

Correlated Equilibrium as an Expression of Bayesian Rationality

Author(s): Robert J. Aumann

Source: *Econometrica*, Jan., 1987, Vol. 55, No. 1 (Jan., 1987), pp. 1-18

Published by: The Econometric Society

Stable URL: <https://www.jstor.org/stable/1911154>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



The Econometric Society is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*

JSTOR

CORRELATED EQUILIBRIUM AS AN EXPRESSION OF BAYESIAN RATIONALITY

BY ROBERT J. AUMANN¹

Correlated equilibrium is formulated in a manner that does away with the dichotomy usually perceived between the "Bayesian" and the "game-theoretic" view of the world. From the Bayesian viewpoint, probabilities should be assignable to everything, including the prospect of a player choosing a certain strategy in a certain game. The so-called "game-theoretic" viewpoint holds that probabilities can only be assigned to events not governed by rational decision makers; for the latter, one must substitute an equilibrium (or other game-theoretic) notion. The current formulation synthesizes the two viewpoints: Correlated equilibrium is viewed as the *result* of Bayesian rationality; the equilibrium condition appears as a simple maximization of utility on the part of each player, given his information.

A feature of this approach is that it does not require explicit randomization on the part of the players. Each player always chooses a definite pure strategy, with no attempt to randomize; the probabilistic nature of the strategies reflects the uncertainty of *other* players about his choice. Examples are given.

KEYWORDS: Correlated equilibrium, Bayesian rationality, information, noncooperative games, strategic equilibrium, Nash equilibrium, common priors, Harsanyi doctrine.

1. INTRODUCTION

O wad some pow'r the giftie gie us
To see oursels as ithers see us!

—Burns

THE EQUILIBRIUM CONCEPT of Nash (1951), together with its refinements,² is without doubt the single game-theoretic tool that is most often applied in economics. Yet, though at first the definition seems simple and natural enough, a little reflection leads to some puzzlement as to why and under what conditions the players in an n -person game might be expected to play such an equilibrium. Recall that it is defined as an n -tuple of strategies in which the component strategy of each player maximizes that player's utility given that the other players are playing their components. Now why should any player assume that the other players will play their components of such an n -tuple, and indeed why should they? This is particularly perplexing when, as often happens, there are multiple equilibria; but it has considerable force even when the equilibrium is unique. Indeed, there are games whose Nash equilibria appear quite unattractive even though they are unique (see Harsanyi, 1977, p. 125); and even without such examples, it does not seem clear why the players would play even a unique Nash

¹ Research supported by the Institute for Mathematical Studies in the Social Sciences, Stanford University, under National Science Foundation Grant SES 83-20453; by the Institute for Mathematics and its Applications, University of Minnesota; and by the Mathematical Sciences Research Institute, Berkeley. We are grateful to Lloyd Shapley for an illuminating conversation, a decade back, that clarified for us the equivalence between the two definitions of correlated equilibrium; to Kenneth Arrow, for motivating us to write this paper; to Bob Weber, for the very apt quotation from Burns; and to Dave Kreps and two anonymous referees, for splendid editorial work.

² E.g., Selten (1975), Myerson (1978), Kreps and Wilson (1982), Kalai and Samet (1984), Kohlberg and Mertens (1986).

equilibrium. In a two-person game, for example, Player 1 would play his component only if he believes that 2 will play his; this in turn would be justified only by 2's belief that 1 will indeed play his component; and so on. To many this will sound like a plain old circular argument: consistent, perhaps, but hardly compelling.

Nash equilibrium does make sense if one starts by assuming that, for some specified reason, each player knows which strategies the other players are using. But this assumption appears rather restrictive.

Another criticism of equilibrium that has been advanced is that it appears inconsistent with the modern subjectivist, Bayesian view of the world. According to the Bayesian view, subjective probabilities should be assignable to every prospect, including that of players choosing certain strategies in certain games. Rather than playing an equilibrium, the players should simply choose strategies that maximize their utilities given their subjective distributions over the other players' strategy choices.

It is the purpose of this paper to provide a simple rationale for equilibrium. Surprisingly, our rationale is precisely in terms of the "criticism" in the previous paragraph. We will show that, far from being inconsistent with the Bayesian view of the world, the notion of equilibrium is an unavoidable consequence of that view. It turns out, though, that the appropriate equilibrium notion is not the ordinary mixed strategy equilibrium of Nash (1951), but the more general notion of correlated equilibrium.

Specifically, we will show the following. Let us be given an n -person game G in strategic (i.e. normal) form. Assume that (i) as in Savage (1954), each player has a subjective probability distribution over the set of all states of the world; and (ii) it is common knowledge³ that each player chooses a strategy that maximizes his expected utility given his information. Then the strategies chosen by the players constitute a correlated equilibrium of G . Conversely, each correlated equilibrium of G can be obtained in this way for an appropriate choice of the parameters. A precise formulation and proof will be found in Section 3.

To some readers, this may seem obvious. Once one has the formulation, it is indeed embarrassingly easy to prove, as we shall see. But it should not be confused with certain characterizations of Nash equilibrium, which look similar but are actually much more transparent. There one assumes not only that each player is a utility maximizer, but that he knows the strategies actually being used by all other players. We make no such assumption; indeed, in our treatment, the players do *not* in general know how others are playing. We assume only that it is common knowledge that all the players are Bayesian utility maximizers, that they are rational in the sense that each one conforms to the Savage theory. Such an assumption underlies most of game theory, and of economic theory as well; we show that it, by itself, is enough to assure that the outcome is a correlated equilibrium.

³ An event is *common knowledge* among a set of agents, if it is known to all, it is known to all that it is known to all, and so on ad infinitum. See Lewis (1969) and, for a precise mathematical formulation, Aumann (1976).

