Coherent Behavior in Noncooperative Games*

ROBERT F. NAU AND KEVIN F. MCCARDLE

*The Fuqua School of Business, Duke University,
Durham, North Carolina 27706*

Received May 20, 1988; revised February 24, 1989

A new concept of mutually expected rationality in noncooperative games is proposed: joint coherence. This is an extension of the "no arbitrage opportunities" axiom that underlies subjective probability theory and a variety of economic models. It sheds light on the controversy over the strategies that can reasonably be recommended to or expected to arise among Bayesian rational players. Joint coherence is shown to support Aumann's position in favor of objective correlated equilibrium, although the common prior assumption is weakened and viewed as a theorem rather than an axiom. An elementary proof of the existence of correlated equilibria is given, and relationships with other solution concepts (Nash equilibrium, independent and correlated rationalizability) are also discussed. *Journal of Economic Literature* Classification Numbers: 021, 022, 026, 213. © 1990 Academic Press, Inc.

1. INTRODUCTION

The central problem in noncooperative game theory is "how to translate the intuitive assumption of mutually expected rationality into mathematically precise behavioral terms (solution concepts)" (Harsanyi [20]). A continuing proliferation of solution concepts in the literature suggests that a consensus on how to perform this translation has not yet been reached. Nash's [29] original concept of an equilibrium among independently randomized strategies is usually taken as a starting point, but it is held to be too weak in some circumstances (leading to proposals of refined forms of equilibrium such as subgame perfect, trembling-hand perfect, sequential, proper, strategically stable, divine, etc.) and too strong in others (leading to proposed coarsenings such as correlated equilibrium, rationalizability, correlated rationalizability, etc.).

In the last few years, "Bayesian rational" axiomatic foundations have been constructed for many of these solution concepts (Tan and Werlang

* The authors gratefully acknowledge the support of the Business Associates Fund at the Fuqua School of Business and Control Data Corporation. We are also indebted to Professor Robert Aumann and two anonymous referees for their helpful comments.

[36], Brandenburger and Dekel [8, 9], Bernheim [6], Aumann [3]). These axiom systems extend Savage's [33] axioms of Bayesian rational individual behavior to include consistency conditions on infinite hierarchies of reciprocal knowledge and beliefs and/or the assumption of a common prior distribution over an exogenous state space. While these constructs are mathematically interesting and have helped to elucidate the relationships among different solution concepts, they have not settled the argument as to which concept is "the" expression of Bayesian rationality in games, nor are any of them entirely satisfactory from a normative viewpoint.

The treatment of prior distributions or even hierarchies of distributions as primitive elements in a theory of rational strategic behavior violates an important distinction between behavior and belief that is observed in the work of Bayesian decision theorists such as Savage and de Finetti. There, rationality is defined by axioms on observable events, consequences, and actions (assertions of preference, acceptance of gambles, etc.). The framing of axioms in terms of observable quantities is not arbitrary or unimportant: it ensures that they can be independently verified or enforced, at least in principle. Theorems are then derived to say that behavior is rational if and only if it is supported by consistent belief and preference structures (probability distributions and utility functions). This demonstration of a duality between external standards of behavior and internal representations of beliefs and preferences is what gives normative force to the Bayesian model.

In this paper we present a definition of rational behavior in non-cooperative games that is, we believe, more consistent in spirit with the decision-theoretic view of individual rationality, and thereby addresses some of the criticisms raised by Kadane and Larkey [22, 23] concerning the schism between subjective probability theory and game theory. The central idea is an extension to the multi-player setting of de Finetti's [12, 13] operational criterion of rationality, namely that choices under uncertainty should be *coherent* in the sense of not presenting opportunities for arbitrage ("Dutch books") to an outside observer who serves as betting opponent. That is, a rational individual should not let himself be used as a money pump.

A noncooperative game is a joint decision problem in which each player's strategies are lotteries whose payoffs depend on the uncertain actions of his opponents. The player's choice of strategy implies a preference for the chosen strategy over any alternative strategy, which can be given the operational interpretation that he will accept a gamble in which the payoffs of any other available strategy are exchanged for the payoffs of the chosen one. Gambles accepted in this way may present arbitrage opportunities to an outside observer, and we suggest that a primitive characterization of rational behavior in the game can be given in terms of the avoidance of such arbitrage opportunities. This is analogous

to the use of no-arbitrage-opportunity assumptions elsewhere in economics and finance (e.g., Varian [37]) to characterize the collective behavior of rational agents in commodities and securities markets. Here, the "securities" being traded are, literally, shares of the players' prospects in the game.

A priori, we can distinguish several levels of assumptions about arbitrage opportunities the players might wish to avoid. In ascending order of strictness:

(I) Individual coherence: Each player should avoid strategies that expose him individually to arbitrage.

(II) Common knowledge of (I): Each player should avoid strategies that *would* expose him individually to arbitrage if the strategies already being avoided by his opponents under (I) and (II) were deleted from consideration.

A joint strategy that violates (II) but not (I) is one in which some player is effectively betting on his opponent to behave incoherently, or betting on his opponent to bet on him to behave incoherently, and so on. This creates an arbitrage opportunity requiring simultaneous transactions with two or more players: from the observer's viewpoint, it is arbitrage against the group rather than a single player. There may also be opportunities for arbitrage against the group other than those characterized by (II). For example, such an opportunity may exist if the players choose strategies in which each is effectively betting on "outguessing" his opponents, and knows that they are betting on outguessing him, and so on. Therefore, we add:

(III) Inductive extrapolation of (II): The players should avoid all strategies that expose the group to arbitrage.

