (1954) development of probabilities themselves! Thus to derive Nash's solution of a given unsymmetric bargaining problem, one *must* consider a whole lot of other problems, including a symmetric one. That is, one must say how the players *would* divide the payoff if they were in this other situation, even though they are not, never were, and never will be. In Savage's theory, to deal with a single "real" choice problem, one must consider a host of other, hypothetical ones.¹⁴ Similarly for Arrow, Shapley, and indeed almost any axiomatic system.

¹³Nevertheless, we disagree with Gul's assertion that "to talk about one player being more ... informed than another ... necessitate(s) a dynamic framework in which information is actually acquired." In fact, no dynamics are required for this, in either representation of information. With partitions, 1 is more informed than 2 at the state ω if and only if the atom of \mathcal{P}^1 containing ω is strictly included in an atom of \mathcal{P}^2 . This nicely captures the intuitive idea of being "more informed:" 1 knows everything that 2 knows, but 2 does not know everything that 1 knows. With hierarchies, 1 is more informed than 2 if and only if 1 knows 2's hierarchy (i.e., 1's hierarchy allows for only one hierarchy of 2, and so ascribes probability 1 to it), but 2 does not know 1's.

¹⁴Savage himself was very much aware of the crucial role played in probability by hypothetical constructs that never existed and never will. In discussing a medical decision problem, Savage (1971) wrote, "It is quite usual in this theory to contemplate acts that are not actually available. These serve something like construction lines in geometry. A typical decision theoretic argument runs, 'If B were available, I would clearly prefer A to B and B to C ; therefore, my momentary impression that C is more attractive than A will not bear inspection.' In particular, I can contemplate the possibility that the lady dies medically and yet is restored in good health to her husband." In referring to a different problem in the same letter, he wrote, "I would regard it as fanciful but not as nonsense to say, 'You experience sunshine if it rains, and rain otherwise.'" Hypothetical, artificial constructs were bread and butter to Savage.

(b) *Since there never was a prior stage, the prior distribution is meaningless. In particular, a common prior—if there is one—is “uninterpretable.”*

Whereas point (a) challenged the conceptual *validity of the theorem* and/or its derivation, here it is the conceptual *meaningfulness of the conclusion* that is challenged.

In response, we note first that the conclusion does have at least one important practical implication not involving the prior at all, and so not requiring “interpretation” of the prior: When knowledge and probability 1 are the same,¹⁵ the action profile at any state must be among those occurring with positive probability in some correlated equilibrium. Since the correlated equilibria often constitute a rather circumscribed set, which can be explicitly calculated by solving a finite family of linear inequalities, this has considerable “predictive” content.

Second, we hold that the conclusion *does* have clear intuitive meaning, even as stated. We do not agree that a common prior has no “meaning” or “interpretation” unless there actually was a time at which some particular person held this prior as his probability distribution. The interpretation of a common prior derives naturally from the axioms: It represents the “correct” or “appropriate” probabilities—what people would or should believe—if there were no private information. Admittedly, the existence of such “neutral” probabilities is not obvious. But that is what the axioms yield; and in any case, the *meaning* or interpretation of this concept is clear.

(c) *A state ω in Ω represents an n -tuple of infinite hierarchies of beliefs of the players over the space X of relevant parameter values. A partition structure is only a “notational device” for representing these hierarchies. Therefore, the CPA can be defended only by defending its implications for these infinite hierarchies. These implications are unclear. Therefore, the CPA cannot be defended; it must be rejected.*

First, Gul’s claim is not quite correct; a state ω represents an n -tuple of *knowledge-belief* hierarchies, not just of belief hierarchies (see Section 2). While this is not very important, it helps us to see why some of Gul’s other claims may be problematic (see Section 8d).

Second, we do not agree that partition structures are merely “notational devices” to represent hierarchies. The two models are simply equivalent. The hierarchy model has the advantage that its tautological nature is more apparent. For other purposes, including ours, the partition model is more useful; it is more transparent and less cumbersome. Since the models are equivalent, one may choose the more convenient one for any particular task that comes to hand; for our task, the partition model is more convenient.

Third, even if we would grant that the hierarchy model is the “primitive,” our argument would not be affected. We could still construct Ω and the various partition structures \mathbf{P} (see the scenario in Section 6), still assume the axioms in Section 4, and still deduce common priors. If the originally given hierarchies do not admit common priors, the axioms are violated. Thus if Gul wishes to challenge the CPA, he *must* come to grips with the axioms.