An important feature of our approach is that it does not require explicit randomization on the part of the players. Each player always chooses a definite pure strategy, with no attempt to randomize; the probabilistic nature of the strategies reflects the uncertainty of *other* players about his choice. This will be discussed in detail in Section 6.

The basic idea of this paper is extraordinarily simple; the main theorem is proved in less than a dozen lines. Some conceptual issues do arise, but still, we cannot grasp how such a simple idea grew into a paper of this length. A reader just interested in the main points, who knows what a correlated equilibrium is, should read Sections 3 and 6 only. Those not familiar with correlated equilibria should read Section 2 first.

2. CORRELATED EQUILIBRIUM

Let G be an n -person game in strategic form. For $i = 1, \dots, n$, denote by S^i the set of pure strategies—henceforth *actions*—of Player i , and by $h^i(s)$ the payoff to i for an n -tuple $s := (s^1, \dots, s^n)$ of actions.⁴ Set $S := S^1 \cdots \times S^n$ and $h := (h^1, \dots, h^n)$. Define a *correlated strategy n -tuple* in G as a function f from a finite⁵ probability space Γ into S ; that is, as a random variable whose values are n -tuples of actions.

Correlation is a more general form of strategy randomization than mixing. In both cases, the players base their choice of an action on observation of a random event; with mixed strategies, the observations are independent, whereas with correlated strategies, they need not be. To see that mixed strategies may be formally viewed as correlated strategies, take Γ to be the product probability space $\Gamma^1 \times \cdots \times \Gamma^n$, where Γ^i is the probability space by means of which i 's mixed strategy is defined—the set of outcomes of his roulette spin. The simplest examples of “true” correlation (nonindependent action choices) occur when all players base their choices on observation of the *same* random variable. For example, in the prisoner's dilemma of Figure 1, the symmetric efficient outcome can be achieved if *both* players simultaneously observe a single toss of a fair coin, then play (top, right) or (bottom, left) according as to whether heads or tails came up. It cannot, however, be achieved by mixed strategies.

In general, the observations of the players need be neither identical nor independent. Chance chooses an element γ of Γ , then suggests to each player i

2, 2	0, 6
6, 0	1, 1

FIGURE 1

⁴ The symbol $:=$ means that the expression on the left is being defined.

⁵ I.e., with a finite number of points (and so also of events). The finiteness of Γ plays no role in the sequel; we assume it only to avoid complicating the model with irrelevant generality.

that he take the action $f^i(\gamma)$ ($:= (f(\gamma))^i$). If all players follow the suggestions, the correlated strategy n -tuple f results.

The function f may be identified with the n -tuple (f^1, \dots, f^n) . If g^i is a different function from Γ to S^i , then as usual,

$$(f^{-i}, g^i) := (f^1, \dots, f^{i-1}, g^i, f^{i+1}, \dots, f^n).$$

Throughout the sequel, E denotes “expectation.” If f is a correlated strategy n -tuple, note that $h^i(f)$ is a real-valued random variable.

DEFINITION 2.1: A *correlated equilibrium* (c.e.) in G is a correlated strategy n -tuple f such that

$$(2.2) \quad Eh^i(f) \geq Eh^i(f^{-i}, g^i)$$

for each player i and each g^i that is a function of f^i .

In form this definition is very similar to that of Nash equilibrium. Equilibrium is achieved when no player can gain by deviating from the suggestions, given that the others obey them. The deviations g^i are restricted to be functions of f^i (i.e., compositions of some other function with f^i) because i is informed of $f^i(\gamma)$ only, and so can distinguish only between members of Γ that are distinguished by f^i .

It will be useful to define the *distribution* of a correlated strategy n -tuple f as the function that assigns to each n -tuple s of actions the number $\text{Prob} \{f^{-1}(s)\}$. Much like mixed strategies, correlated strategy n -tuples can for most practical purposes be identified with their distributions. Thus the correlated strategy n -tuple discussed above (Figure 1) can, using a familiar notation, be written $\frac{1}{2}(\text{top, right}) + \frac{1}{2}(\text{bottom, left})$. Another way to represent the distribution is to insert the appropriate probability in each square of the payoff matrix; for the above example, the distribution is that of Figure 2.

This distribution is not a *correlated equilibrium distribution* (c.e.d.), i.e., it does not represent a correlated equilibrium. Indeed, in the prisoner’s dilemma, it is *always* worthwhile for Player 1 to play bottom, also if it is suggested that he play top. Figure 3 contains an example of a game (“the battle of the sexes”) and a correlated strategy distribution that *is* in equilibrium. If top is suggested to 1, then he knows that left was suggested to 2, so he would lose by deviating to bottom; similarly if bottom is suggested, and for Player 2.

By similar methods it may be seen that any convex combination of Nash equilibria is a correlated equilibrium. But there are also other kinds of correlated

0	$\frac{1}{2}$
$\frac{1}{2}$	0

FIGURE 2

2, 1	0, 0	$\frac{1}{2}$	0
0, 0	1, 2	0	$\frac{1}{2}$

FIGURE 3

6, 6	2, 7
7, 2	0, 0

FIGURE 4

equilibria. For example, Figure 4 illustrates a two-person game (“chicken”) with three Nash equilibria, whose payoffs are (2,7), (7,2), and $(\frac{4}{3}, \frac{4}{3})$. The distribution of Figure 5 may be seen to be a correlated equilibrium, whose payoff (5,5) is outside the convex hull of the Nash equilibrium payoffs. It follows that the distribution of Figure 5 is itself outside the convex hull of the Nash equilibria. Other examples of correlated equilibria that are outside the convex hull of the Nash equilibria may be found in Aumann (1974); for a particularly beautiful example, see Moulin and Vial (1978).