A strategy that satisfies (III) will be defined to be *jointly coherent*. Players who subscribe to the standard of joint coherence are those who do not let themselves be used *collectively* as a money pump.

Our main result is that a strategy is jointly coherent if and only if it occurs with positive probability in some correlated equilibrium.¹ Thus, correlated equilibrium is the dual of joint coherence, in the same sense that a supporting probability distribution is the dual of a system of coherent preferences in a nonstrategic decision problem, and a system of equilibrium prices is dual to the absence of arbitrage opportunities in a competitive market. This result lends support to Aumann's [3] contention that correlated equilibrium is the (strongest) natural expression of Bayesian

¹ A self-enforcing specification of randomized strategies that, unlike Nash equilibrium strategies, may be correlated between players (Aumann [2]).

rationality in noncooperative games. However, Aumann's formulation depends on, among other things, the assumption of a common prior distribution over an exogenous state space to which the players' strategies are pegged. In our formulation, the avoidance of all arbitrage opportunities against the group is viewed as an axiom from which it follows as a theorem that the players should behave *as if* they held *some* such common prior. It will also be seen that strategies satisfying (II) are those participating in *a posteriori* equilibria (Aumann [2], Brandenburger and Dekel [10]), a generalization of correlated equilibria in which the players may have different priors. Thus, the difference in implications between (II) and (III) is the existence of a supporting common prior.

2. JOINT COHERENCE

We consider the case of finite games among players whose utilities for consequences are known—or, equivalently, whose utility for money is linear. Our starting point for characterizing rational play will be a set of axioms for rational individual choice under uncertainty with respect to a finite set of states of nature. As is well known, such axioms can be formulated either in terms of acceptance of gambles or in terms of binary preferences. We will emphasize the former, while pointing out links with the latter. Let S be a finite set of states of nature, and let G , G_1 , G_2 , etc., denote gambles (lotteries) on S .

DEFINITIONS. A gamble G is *acceptable* for an individual if he agrees to a transaction in which he will receive the payoff $\beta G(s)$ when state $s \in S$ obtains, where $\beta \in \{0, 1\}$ is to be chosen by an opponent prior to the realization of s . G is *conditionally acceptable* given the occurrence of an event H if $1_H G$ is acceptable, where 1_H denotes the indicator function of H .² G_1 *dominates* G_2 if $G_1(s) \geq G_2(s) \forall s \in S$.

Let $\mathcal{A}$ denote the set of all acceptable gambles for a particular individual. The following axioms are considered to characterize rational gamble acceptance:

- A1 (Dominance). G dominates $0 \Rightarrow G \in \mathcal{A}$
- A2 (Linearity). $G \in \mathcal{A} \Rightarrow \alpha G \in \mathcal{A} \quad \forall \alpha \geq 0$
- A3 (Additivity). $G_1, G_2 \in \mathcal{A} \Rightarrow G_1 + G_2 \in \mathcal{A}$
- A4 (No arbitrage opportunities). $-1 \notin \mathcal{A}$.

² Addition and multiplication are defined pointwise: if G_1 and G_2 are gambles and α_1 and α_2 are scalars, then $\alpha_1 G_1 + \alpha_2 G_2$ is the gamble whose payoff in state s is $\alpha_1 G_1(s) + \alpha_2 G_2(s)$. $1_H G$ yields the payoff $G(s)$ if $s \in H$, otherwise 0. When used in expressions referring to gambles, the numbers 1, 0, and -1 denote constant gambles having these values. A positive payoff is considered as a gain for the individual, a negative payoff is a loss.

A set of acceptable gambles satisfying A4 given A1–A3 is one in which the individual neither explicitly nor implicitly accepts a uniformly negative payoff: such a set of gambles is defined to be *coherent*.

Equivalently, define a binary preference relation “ $\succsim$ ” on lotteries, subject to the following axioms:

B1 (Dominance). G_1 dominates $G_2 \Rightarrow G_1 \succsim G_2$

B2 (Linearity). $G_1 \succsim G_2 \Rightarrow \alpha G_1 \succsim \alpha G_2 \forall \alpha \geq 0$

B3 (Transitivity). $G_1 \succsim G_2, G_2 \succsim G_3 \Rightarrow G_1 \succsim G_3$

B4 (Independence or cancellation). $G_1 \succsim G_2 \Rightarrow G_1 + G \succsim G_2 + G \forall G$

B5 (No arbitrage opportunities). $-1 \not\succsim 0$.

With the identification $G \in \mathcal{A} \Leftrightarrow G \succsim 0$, the A and B axioms imply each other (de Finetti [12], Buehler [11], Walley [38]).

The well-known coherence theorem (de Finetti [13], Buehler [11]) states that a set of acceptable gambles [preferences] is coherent if and only if there exists some probability distribution on S that assigns nonnegative expected value to every acceptable gamble. This result follows from a linear duality argument. Thus, an individual must act *as if* his preferences were derived from an internal probability distribution if he is to avoid being used as a money pump. This is generally considered to be the most persuasive normative argument in favor of the existence of subjective probabilities.