Fourth, the idea that interpretations must be directly in terms of primitives is strange. We do not go back to Dedekind cuts every time we think about real numbers, nor to ϵ , δ every time we think of continuity. Savage derives probabilities from preferences, but surely we do not return to preferences every time we formulate an idea about probability.

¹⁵I.e., at each state, each player assigns probability 1 to an event if and only if he knows it. When this is not the case, a similar but more complicated conclusion holds.

Try “interpreting” covariance 0, or even just independence, directly in terms of preferences!

(d) *How can common priors reflect the view that “differences in probabilities reflect differences in information only” if the beliefs of player 1 alone can preclude a common prior?*

The force of this argument eludes us. First, it is *always* the case that the hierarchy of one player alone determines the entire associated partition structure, and therefore also whether or not the CPA applies; see footnote 8. Second, our position can be roughly summarized by “CPA iff (differences in probabilities imply differences in information).” Gul’s example has “*not* CPA.” How is our position contradicted?

(e) *The CPA does not imply that the beliefs of an “outside observer” over the players’ action profiles constitute a correlated equilibrium.*

This sounds incorrect, at least when the term “outside” has its straightforward, usual meaning, as in [A]—that the observer has no private information.¹⁶ If, as Gul suggests, we take him as an ordinary person with his own beliefs, then our axioms (i.e., the CPA) apply to him as well as the players; so since he has no private information, his probability must be a common prior of the players. By [A], this is a correlated equilibrium.

In any case, outside observers are not central to our view; in [A] they appear only once, and then not in connection with the CPA.

(f) *Arguing for common priors on the grounds that people with the same information should have the same beliefs is like arguing for risk aversion on grounds of decreasing marginal utility.*

The comparison with risk aversion is quite apt in a number of ways. In both cases there are three modes of formulation:

the purely conceptual, as in “people dislike taking risks” and “beliefs are based on information;”

the abstract formal, as in “a sum of money is preferred to a lottery whose expectation is that sum” and Axiom (1) in Section 4; and

the concrete formal, as in “the marginal utility of money is decreasing” and “there exists a common prior.”

We don’t believe in one formulation “because” we believe in another; the three modes simply express the same ideas in different ways. Indeed, we don’t “believe” in these axioms at all; they are not articles of faith. Like Euclid’s parallel postulate, they have intuitive appeal, but are not “compelling.” Their importance stems from their pulling together a large body of theory, while maintaining a certain simplicity and spareness. Like other scientific hypotheses, they must be judged by where they lead rather than by considerations of innate plausibility.

Institute of Mathematics and Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel.

Manuscript received July, 1994; final revision received July, 1997.

¹⁶By this we mean that it is common knowledge that each player knows what the observer knows; i.e., the observer’s partition coarsens each player’s partition.

REFERENCES

- ARROW, K. J. (1951): *Social Choice and Individual Values*. New York: John Wiley and Sons.
- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1-18 ([A] in the text).
- GALE, D. (1979): "The Game of Hex and the Brouwer Fixed Point Theorem," *American Mathematical Monthly*, 86, 818-827.
- GUL, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66, 923-927.
- HARSANYI, J. C. (1967-68): "Games with Incomplete Information Played by Bayesian Players," *Management Science*, 14, 159-183, 320-334, 486-502.
- MERTENS, J.-F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1-29.
- NASH, J. F. (1950): "The Bargaining Problem," *Econometrica*, 18, 155-162.
- SAVAGE, L. J. (1954): *The Foundations of Statistics*. New York: John Wiley and Sons.
- (1971): "Letter from Leonard Savage to Robert Aumann," in *Essays on Economic Decisions Under Uncertainty*, by J. H. Dreze. Cambridge: Cambridge University Press (1987), 78-81.
- SHAPLEY, L. S. (1953): "A Value for n -person Games," in *Contributions to the Theory of Games*, Vol. II, ed. by H. W. Kuhn and A. W. Tucker. Princeton: Princeton University Press, 305-317.
- ZERMELO, E. (1913): "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," *Proceedings of the Fifth International Congress of Mathematicians (1912)*, 2, 501-504.