Computationally, correlated equilibria are simpler objects than Nash equilibria. The correlated equilibrium distributions of a given game G constitute a compact convex polyhedron whose defining linear inequalities can be explicitly written down. Here we treat only the two-person case, which is notationally less cumbersome; the principle, however, is no different in the general case.

Let $n = 2$; set $l := |S^1|$, $m := |S^2|$, $h_{jk} := h(j, k)$ for $j \in S^1$ and $k \in S^2$. Let \sum_j and \sum_k denote summation over all j in S^1 and all k in S^2 respectively. The distribution of a correlated strategy pair is an lm -tuple (p_{jk}) , where j and k range over S^1 and S^2 respectively,

$$p_{jk} \geq 0 \quad \text{for all } j \text{ and } k, \quad \text{and} \quad \sum_j \sum_k p_{jk} = 1.$$

We will call such objects simply *distributions* for short.

$\frac{1}{3}$	$\frac{1}{3}$
$\frac{1}{3}$	0

FIGURE 5

PROPOSITION 2.3: A distribution (p_{jk}) is a correlated equilibrium distribution if and only if

$$(2.4) \quad \sum_k (h_{jk}^1 - h_{qk}^1) p_{jk} \geq 0 \quad \text{for all } j, q \text{ in } S^1, \text{ and}$$

$$\sum_j (h_{jk}^2 - h_{jr}^2) p_{jk} \geq 0 \quad \text{for all } k, r \text{ in } S^2.$$

PROOF: For (2.2) it is necessary and sufficient that for each i , the expectation still obeys the inequality when conditioned on each possible value of f^i , i.e. on each possible “suggestion” to Player i (a suggestion is called *possible* if it has positive probability). Suppose j is a possible suggestion to Player 1, i.e., $\sum_k p_{jk} > 0$. Given this suggestion, the conditional probability of 1 that k was suggested to 2 is $p_{jk}/\sum_k p_{jk}$; hence 1’s conditional expected payoff is $\sum_k h_{jk}^1 p_{jk}/\sum_k p_{jk}$, which we denote $H^1(j|j)$. If 1 were to deviate to the action q , his conditional expected payoff would be $\sum_k h_{qk}^1 p_{jk}/\sum_k p_{jk}$, since the payoffs change, but 1’s information remains the same; we denote this by $H^1(q|j)$. The equilibrium condition (2.2) for $i = 1$ says that deviation to q is not worthwhile, i.e., that $H^1(j|j) \geq H^1(q|j)$; if we multiply through by $\sum_j p_{jk}$, we obtain (2.4). Similarly, (2.5) is seen to express the equilibrium condition (2.2) for $i = 2$.

Definition 2.1 appears a little different from our 1974 definition of correlated equilibrium; there, the players may get more information about γ than just the suggested action $f^i(\gamma)$. It turns out, though, that practically speaking, the two definitions are equivalent. This will be discussed again below (Section 4g).

3. BAYESIAN RATIONALITY IN GAMES

We assume exogenously given an n -person game as in the previous section. Also exogenously given are the following: (i) a finite⁶ set Ω , with generic element ω ; (ii) for each player i in G , a probability measure p^i on Ω ; (iii) for each player i in G , a partition \mathcal{P}^i of Ω . The set Ω represents the set of all possible *states of the world*, and each ω denotes a specific such state. The probability measure p^i is i ’s *prior* on Ω . The partition \mathcal{P}^i is i ’s *information partition*; if the true state of the world is $\omega \in P \in \mathcal{P}^i$, then i knows that some element of P is the true state of the world, but he does not know which one it is.

The term “state of the world” implies a definite specification of all parameters that may be the object of uncertainty on the part of any player of G . In particular, each ω includes a specification of which action is chosen by each player of G at that state ω . Conditional on a given ω , everybody knows everything; but in general, nobody knows which is really the true ω . Taking the “atoms” of Ω to represent complete specifications of all possible variables enables us to represent *all* aspects of uncertainty on the part of any player—including uncertainty about the uncertainty of other players—by means of the partitions \mathcal{P}^i . This constitutes the standard model for “differential” (“incomplete”) information in multi-person contexts; its implications will be discussed in more detail in the next section.

⁶ See Section 4c.

We will use the following assumption:

COMMON PRIOR ASSUMPTION: All the priors p^i are the same; that is, there is a probability measure p on Ω such that $p^1 = p^2 = \dots = p^n = p$.

This is sometimes called the *Harsanyi Doctrine*. It does *not* imply that all players have the same subjective probability. The subjective probability of a player is his posterior given his information; these may well be different. Roughly speaking, the common prior assumption says that differences in probability estimates of distinct individuals should be explained by differences in information and experience. This, too, will be discussed further in the sequel, and we will also indicate how our results must be modified when the common prior assumption is relaxed or abandoned.

Denote by $s^i(\omega)$ the action chosen by Player i in the game G at the state ω of the world. Formally, each s^i is simply an exogenously specified function from Ω to S^i ; it may be considered item (iv) in the informational model of the world specified at the beginning of this section. We make the modest assumption that each player i knows which action he chooses, i.e., that s^i is measurable⁷ w.r.t. (with respect to) \mathcal{P}^i .

Let $\mathbf{s}(\omega) := (s^1(\omega), \dots, s^n(\omega))$ be the n -tuple of actions chosen at state ω . Call a player i *Bayes rational at ω* if his expected payoff given his information, $E(h^i(\mathbf{s}) | \mathcal{P}^i)(\omega)$, is at least as great as the⁸ amount $E(h^i(\mathbf{s}^{-i}, s^i) | \mathcal{P}^i)(\omega)$ that it would have been had he chosen an action s^i other than the action $s^i(\omega)$ that he did in fact choose; in brief, if he chooses an action that maximizes his payoff given his information.