Unlike de Finetti, we have not made any assumption about the completeness of preferences over a continuous spectrum of gambles—that is, the existence of exact indifference points. For example, we have not assumed that, for any gamble G , either G or $-G$ must be acceptable, nor that there exists a constant p such that both $G - p$ and $p - G$ are acceptable. Such assumptions distinguish the conventional theory of sharp (point-valued) probabilities and expectations from theories of lower and upper (interval-valued) probabilities and expectations (Koopman [24], Smith [34], Good [17], Suppes [35], Williams [39], Walley [38]). Arguably, they are excess baggage in a decision-making context with only finitely many states and decisions, such as an n -person matrix game. More importantly, there appears to be a growing consensus among Bayesians that the assumption of completeness on infinite sets is not merely a harmless idealization; that subjective probabilities are in practice ambiguous or numerically indeterminate; and that decision models and statistical methods should accordingly be robust against this indeterminacy.³ If the

³ See, for example, Giron and Rios [16], Berger [4], Leamer [25], or Bewley [7]. The completeness assumption with respect to utility measurement has also been questioned, e.g., by Aumann [1].

idea of sharply defined first-order subjective probabilities cannot be taken seriously, the credibility of game-theoretic results built on infinite hierarchies of subjective distributions is obviously compromised. Therefore, one goal of the present paper is to establish a foundation for discussing rationality in finite games that does not depend on the completeness assumption.

We now extend the axiom system given above to a strategic setting. Consider a noncooperative game of complete information defined by a set $N = \{1, \dots, n\}$ of players; a finite set S of joint strategies, with S_i denoting the strategy set of player i , and S_{-i} denoting the joint strategy set of all players other than i ; and a function $u_i(s_i, s_{-i})$ denoting the payoff to player i when he plays $s_i \in S_i$ and his opponents play $s_{-i} \in S_{-i}$. To characterize mutually expected rationality in the game, we will adopt the viewpoint of an observer of the game ("she") who serves as the common betting opponent of the players. That is, rather than focusing on the introspective processes used by the players in formulating their strategies, we will consider only the implications of their overt behavior for the observer. The observer's view is the lowest common denominator among views of the game that the players themselves can adopt, and this fact is common knowledge. The definition of an "acceptable" gamble given above is henceforth reinterpreted from the observer's viewpoint: an acceptable gamble is one that has been made available to the observer by any player. Provided that the currency in which gambles are expressed is transferable between the players and the observer, the identity of the player who accepts a particular gamble is unimportant: as far as the observer is concerned, the gamble has been accepted by the group.

We embed the game in a larger universe which includes the observer by assuming the players accept conditional gambles which "reaffirm" their choices of strategies in the game. For player i , whose state space is S_{-i} , this means that his choice of strategy j is assumed to imply acceptance of a gamble whose payoff in state s_{-i} is $u_i(j, s_{-i}) - u_i(k, s_{-i})$, for every $k \in S_i$, $k \neq j$. By de Finetti's theorem, the acceptance of these gambles is individually coherent if and only if there is a probability distribution on S_{-i} under which each of them has nonnegative expected value—that is, a distribution on S_{-i} against which j is a "best response." Therefore, Assumption (I) from Section 1 implies that each player should choose a strategy which is a best response to some distribution over the actions of his opponents in the original game. Assumption (II) is that each player should choose a strategy that is individually coherent in every subgame that remains following the iterative deletion of individually incoherent strategies. By the preceding result, these are the strategies that remain following the iterative deletion of all players' strategies that are not best responses to any distributions over their opponents' strategies. Such

strategies are, by definition, those that satisfy the condition of *correlated rationalizability* (Brandenburger and Dekel [8, 10]). Thus, assumption (II) is the dual definition of correlated rationalizability. Brandenburger and Dekel also show that the correlated rationalizable strategies are those which participate in *a posteriori equilibria* (Aumann [2]), subjective correlated equilibria that are self-enforcing ex post as well as ex ante.

To characterize the strategies satisfying assumption (III), the gambles implicitly accepted by the players must be rewritten in the form of unconditional gambles on the observer's state space, S . Let u_{ij} denote a gamble whose payoff in state $s = (s_i, s_{-i})$ is the payoff that player i would have received by playing strategy $j \in S_i$ given that the other players had chosen s_{-i} , regardless of the actual value of s_i :

$$u_{ij}(s) \equiv u_i(j, s_{-i}).$$

Also, let $1_{ij}(s)$ denote the indicator function for player i choosing strategy j :

$$1_{ij}(s) \equiv \begin{cases} 1 & \text{if } s_i = j \\ 0 & \text{otherwise.} \end{cases}$$

We now assume

A5 (Implicitly accepted gambles). $(u_{ij} - u_{ik}) 1_{ij} \in \mathcal{A}$ for all $i \in N$ and all $j, k \in S_i$.

In other words, conditional on the "event" that player i chooses strategy j , we assume the acceptability [for him] of a gamble in which the payoffs of any other strategy k are exchanged for those of j .

From the observer's viewpoint, A5 together with A2 and A3 implies that the players will accept any nonnegative linear combination of the gambles $\{(u_{ij} - u_{ik}) 1_{ij}\}$. The observer's decision problem is to determine which of these she will enforce. (A1 implies that the players will also accept any gamble that dominates another acceptable gamble, but these are inefficient for the observer.) Accordingly, she chooses a nonnegative coefficient α_{ijk} for the acceptable gamble $(u_{ij} - u_{ik}) 1_{ij}$ for every $i \in N$ and every $j, k \in S_i$. For concreteness, we might suppose that these coefficients are deposited with the players in sealed envelopes to be opened conditional on their individual choices of strategies. Thus, player i receives an envelope labeled "to be opened in the event you choose strategy j ," for each $j \in S_i$. Inside this envelope are the coefficients α_{ijk} for all $k \in S_i$. After they have chosen their strategies in the game, the players open the appropriate envelopes and the gambles are settled.