MAIN THEOREM: *If each player is Bayes rational at each state of the world, then the distribution of the action n -tuple \mathbf{s} is a correlated equilibrium distribution.*

PROOF: We must exhibit a probability space Γ and a function $f: \Gamma \rightarrow S$ that is a correlated equilibrium and that has the same distribution as \mathbf{s} . Indeed, if we set $\Gamma := (\Omega, p)$, then \mathbf{s} itself is such a function. To see this, let $i \in N$, and let $g^i: \Omega \rightarrow S^i$ be any function of s^i . Since s^i is \mathcal{P}^i -measurable, g^i also is; that is, it is constant on each P in \mathcal{P}^i . The Bayes rationality of i at each state implies that for each such P and each s^i in S^i ,

$$E(h^i(\mathbf{s}) | P) \geq E(h^i(\mathbf{s}^{-i}, s^i) | P).$$

⁷ A random variable is *measurable* w.r.t. a partition if it is constant on each element of the partition. Intuitively, this means that the r.v. contains no more information than the partition; a person who knows which event in the partition obtains, also knows the value of the r.v.

⁸ Recall that if x is a random variable, then $E(x | \mathcal{P}^i)$ is defined as the r.v. whose value at a specific ω is $E(x | P)$, where P is the event in \mathcal{P}^i that contains ω . (P represents i 's knowledge, the smallest event that he knows to have occurred.)

Taking s^i to be the constant value of g^i throughout P yields

$$E(h^i(\mathbf{s})|P) \geq E(h^i(\mathbf{s}^{-i}, g^i)|P);$$

multiplying both sides by $\text{Prob } P$ and summing over all P in \mathcal{P}^i yields

$$E(h^i(\mathbf{s})) \geq E h^i(\mathbf{s}^{-i}, g^i).$$

For $f = \mathbf{s}$, this is precisely the condition (2.2) that defines correlated equilibrium.

4. DISCUSSION

(a) *Personal Choice as a State Variable.* The chief innovation in our model is that it does away with the dichotomy usually perceived between uncertainty about acts of nature and of personal players. Of course, a player always knows which decision he himself takes; here, this information is not treated differently from private information he may have about other aspects of the state of the world.

This point is a little subtle and is worth some discussion. In traditional Bayesian decision theory, each decision maker is permitted to make whatever decision he wishes, after getting whatever information he gets. In our model this appears not to be the case, since the decision taken by each decision maker is part of the description of the state of the world. This sounds like a restriction on the decision maker's freedom of action; at a given state ω , it is as if the model forces him to take the decision dictated by ω .

But closer examination reveals that "freedom of choice" is not an issue. The model describes the point of view of an outside observer. Such an observer has no a priori knowledge of what the players will choose; for him, the choices of the players are part of the description of the states of the world. This does not mean that the players cannot choose whatever they want, but only that the observer will not in general know what they want.

This "outside observer" perspective is common to all differential information models in economics (as well as to all extensive games that are not of perfect information). In such models, each player gets some information or "signal"; he hears only the signal sent to him, not that of others. In analyzing this situation, each player i must first look at the whole picture as if he were an outside observer; he cannot ignore the possibility of his having gotten a signal other than the one he actually got, even though he knows that he did not actually get such a signal. This is because the *other* players do not know what signal he got. Player i must take the ignorance of the other players into account when deciding on his own course of action, and he cannot do this if he does not explicitly include in the model signals other than the one he *knows* he got. The "outside observer" referred to above is thus a surrogate for the ignorance of the system as a whole—the lack of common knowledge—of the signals received by each player.

Similarly, each player is of course free to choose whatever action he wishes. That we include his action as part of the description of the state of the world is

only a convenient way of expressing the fact that the other players do not know which action he wishes to choose.

The reader may still be puzzled by the fact that in deriving his posterior about other players' choices, each player conditions on his own information, which includes his own choice. How can a player's own choice help him to guess what other players will do?

But again, this paradox is illusory. The player is not really conditioning on his choice, but on the substantive information that leads him to make this choice. This substantive information leads to a posterior on the other players' choices, which in turn leads to an optimal choice, or to a set of such choices. Intuitively, the fact that a player's choice is part of his information is used not so much in deriving his own posterior about what others will do, but rather in estimating the others' posteriors about what he will do. They cannot simply make arbitrary "guesses" about this, but must take into account that he is maximizing given his own information. Since he is reasoning similarly about them, it appears that one might be led to troubling circular reasoning. This precisely is avoided by the model of the previous section.

To sum up, a game in which the players do not know what the other players do should be treated like any other model of differential information, where the possible values of the parameters about which some of the players are ignorant are used as specifications of the state of the world.

(b) *Common Knowledge of Information Partitions and Priors.* A question that often crops up when models of differential information are discussed is whether there can be uncertainty on the part of one player about the information partitions \mathcal{P}^i of other players.

The answer is "no". While Player 1 may well be ignorant of what Player 2 knows—i.e., of the element of \mathcal{P}^2 that contains the "true" state ω of the world—1 cannot be ignorant of the *partition* \mathcal{P}^2 itself. Indeed, \mathcal{P}^2 is part of the description of the model, and does not enter the description of any particular ω ; it therefore cannot be the object of uncertainty, it must be common knowledge.

This is not an assumption, but a "theorem", a tautology; it is implicit in the model itself. Since the specification of each ω includes a complete description of the state of the world, it includes also a list of those other states ω' of the world that are, for Player 2, indistinguishable from ω . If there were uncertainty about this list on the part of Player 1 (or of any other player), then the description of ω would not be complete; one should then split ω into several states, depending on which states are, for 2, indistinguishable from ω . Therefore the very description of the ω 's implies the structure of \mathcal{P}^2 , and similarly of all the \mathcal{P}^i . The description of the ω 's involves no "real" knowledge; it is only a kind of code book or dictionary. The structure of the \mathcal{P}^i also involves no real knowledge; it simply represents different methods of classification in the dictionary.