Let α denote the triply indexed vector whose (i, j, k) th element is α_{ijk} , and let A denote the matrix whose columns are the payoff vectors for the acceptable gambles defined in A5. That is, A is the matrix whose columns

are indexed by (i, j, k) , whose rows are indexed by s , and whose element in column (i, j, k) and row s is

$$A_{ijk}(s) \equiv [u_{ij}(s) - u_{jk}(s)] 1_{ij}(s). \quad (1)$$

The s th row of A will be denoted $A(s)$. The vector of payoffs from observer to players under all states may now be expressed as $A\alpha$, and the specific payoff in state s is $A(s)\alpha$.⁴ Then, given A1–A3 and A5, $G \in \mathcal{A}$ if G dominates $A\alpha$ for some nonnegative α . Formally, the matrix A describes the observer's view of the game, and so may be considered to summarize the features of the game's structure that are common knowledge among the players. Since the elements of A depend only on differences in utility between strategies of the same player, and since the scaling of the elements of α is arbitrary, it follows that the payoff structure is unaffected by non-negative linear transformations of the individual utility functions.

To complete the characterization of strategies satisfying (III), it will be necessary to modify our original no-arbitrage-opportunity axiom (A4). In the discussion of individual coherence, we held the individual responsible for the set of gambles he accepted, not the specific outcome which resulted, and the no-arbitrage-opportunity axiom was therefore a restriction on the set of acceptable gambles. Here the situation is reversed: we hold the players responsible for the outcome (i.e., the joint strategy that is played) rather than the set of gambles that were accepted, so the no-arbitrage-opportunity axiom must refer to the outcome of the game. It will be proved later (as Proposition 1) that the acceptance of the gambles defined in A5 can never imply, via A1–A3, the acceptance of a uniformly negative gamble in violation of A4. However, these assumptions will generally imply the acceptance of some *semi*-negative gambles (negative in some outcomes of the game, zero in others), and the observer may be expected to enforce these gambles since they carry no risk for her. If the players choose an outcome that is one of the losing outcomes in an acceptable semi-negative gamble, they will have “deliberately” allowed arbitrage to be achieved against them. Our modified no-arbitrage-opportunity assumption is precisely that this should not happen. Let d denote the outcome of the game (the strategy actually played), and let 1_d denote the indicator

⁴ In a discussion of refinements of correlated equilibrium, Myerson [27] introduces a vector α whose (i, j, k) element is called a “shadow price” for the incentive constraint that player i should not expect to gain by using strategy k instead of j when j is recommended to him. The quantity we have defined as the total payoff from observer to players is described by Myerson as the “aggregate incentive value of s for [the set of all players] with respect to α .” It will be seen that α ultimately plays the same role in our analysis. The difference is that we have a prior interpretation of it as a vector of coefficients for gambles between the players and an observer. From this and the joint coherence assumption, the constraints defining a correlated equilibrium will be derived.

function of d on the set S (i.e., $1_d(s) = 1$ if $s = d$, zero otherwise). Then we assume:

A4' (No arbitrage opportunities). $-1_d \notin \mathcal{A}$.

DEFINITION. An outcome d of the game is *jointly coherent* if it satisfies A4' given A1–A3 and A5.

We will say that α *achieves arbitrage* against an outcome d if α is non-negative and -1_d dominates $A\alpha$ or some positive multiple thereof—i.e., if $A(s)\alpha \leq 0$ for all $s \in S$, and $A(d)\alpha < 0$. In these terms, the jointly coherent outcomes of the game are those against which arbitrage cannot be achieved. Note that if α and α' achieve arbitrage against distinct strategies d and d' , respectively, then in view of the linearity of the payoff function, $\alpha + \alpha'$ achieves arbitrage against both d and d' .⁵ It follows that there is a single α that achieves arbitrage against all jointly incoherent strategies. Therefore the players might reasonably assume that they will certainly lose money (as a group) to the observer if any jointly incoherent strategy is played.

3. RESULTS

It is intuitively clear that the set of jointly coherent strategies is not empty. If strategy d is played, and player i finds that $\alpha_{ijk} > 0$ for $j = d_i$ and some $k \neq d_i$, this can be interpreted as a judgment by the observer that the player should have chosen k instead of j . This judgment is vindicated if α in fact achieves arbitrage against d . Yet, in a finite game, it cannot be true that, no matter what the players do, some of them should have done otherwise—or else life would be terribly unfair! Formally, we have the following:

PROPOSITION 1. *There is at least one jointly coherent strategy.*

Proof. If all strategies were jointly incoherent, then there would exist α such that $A(s)\alpha < 0$ for all s . Even if this α were revealed to the players in advance, they would be unable to find objectively randomized strategies (independent or correlated) yielding nonnegative expected gain to every player. We will demonstrate a contradiction, namely that for any α there exists an independently randomized strategy for each player yielding an

⁵ If α and α' achieve arbitrage against d and d' , respectively, then $A\alpha$ and $A\alpha'$ are semi-negative vectors with strictly negative elements in positions d and d' , respectively. Since $A(\alpha + \alpha') = A\alpha + A\alpha'$, it follows that $A(\alpha + \alpha')$ is semi-negative with strictly negative elements in both positions d and d' ; i.e., $\alpha + \alpha'$ achieves arbitrage against both d and d' .