The situation with priors is similar. Once one accepts the Bayesian viewpoint, that each player has a prior on Ω , it follows that there cannot be any uncertainty on the part of one player about other players' priors. Each player's prior must be common knowledge among all players. (This in itself does not imply the

further assumption, which we have made, that all the priors are the same; this further assumption will be discussed separately below.)

The reasoning leading to the conclusion that the priors are common knowledge is similar to that leading to the conclusion that the \mathcal{P}^i are common knowledge. It is enough to convince ourselves that Player 1's prior $p^1(\omega)$ for each given state ω of the world is common knowledge. If it were not, then the description of ω would be incomplete; we would be able to split ω into several states, depending on the various possibilities for $p^1(\omega)$.

The reader may ask why we adduce verbal, informal arguments for these assertions, rather than proving them as formal theorems. The answer is that the assertions have no formal content. Within the model, knowledge refers to events, i.e., sets of states. One can ask whether at a given state ω , a player i knows an event A ; this is the case if and only if A includes that element of \mathcal{P}^i that contains ω . But neither a partition nor a prior is an event; formally speaking, the concepts of knowledge and common knowledge do not apply to them. In the above informal discussion, when we refer to the partitions, say, as "common knowledge", we mean that all players know them, all know that all know them, and so on, where "know" has its informal, everyday meaning.

It is important to realize that while the above assertions (that the partitions and priors are common knowledge) aid us in understanding the model, they do not affect the conclusions. Each player uses only his own prior and his own partition in reaching a decision; within the model, it does not matter that he "knows" the partitions and priors of the others. It is only in interpreting the results that these points become significant.

(c) *Finiteness of the Model.* Unlike the finiteness of Γ assumed in Section 2, the finiteness of Ω does involve some loss of generality. For example, in describing uncertainty about other people's uncertainty, it is perhaps more natural to allow a continuum of states. Models that do this in full generality and entirely explicitly are, in fact, available; e.g., Mertens and Zamir (1985).

We nevertheless chose the finite model, because the point we make is basically very simple, and has nothing to do with whether the model is finite or infinite. This simplicity might be obscured by an infinite model, with its relatively heavy apparatus of σ -fields, measures, integrals, Radon-Nikodym derivatives, and so on. Such an analysis can probably be carried out in a manner analogous to that presented here; but we have not actually done this, and the possibility of unexpected difficulties cannot be excluded.

In the last analysis, the world *is* usually considered finite; in a sense, finite models do appear more natural. An infinite model is appropriate when there is something to be gained from it, when we are dealing with a phenomenon that is not conveniently expressible in a finite framework. The choice of model is a matter of convenience; here, the finite model is the more convenient and transparent one.

(d) *The Converse.* Let us call the system consisting of Ω , p , the \mathcal{P}^i , and s an *information system*. The main theorem asserts that under Bayesian rationality, every information system corresponds to a correlated equilibrium. The converse

also holds: Under Bayesian rationality, every correlated equilibrium corresponds to some information system. More precisely, for each game G and each correlated equilibrium f of G , there is an information system for which it is Bayes rational for the players to play in accordance with s , and the resulting distribution is the same as that of f .

For the proof, note that f itself provides an example of such an information system. We need only define $(\Omega, p) := \Gamma$, $s := f$, and \mathcal{P}^i to be the partition of Ω generated by f^i (i.e., ω and ω' are in the same element of \mathcal{P}^i iff $f^i(\omega) = f^i(\omega')$).⁹ Bayesian rationality is then a restatement of (2.2). Thus under Bayesian rationality, the set of all information systems corresponds precisely to the set of all correlated equilibria.

(e) *Exogeneity and Endogeneity.* As we indicated at the beginning of the previous section, it is probably best to think of the information system as exogenous. But a case could be made for another view. While a part of each player's information is undoubtedly generated by overt signals from the outside world, another part is obtained by reasoning about how other players reason. For example, this is so for the information that a strictly dominated strategy will not be used. Some readers might prefer to call that kind of information "endogenous." In actuality, much information is a mixture of both kinds, which it is not easy to untangle.

The distinction between "exogenous" and "endogenous" is often useful in economics, but it should not be pushed too far. In the natural sciences, such a distinction is little used.¹⁰ When discussing the motion of the planets, should the motion of Uranus be considered "exogenous" and that of Neptune "endogenous", or the other way around? Perhaps it is not so terribly important. What *is* important is the relationship between the motions.

The purpose of this paper is to record the observation that common knowledge of rationality implies correlated equilibrium. For this, it is not very important whether the information of the players—or even their actions—are considered exogenous or endogenous.

(f) *Mixed Strategies.* In view of the fact that the players may wish to use mixed strategies, the reader might question the assumption that each player knows his own action. There is no problem if, as usual, one thinks of a mixed strategy simply as a random device for helping the player make up his mind, so that in the end he does know which action he takes. But it is conceivable that a player can actually commit himself to a mixed strategy before he knows its pure outcome; in that case, the assumption in question seems formally incorrect.

The simplest way to circumvent the difficulty is to say that if indeed it is possible for a player i to commit himself irrevocably to a definite mixed strategy, then such a commitment should itself be considered an action—a member of S^i —and then, of course, there is no difficulty. It should be noted that this will

⁹ Note that the definitions of Ω , p , and s go in exactly the opposite direction from the corresponding ones in the proof of the main theorem, where we defined $\Gamma := (\Omega, p)$, $f := s$. But the proof of the main theorem is not simply a mirror image of this one, since there p' was not necessarily generated by f' .

¹⁰ We are aware of the pitfalls of blindly applying the methodology of one science to another. But neither should one go to the other extreme, of blindly rejecting all parallels.

not change the set of possible outcomes; adding an action whose payoffs are a convex combination of the payoffs to other actions does not change the set of correlated equilibria.

(g) *Alternative Definitions of Correlated Equilibrium.* In our 1974 definition of correlated equilibrium, each player i is endowed with an information partition \mathcal{I}_i on the underlying probability space (the space of possible outcomes of the randomizing device). After finding out which element of \mathcal{I}_i took place, i may take any action he wishes. Thus an equilibrium is a function from outcomes of the randomizing device to action n -tuples, whose i th component is \mathcal{I}_i -measurable (each player knows which action he takes). The equilibrium condition is as in Section 2 above.

At first blush, this definition appears different from that of this paper (2.1), since here a player is told *only* what action he takes, and is given no further information. For fixed \mathcal{I}_i , the definitions are indeed different. But if the \mathcal{I}_i are allowed to vary, the set of correlated equilibrium distributions obtained under the 1974 definition is the same as that of this paper.

To see this, note that the 1974 set-up is formally identical with that of Section 3 above, with \mathcal{I}_i instead of \mathcal{P}^i . Thus the fact that every c.e.d. in the 1974 sense is also a c.e.d. in the current sense is simply a restatement of our main theorem. The converse is immediate, since we may, if we wish, take \mathcal{I}_i to consist of the coarsest partition of the underlying probability space for which i 's action is measurable.

The reason that in this paper we chose the definition without the \mathcal{I}_i is that it makes the explicit calculation of the set of all c.e.d.'s more transparent.

(h) *Common Knowledge and Universality.* The formulation of our result in the introduction assumes common knowledge (c.k.) of rationality. The formal treatment of Section 3 assumes *universal* rationality, i.e., rationality at each point of the state space Ω . These two formulations are equivalent. Indeed, our 1976 treatment of c.k. implies that the universal event Ω is itself c.k. (i.e., anything that is true at all possible states of the world is common knowledge). Conversely, a c.k. event can without loss of generality be assumed universal, because one can always restrict the universe to the smallest c.k. event (i.e., to that member of the meet $\wedge \mathcal{P}^i$ of the information partitions that contains the true state ω of the world).

5. THE COMMON PRIOR ASSUMPTION

This section is devoted to a discussion of the only element of our model that is not a tautological consequence of the Bayesian approach: the common prior assumption (CPA for short; see Section 3).

Common priors are explicit or implicit in the vast majority of the differential information literature in economics and game theory. As soon as one writes, "let p (rather than p^i) be the probability of . . .", one has assumed common priors. The assumption is pervasive in the enormous literature on rational expectations, trading in securities, bargaining under incomplete information, auctions, repeated games, signalling, discrimination, insurance, principal-agent, moral hazard,

search, entry and exit, bankruptcy, what have you. Citing the relevant papers would make our references longer than our text. Occasionally the definitions do pay lip-service to the possibility of distinct priors p^i ; but usually this is quickly abandoned, and in the theorems and examples, one returns to common priors.¹¹

In game theory, the standard representation of extensive games¹² involves moves of “chance” or “nature” where each alternative has a single probability, common to all players. The standard formulation of Nash equilibrium (1951) uses mixed strategies for which the probabilities are common to all players. Both these definitions have been used hundreds—perhaps thousands—of times in economic applications, almost without question.

Why is this so? Why has the economic community been unwilling, in practice, to accept and actually use the idea of truly *personal* probabilities, in much the same way that it did accept the idea of personal utility functions? After all, in the development of Savage (1954), both the utilities and the probabilities are derived separately for each decision maker. Why were utilities accepted as personal, but probabilities not?

Perhaps the most basic reason is that utilities directly express tastes, which are inherently personal. It would be silly to talk about “impersonal tastes”, tastes that are “objective” or “unbiased.” But it is not at all silly to talk about unbiased probability estimates, and even to strive to achieve them. On the contrary, people are often criticized for wishful thinking—for letting their preferences color their judgment. One cannot sensibly ask for expert advice on what one’s tastes should be; but one may well ask for expert advice on probabilities.¹³

On a more pragmatic level, the CPA expresses the view that probabilities should be based on information; that people with different information may

¹¹ A notable exception is Harrison and Kreps (1978).

¹² von Neumann and Morgenstern (1944), Kuhn (1953).

¹³ That Bayesian decision theory à la Savage derives both utilities and probabilities from preferences does not imply that it does not discriminate conceptually between these two concepts. Savage used preferences to get a rigorous handle on probabilities, a coherent formulation free of inner contradictions. At the same time, he strove, to the extent possible, to rid his probabilities of personal elements, to make them fit for use not only in decision theory but also for scientific induction. He was fond of pointing out that after sufficiently many observations, all “reasonable” priors lead to similar conclusions. His postulate P4 (roughly speaking, that the probability of an event is independent of the prize offered contingent on that event) can only be understood in terms of a probability concept that has an existence of its own in the decision maker’s mind, quite apart from preferences on acts. He wrote that “. . . the . . . view sponsored here does not leave room for optimism or pessimism . . . to play any role in the person’s judgment” (1954, p. 68). And again, “All views of probability are rather intimately connected with one another. For example, any necessary view can be regarded as an extreme personalistic view in which so many criteria of consistency have been invoked that there is no role left for individual judgment” (1954, p. 60).