expected gain of exactly zero. Consider the situation that α presents to player i : each strategy available to him defines a lottery over the uncertain actions of his opponents, and α_{ijk} is the multiple of the lottery defined by strategy k that the observer will ask him to give up in return for the same multiple of the lottery defined by strategy j if he should choose strategy j in the game. Without loss of generality, assume $\sum_{k:k \neq j} \alpha_{ijk} < 1$, and define $\alpha_{ijj} = 1 - \sum_{k:k \neq j} \alpha_{ijk}$, whence $\sum_k \alpha_{ijk} = 1$ for all $j \in S_i$. (The exchange of a lottery for itself yields no net transaction, so α_{ijj} is arbitrary.) Suppose now that player i employs an independently randomized strategy characterized by a probability vector p . Then the *expected* multiple of the lottery defined by strategy k that he will give up is $\sum_j \alpha_{ijk} p_j$, and the expected multiple of it that he will receive is simply p_k ($= \sum_j \alpha_{ikj} p_k$), independent of the actions of his opponents. The player's net expected gain or loss in multiples of the lottery defined by strategy k will be zero if these are equal. That is, if p satisfies

$$\sum_j \alpha_{ijk} p_j = p_k \quad \forall k \in S_i, \quad (2)$$

his net expected transaction in every strategy will be zero, and hence his overall expected gain will be zero regardless of the strategies chosen by his opponents. It remains to show that such a distribution p always exists. Let M be the square matrix whose element in row k and column j is α_{ijk} . Note that M is a stochastic matrix (i.e., nonnegative with columns summing to unity) with positive diagonal elements; hence it may be considered as the transition matrix for a finite, aperiodic Markov chain. The system of equations (2) can now be written as $Mp = p$, which is seen to be the equation defining p as a stationary distribution of the Markov chain whose transition matrix is M . As is well known, every finite, aperiodic Markov chain has a stationary distribution, which can be constructed as $p = \lim_{m \rightarrow \infty} M^m p^*$ for an arbitrary initial distribution p^* .⁶ ■

We now come to our main result, namely that the existence of a correlated equilibrium distribution supporting the strategy actually played

⁶ The Markov chain argument has the following intuitive interpretation: player i begins with an initial distribution p^* and then observes α . For each j and k he interprets α_{ijk} as a fraction of the time that (in the observer's judgment) he should have played k instead of j , and he attempts to reallocate probability mass accordingly. Thus, an amount $\alpha_{ijk} p_j^*$ of probability mass is shifted from strategy j to strategy k , for every j and k , yielding the new distribution Mp^* . He then repeats this process on the new distribution, obtaining $M^2 p^*$, and so on. As he continues to make transitions to new distributions in this way in response to the same set of judgments by the observer (i.e., the same α), he eventually achieves an equilibrium in which the net inflow of probability mass into any strategy is equal to the net outflow. His distribution now "agrees" with the observer's judgments, and there is no net transaction.

is the dual of the requirement that it be jointly coherent, in the same sense that the existence of a probability distribution supporting a set of acceptable gambles is dual to the requirement that it be individually coherent. It is based on a separating-hyperplane lemma (Gale [15 Theorem 2.8]) that is a variant of Farkas' lemma, the basis of the duality theorem of linear programming:

LEMMA. *Exactly one of the following two systems of linear inequalities has a solution:*

$$\begin{aligned} \text{(i)} \quad & Ax \leq b, \quad x \geq 0, \\ \text{(ii)} \quad & \pi' A \geq 0, \quad \pi' b < 0, \quad \pi \geq 0, \end{aligned}$$

where, without loss of generality, the vector π satisfying (ii) is scaled so that it is a probability distribution.

The coherence theorem for noncooperative games now follows:

PROPOSITION 2. *A strategy is jointly coherent if and only if there exists a correlated equilibrium distribution in which it has positive probability.*

Proof. A strategy d is jointly coherent if and only if there does not exist a nonnegative vector α satisfying $A\alpha \leq b$, where A is the matrix defined in (1) and b is a vector whose d th element is negative and whose remaining elements are zero. By the lemma, this is true if and only if (ii) has a solution for this matrix A and vector b . Such a solution is a distribution π satisfying the system of inequalities

$$\begin{aligned} \sum_{s \in S} \pi(s) [u_{ij}(s) - u_{ik}(s)] 1_{ij}(s) &\geq 0 \quad \forall i \in N, j \in S_i, k \in S_i, \\ \pi(d) &> 0. \end{aligned}$$

This is merely the system of inequalities defining a correlated equilibrium distribution⁷ that assigns positive probability to d . Conversely, if there

⁷ Letting P_π denote the probability measure on subsets of S induced by π , the system of inequalities implies that either $P_\pi(s_i = j) = 0$ or else, on dividing through by $P_\pi(s_i = j)$, we obtain

$$\sum_{s_{-i} \in S_{-i}} P_\pi(s_{-i} | s_i = j) (u_i(j, s_{-i}) - u_i(k, s_{-i})) \geq 0 \quad \forall i \in N, \forall j, k \in S_i.$$

Hence, if π is a commonly known joint distribution for generating strategy recommendations in which strategy j of player i occurs with positive probability, then upon receiving the recommendation $s_i = j$ and updating his distribution over his opponent's strategies to $P_\pi(s_{-i} | s_i = j)$, he finds that j has at least as great an expected payoff as any alternative strategy k —i.e., the recommendation is self-enforcing.

exists a correlated equilibrium distribution assigning positive probability to d , it constitutes a solution to system (ii), ruling out the existence of a solution to system (i)—i.e., ruling out the achievement of arbitrage against d . ■

From Propositions 1 and 2, it follows that the set of correlated equilibria is nonempty. Of course this is also guaranteed by the existence of Nash equilibria, which are special cases of correlated equilibria. But, since correlated equilibria are “computationally simpler objects,” it is of interest to note that their existence can be established by more elementary methods than those generally used to prove the existence of Nash equilibria.^{8,9}

The correlated equilibrium distribution supporting a jointly coherent outcome of the game will in general not be unique—which is to say, the players’ beliefs may be to some extent indeterminate—unless an additional assumption is made concerning the completeness of preferences among gambles. For the reasons noted earlier, we do not feel that such an assumption is either necessary or desirable in the context of finite games. However, it can be easily incorporated into our framework as follows. In a first-order description of the game, player i is uncertain about the joint strategy s_{-i} that will be played by his opponents. If he is required to reveal an exact probability distribution on S_{-i} for purposes of betting, we must be careful that this revelation does not perturb the game. For example, he may have a prior distribution that will be revised on receipt of private information before the start of play. In this case, he will not wish to gamble with the observer according to his prior, since these gambles are to be conditioned on the strategy that he plays in light of his posterior—but neither will he wish to reveal his actual posterior, since this would divulge his information (and perhaps his strategy selection). We should therefore restrict ourselves to asking for a set of *conditional* distributions, given the strategies he might play.¹⁰

⁸ The proof of Proposition 1 depends on the existence of a stationary distribution for a finite Markov chain, which is a fixed point of a linear mapping of the simplex into itself. This is a weaker result than the existence of a fixed point for an arbitrary continuous mapping of a compact set into itself, on which the Nash proof is commonly based. Of course, the existence of Nash equilibria could have been used directly in the proof of Proposition 1: any Nash equilibrium will yield nonnegative expected gain for every player against any α chosen by the observer.

⁹ Since the first draft of this paper was completed in December 1987, it has come to our attention that another elementary proof of the existence of correlated equilibria was developed independently and somewhat earlier by Hart and Schmeidler [21]. Their proof is structurally similar, insofar as it introduces an outside observer (who plays a 2-person game against the set of all players), invokes linear duality (via the minimax theorem), and relies on a technical lemma similar to our Proposition 1.

¹⁰ Note that by asking for unconditional distributions, we would implicitly be imposing the Nash solution concept.

Let h denote an index for the elements of S_{-i} , and let $1_{-ih}(s)$ denote the indicator function for the event that $s_{-i} = h$:

$$1_{-ih}(s) \equiv \begin{cases} 1 & \text{if } s_{-i} = h \\ 0 & \text{otherwise.} \end{cases}$$

Following de Finetti [12, 13], we can now introduce:

A6 (Completeness). $\exists p_{ijh}$ such that both $(1_{-ih} - p_{ijh}) 1_{ij} \in \mathcal{A}$ and $(p_{ijh} - 1_{-ih}) 1_{ij} \in \mathcal{A} \quad \forall i \in N, j \in S_i, h \in S_{-i}$.

Here, the constant p_{ijh} is, by definition, player i 's subjective conditional probability for his opponents choosing $s_{-i} = h$ given that he has chosen $s_i = j$, and this is to be articulated prior to the start of play. The observer may now choose a coefficient β_{ijh} , positive or negative, for the acceptable gamble $(1_{-ih} - p_{ijh}) 1_{ij}$ for every $i \in N, j \in S_i$, and $h \in S_{-i}$. Let p and β denote the triply indexed vectors whose (i, j, h) th elements are p_{ijh} and β_{ijh} , respectively; and let B be the matrix whose (i, j, h) th column is the payoff vector $(1_{-ih} - p_{ijh}) 1_{ij}$ introduced in A6. (Note that B depends on p). Then the total payoff from observer to players due to all gambles is $A\alpha + B\beta$. Joint coherence can now be defined for a strategy-probability pair (d, p) , and the obvious generalization of Proposition 2 is that (d, p) is jointly coherent if and only if there exists a correlated equilibrium distribution π in which d has positive probability *and* in which, for all (i, j, h) , either $P_\pi(s_i = j) = 0$ or else $P_\pi(s_{-i} = h | s_i = j) = p_{ijh}$.

Technically there is no loss of generality in making the completeness assumption, insofar as it does not change the set of correlated equilibria supporting the jointly coherent strategies in the original game. But if it were enforced in practice, it would change the dynamics of the game by increasing the demands on the players for information processing and pre-play communication. They would have to reach agreement on which correlated equilibrium they were playing, even if this were unimportant in the original game, in order to avoid arbitrage opportunities due merely to inconsistent beliefs. They might well expend more effort in articulating and coordinating their choices of distributions than in determining their choices of strategies.

4. COMPARISON WITH AUMANN'S FORMULATION

The results of section 3 provide support for Aumann's [3] contention that correlated equilibrium is the strongest natural expression of Bayesian rationality in noncooperative games—with some subtle but important qualifications. Aumann assumes that there exists an exogenous set Ω of

states of nature on which the players have a common prior distribution, and that player i 's strategy is a commonly known function of ω that is measurable with respect to a commonly-known partition $\mathcal{P}_i$. This arrangement is then shown to be Bayesian rational (payoff-maximizing for each player given his information) if and only if the distribution of outcomes of the game is a correlated equilibrium distribution.