Savage is dead—so much the worse for us—and one can only speculate as to how he would have regarded the CPA. As far as we know, he never dealt formally with more than one decision maker at a time. The CPA can be brought formally into his framework via an axiom that says that under common knowledge of rationality, risk averse agents with precisely the same information will never bet against each other (one must, in some appropriate way, exclude hedging). Contrary to the modern vogue, Savage was not a minimalist; he did not try to make his axioms as few and weak as possible, but as useful as possible. At one point he wrote that “the personalistic view incorporates all . . . criteria for reasonableness in judgment known to me, and . . . when any criteria that may have been overlooked are brought forward, they will be welcomed . . .” (1954, p. 67). It’s just possible that he would have welcomed the CPA.

1	2	2
-2	9	-9
-2	-9	9

FIGURE 6

legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so.

Differently put, a major message of modern economic theory is the vital importance of information in economic activity, and of differences in information. Under the CPA, differences in probabilities express differences in information *only*. Thus the CPA enables one to zero in on purely informational issues in analyzing economic (and other interactive) models with uncertainty.

If, in spite of all this, we abandon the CPA, all is not lost. One can define a *subjective correlated equilibrium* as in Section 2, replacing the single probability p on Γ by n different probabilities p^1, \dots, p^n . In the main theorem, then, one must only replace “correlated equilibrium distribution” by “subjective correlated equilibrium distribution.”

While such an approach is mathematically perfectly consistent, it yields results that are far less sharp than those obtained with common priors. For example, if a 2-person 0-sum game G in strategic form has a unique pair s of optimal (maxmin and minmax) strategies, which are moreover pure, then one would certainly expect rational players to play s ; the reasoning leading to this conclusion, which was spelled out in detail by von Neumann and Morgenstern (1944), is extremely compelling. And indeed, s is the only correlated equilibrium point¹⁴ in g . But, it is far from the only subjective correlated equilibrium point. For example, in the game of Figure 6, the unique pair of optimal strategies is (top, left); it leads to the payoff +1 for the row player, and -1 for the column player. But the subjective distribution indicated in Figure 7 is a subjective correlated equilibrium; its payoff is +3 for *each* player (in spite of the game being zero sum!), and it makes no use at all of the unique optimal strategies.

0, 0	0, 0	0, 0
0, 0	$\frac{1}{3}, \frac{1}{6}$	$\frac{1}{6}, \frac{1}{3}$
0, 0	$\frac{1}{6}, \frac{1}{3}$	$\frac{1}{3}, \frac{1}{6}$

FIGURE 7

¹⁴ More generally, in any 2-person 0-sum game, all correlated equilibria are convex combinations of pairs of optimal strategies.

This game is not a fluke. Quite generally, the concept of subjective correlated equilibrium places very few restrictions on the possible outcomes. To get the flavor of this, note that for *any* two-person game, not necessarily zero-sum, in which there are no weakly dominated actions, there is a subjective correlated equilibrium at which both players assign positive probability to each outcome (action pair); and there is also a subjective correlated equilibrium that gives each of the two players an expected payoff as close as you wish to his maximum possible payoff. Such results indicate that the subjective correlated equilibrium is a relatively “weak” concept, giving little information; and that while logically consistent, it involves a conceptual inconsistency between the players, which distorts and hides the conflict of interests that is the subject of game-theoretic analysis.

The case of distinct priors was considered by John Harsanyi (1967) in his path-breaking papers on games of incomplete information; he called it the “inconsistent case.” The idea of subjective randomization was introduced and developed in our 1974 paper; this idea depends essentially on distinct priors. Neither idea had any considerable echo. Apparently, economists feel that this kind of analysis is too inconclusive for practical use, and side-steps the major economic issues.

6. RANDOMNESS AS AN EXPRESSION OF IGNORANCE

In the traditional view of strategy randomization, the players use a randomizing device, such as a coin flip, to decide on their actions. This view has always had difficulties. Practically speaking, the idea that serious people would base important decisions on the flip of a coin is difficult to accept. Conceptually, too, there are problems. The reason a player must randomize in equilibrium is only to keep others from deviating; for himself, randomizing is unnecessary.

The first to break away from the idea of explicit randomization was J. Harsanyi (1973). He showed that if the payoffs to each player i in a game are subjected to small independent random perturbations, known to i but not to the other players, then the resulting game of incomplete information has *pure* strategy equilibria that correspond to the mixed strategy equilibria of the original game. In plain words, nobody really randomizes. The appearance of randomization is due to the payoffs not being exactly known to all; each player, who does know his own payoff exactly, has a unique optimal action against his estimate of what the others will do.

Here we take this reasoning one step further. Even without perturbed payoffs, the players simply do not know what actions the other players will take. In “matching pennies”, each player knows very well what he himself will do, but ascribes $\frac{1}{2}$ - $\frac{1}{2}$ probabilities to the other’s actions, and knows that the other ascribes those probabilities to his own actions.

With this view, mixed strategy equilibria appear quite special and rather unnatural. They imply that the players always act as if they all had the same beliefs about what all other players will do, and as if these beliefs were common

.3	.3
.3	.1

FIGURE 8

knowledge. (The actions themselves, of course, need not be common knowledge, except in the case of pure strategy equilibria.) In particular, they act as if¹⁵ each player *i* always knew exactly what each other player believes about his (*i*'s) actions.

Correlated strategy *n*-tuples, on the other hand, do allow for varying beliefs about the beliefs of other players. In the distribution of Figure 8, for example, Player 1 may ascribe probabilities that are either $\frac{1}{2}-\frac{1}{2}$ or $\frac{3}{4}-\frac{1}{4}$ to Player 2's actions; and 2 can never know what 1 believes about his (2's) actions. Such phenomena are often of the essence in the analysis of real situations, but they are ruled out by the mixed strategy paradigm.

In games with more than two players, correlation may express the fact that what 3, say, thinks that 1 will do may depend on what he thinks 2 will do. This has no connection with any overt or even covert collusion between 1 and 2; they may be acting entirely independently. Thus it may be common knowledge that both 1 and 2 went to business school, or perhaps to the same business school; but 3 may not know what is taught there. In that case 3 would think it quite likely that they would take similar actions, without being able to guess what those actions might be.