Aumann justifies the common prior assumption (CPA) by appeal to the arguments originally presented by Harsanyi [19] and the large body of economic literature in which the CPA has subsequently appeared (rational expectations, signalling, etc.). But the no-arbitrage-opportunity assumption is also a familiar tool in economic analysis (e.g., in deriving the existence of competitive equilibria), and we have shown that if it is applied directly to the original game with a minimum of other structural assumptions, it follows immediately that the players must choose their strategies *as if* they had implemented *some* correlated equilibrium. The distinctions "as if" and "some" are important: they allow that the game may be played rationally without the players inquiring more deeply into each other's beliefs (or even their own) than is minimally necessary to determine strategy selections from within their respective finite sets. Indeed, our framework suggests a somewhat different concept of a "solution" of the game: a rational solution is any arrangement among the players that leads to a jointly coherent outcome, with or without a precise articulation of probabilities.

For the sake of comparison, we can recast our results in Aumann's formulation with the exogenous state space, and infer both the common prior and the correlated equilibrium distribution. Suppose that the players commit themselves in advance to letting their strategy choices in the game be determined by commonly known functions of their information about ω . Let $\mathcal{P}_{ij}$ denote the j th element of player i 's information partition, and let J_i denote the set of index numbers for elements of $\mathcal{P}_i$. Let $s(\omega) = (s_1(\omega), \dots, s_n(\omega))$ be a function that maps states into strategies, where $s_i(\omega)$ is measurable with respect to $\mathcal{P}_i$ for all i . It is desired to characterize the conditions under which the players' commitment to the strategy function $s(\omega)$ is rational. As before, we assume that the players' participation in the game implies that certain gambles will be accepted. In this case, we assume that player i accepts a gamble in which, conditional on observing $\omega \in \mathcal{P}_{ij}$, he receives the payoffs of his recommended strategy $s_i(\omega)$ in exchange for those of any other strategy $k \in S_i$, assuming that his opponents have played their recommended strategies. This is formalized as follows, where $\mathcal{A}$ now denotes the set of acceptable gambles on Ω , and 1_{ij} now denotes the indicator function for the event that $\omega \in \mathcal{P}_{ij}$:

A5' (Implicitly accepted gambles). $[u_i(s(\omega)) - u_i(k, s_{-i}(\omega))] 1_{ij}(\omega) \in \mathcal{A} \quad \forall i \in N, j \in J_i, k \in S_i$.

Since the players are now gambling with respect to states of nature they do not control, the relevant no-arbitrage-opportunity assumption is now A4 rather than A4': they should not accept a gamble with payoffs uniformly negative in ω .

DEFINITION. The strategy function $s(\omega)$ is *jointly coherent* if A4 is satisfied given A1 A3 and A5'.

PROPOSITION 3. $s(\omega)$ is jointly coherent if and only if there exists a distribution on Ω under which:

- (a) Each player's strategy function maximizes his expected payoff given his information, and given that the other players adhere to their strategy functions; and
- (b) The distribution of outcomes in the original game is a correlated equilibrium distribution.

Proof. Let A be the matrix with rows indexed by ω and columns indexed by (i, j, k) whose columns are payoff vectors for the acceptable gambles in A5', and let b be the column vector indexed by ω whose elements are all -1 . Then the strategy functions are jointly coherent if and only if system (i) of the separating-hyperplane lemma has no solution for this A and this b . By the lemma, this is true if and only if system (ii) has a solution, which is a probability distribution π satisfying

$$\sum_{\omega \in \Omega} \pi(\omega) [u_i(s(\omega)) - u_i(k, s_{-i}(\omega))] 1_{ij}(\omega) \geq 0 \quad \forall i \in N, j \in J_i, k \in S_i.$$

Let P_π denote the probability measure on subsets of Ω induced by π . Then this system of inequalities implies that, for all (i, j, k) , either $P_\pi(\omega \in \mathcal{P}_{ij}) = 0$ or else $s_i(\omega)$ yields an expected payoff as great as any other strategy $k \in S_i$, conditional on observing $\omega \in \mathcal{P}_{ij}$. This proves part (a). For part (b), note that since the strategy functions $\{s_i(\omega)\}$ are payoff-maximizing, it suffices for player i to be informed only of the value of $s_i(\omega)$, not the element of his partition in which ω fell: his expected payoff is still maximized by following the recommended strategy. This means that, by definition, the distribution of recommended strategies is a correlated equilibrium distribution. ■

5. EXAMPLES

A strategy d is obviously and trivially jointly incoherent if any player's strategy d_i is individually incoherent, in violation of assumption (I) in

Section 1. The following example shows that a strategy is also jointly incoherent if some players' strategies are rational only on the presumption that others are individually incoherent, in violation of assumption (II).

EXAMPLE 1. The numbers in parentheses denote payoffs to the row player and column player, respectively:

	<i>L</i>	<i>R</i>
<i>T</i>	(2, 2)	(1, 1)
<i>B</i>	(1, 1)	(0, 2)

Here, if column chooses *R*, he is effectively betting on row to behave incoherently and choose *B*, which is strongly dominated by *T*. Let the observer choose $\alpha_{\text{row}, B, T} = 2$, $\alpha_{\text{column}, R, L} = 1$, and $\alpha_{ijk} = 0$ for all other (i, j, k) . She then wins 2 units from row if *B* is played, regardless of column's choice; loses 1 unit to column if *BR* is played; wins 1 unit from column if *TR* is played; and there is no transaction if *TL* is played. This is arbitrage against every strategy except *TL*. In particular, any joint strategy in which column plays *R* is jointly incoherent because in this case either row has behaved incoherently (so that the observer can win more from row than she loses to column), or else column would have done better by playing *L* (so that the observer has no transaction with row but can win something from column).