For example, consider the game of Figure 9 (Aumann, 1974). Players 1, 2, and 3 pick the row, column, and matrix respectively. No mixed strategy equilibria

0, 0, 3	0, 0, 0	2, 2, 2	0, 0, 0	0, 0, 0	0, 0, 0
1, 0, 0	0, 0, 0	0, 0, 0	2, 2, 2	0, 1, 0	0, 0, 3

FIGURE 9

¹⁵ The phrase "act as if" means that we ignore irrelevant information, information that does not affect the actions of the players. Technically speaking, the corresponding assertions hold when each player conditions on his own action only. Without this caveat, they are false. For example, consider a three-person game in which each player chooses 0 or 1, and then is paid 0 or 1 according as to whether the sum of the chosen numbers is even or odd. It is a Nash equilibrium for each player to play 0 or 1 with $\frac{1}{2}-\frac{1}{2}$ probabilities. One explicit scenario for this is that each player decides on his action in accordance with a completely private coin toss. A second scenario is like the first, except that 2 observes 1's toss. In that case, 2 believes with certainty that 1 will do what he in fact will do, but this belief is not shared by 3, who still ascribes probabilities $\frac{1}{2}-\frac{1}{2}$ to 1's actions. A third scenario is like the second, except that 2's observation of 1's toss is transmitted through a noisy channel; in that case, 1 will not know what 2 believes about 1's action. (In each case, the scenario is assumed common knowledge.) The point is that for practical purposes, all three scenarios are equivalent, since the additional information never causes any player to change his action.

yield any player more than 1. But there is a correlated equilibrium yielding (2, 2, 2): Player 3 picks the middle matrix with certainty, while 1 and 2 pick (top, left) with probability $\frac{1}{2}$, and (bottom, right) with probability $\frac{1}{2}$. This does *not* imply collusion between 1 and 2; it simply means that it is common knowledge that 3 has no idea what they will do, but does think with certainty that they will either play (top, left) or (bottom, right). This, of course, is an extreme situation; it will perhaps be more common to find situations where in 3's estimate, the actions of 1 and of 2 are neither perfectly correlated nor independent. Again, such situations are plentiful, but cannot be described by means of mixed strategies.

7. THE LITERATURE

The idea presented here was briefly outlined in Appendix 4 of Aumann (1981). Kadane and Larkey (1982) also expressed the idea that the theory of games should be consistent with Bayesian decision theory. But they ignored the fact that a rational player must take into account how other players reason about him, and concluded that under Bayesian theory, nothing at all could be said about what each player thinks others would do. One must, they said, turn to disciplines such as psychology to gain insight into how rational agents should or would play a game.

More recently, Bernheim (1984) and Pearce (1984) have independently developed the theory of rationalizability. This is related to the notion of subjective correlated equilibrium mentioned in Section 5 above; a rationalizable outcome is much like one that occurs with positive probability in a subjective correlated equilibrium (see Tan and Werlang, 1984, and Brandenburger and Dekel, 1985). Nevertheless, there are differences; for example, rationalizability as currently defined does not permit one player to perceive the strategies of other players as correlated (see Section 6 above).

Institute of Mathematics, The Hebrew University, 91904 Jerusalem, Israel

Manuscript received July, 1985; final revision received April, 1986.

REFERENCES

- AUMANN, R. J. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67-96.
- (1976): "Agreeing to Disagree," *Annals of Statistics*, 4, 1236-1239.
- (1981): "Survey of Repeated Games," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Mannheim, Wien, Zurich: Wissenschaftsverlag, Bibliographisches Institut, 11-42.
- BERNHEIM, B. D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007-1028.
- BRANDENBURGER, A., AND E. DEKEL (1985): "Rationalizability and Correlated Equilibria," unpublished preprint, Stanford University and Harvard University.
- HARRISON, J. M., AND D. M. KREPS (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323-336.
- HARSANYI, J. C. (1967): "Games with Incomplete Information Played by Bayesian Players," Parts I, II, III, *Management Science*, 14, 159-182, 320-334, 486-502.

- HARSANYI, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1-23.
- (1977): *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge: Cambridge University Press.
- KADANE, J. B., AND P. D. LARKEY (1982): "Subjective Probability and the Theory of Games," *Management Science*, 28, 113-120.
- KUHN, H. W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games II*, ed. by H. W. Kuhn and A. W. Tucker. Princeton: Princeton University Press, pp. 193-216.
- KALAI, E., AND D. SAMET (1984): "Persistent Equilibria in Strategic Games," *International Journal of Game Theory*, 13, 129-144.
- KOHLBERG, E., AND J. F. MERTENS (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1037.
- KREPS, D., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863-894.
- LEWIS, D. K. (1969): *Convention*. Cambridge: Harvard University Press.
- MERTENS, J. F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1-29.
- MOULIN, H., AND J. P. VIAL (1978): "Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon," *International Journal of Game Theory*, 7, 201-221.
- MYERSON, R. B. (1978): "Refinement of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7, 73-80.
- NASH, J. F. (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286-295.
- VON NEUMANN, J., AND O. MORGENSTERN (1944): *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- PEARCE, D. G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029-1050.
- SAVAGE, L. J. (1954): *The Foundations of Statistics*. New York: Wiley.
- SELTEN, R. (1975): "Re-examination of the Perfectness Concept For Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25-55.
- TAN, T. C.-C., AND S. R. DA COSTA WERLANG (1984): "The Bayesian Foundations of Rationalizable Strategic Behavior and Nash Equilibrium Behavior," unpublished preprint, University of Chicago.