The following is an example of a game in which some strategies are jointly incoherent due to a violation of assumption (III), and is taken from Bernheim [5]:

EXAMPLE 2. The payoff matrix is

	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	(0, 7)	(2, 5)	(7, 0)
<i>M</i>	(5, 2)	(3, 3)	(5, 2)
<i>B</i>	(7, 0)	(2, 5)	(0, 7)

Here, the unique jointly coherent strategy (and hence the unique Nash equilibrium) is *MC*: if the observer chooses $\alpha_{ijk} = 1$ for all (i, j, k) in which $j \neq 2$ and $k \neq j$, and $\alpha_{ijk} = 0$ otherwise, she wins 3 if both players choose their first or third strategies, and wins 1 if exactly one player chooses his first or third strategy. However, the strategies *T*, *B*, *L*, and *R* would also be permitted under the alternative solution concepts of either independent

or correlated rationalizability: row would play T if he thought column would play R , which column would play if he thought row would play B , which row would play if he thought column would play L , which column would play if he thought row would play T . Since each of these four strategies is a best response to one of the others, none can be the first eliminated. These are examples of strategies that can only be justified by each player believing he is capable of outguessing the other, which can occur if their strategies are known to be pegged to events to which they assign different prior probabilities, as in an a posteriori equilibrium.

In the special case of two-person games, there is no distinction between correlated and independent rationalizability; hence independent rationalizability, like correlated rationalizability, is weaker than joint coherence. However, in some games with three or more players, the set of independently rationalizable strategies may be strictly smaller than the set of jointly coherent strategies, due to the restrictiveness of the independence assumption. Consider the following adaptation of a game presented by Aumann [2]:

EXAMPLE 3. The payoff matrix is

	l		c		r	
	L	R	L	R	L	R
T	(0, 1, 3)	(0, 0, 1)	(2, 2, 2)	(0, 0, 0)	(0, 1, 0)	(0, 0, 1)
B	(1, 1, 1)	(1, 0, 0)	(0, 0, 0)	(2, 2, 2)	(1, 1, 1)	(1, 0, 3)

Here no profile of independently randomized strategies for row and column renders c a best response for matrix. It follows by simple domination arguments that the only Nash strategies are those in which row and column play BL and matrix plays either l or r . B , L , and $\{l, r\}$ are also the only rationalizable individual strategies. However, all strategies are jointly coherent, and an obviously desirable solution is for matrix to play c while row and column choose randomly (from matrix's perspective) between TL and BR with probability at least $\frac{1}{3}$ each. Such a correlated strategy yields a higher *certain* payoff to all players than any Nash strategy, and it is therefore difficult to believe that players who were capable of informing each other that they understood the game and knew each other to be rational would not consider it. For example, the matrix player might simply say to the others: "I am going to play c . You two go off and decide whether to play TL or BR ." It is not necessary for row and column to explicitly randomize their choice, as long as matrix is given no reason to subjectively assign either alternative a probability less than $\frac{1}{3}$. This is a "solution" of the game that leaves some probabilities partly indeterminate.

The preceding example illustrates that a game may have jointly coherent strategies that do not participate in any Nash equilibrium. This can occur even in two-person games, such as:¹¹

EXAMPLE 4. The payoff matrix is

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>A</i>	(0, 0)	(10, 5)	(5, 10)	(6, 6)
<i>B</i>	(5, 10)	(0, 0)	(10, 5)	(6, 6)
<i>C</i>	(10, 5)	(5, 10)	(0, 0)	(6, 6)
<i>D</i>	(6, 6)	(6, 6)	(6, 6)	(7, 7)

Here, the unique Nash equilibrium is *Dd*, yet all strategies are jointly coherent, and there is a Pareto superior correlated equilibrium assigning probability $\frac{1}{6}$ to each of *Ab*, *Ac*, *Ba*, *Bc*, *Ca*, and *Cb*.

The preceding two examples illustrate the general undesirability of assuming statistical independence, since it may eliminate solutions that are otherwise attractive, self-enforcing, and even focal. We therefore share Aumann's [3] puzzlement that Nash equilibrium rather than correlated equilibrium has heretofore been accepted as the "fundamental" solution concept for noncooperative games, and also take issue with Bernheim's [5, 6] and Pearce's [32] emphasis on statistical independence in defense of their alternative concept of rationalizability. The independence assumption concentrates on the fact that the players are free to vary their own strategies independently of what their opponents are doing, while ignoring the fact that they may not always wish to do so, and that they will usually have opportunities for preplay communication—otherwise, they could never arrive at a state of common knowledge concerning the structure of the game and each others' rationality. This remains true even if communication is modeled as part of the formal game: there must still be an informal communication stage in which the formal communication mechanism is selected.

On a deeper level, regardless of whether the players have opportunities for deliberate correlation, the independence assumption is incompatible with a truly subjectivistic view of probability. Events which are physically independent or unrelated may still be perceived as subjectively dependent, which is the basis of de Finetti's [12] concept of *exchangeability* and Harrison's [18] critique of independence assumptions in the context of single-person decision analysis.

¹¹ This example was provided by Professor Aumann, and is based on Moulin and Vial [26].

6. CONCLUSIONS AND FUTURE DIRECTIONS

We have shown that the strategies which avoid arbitrage opportunities against the group of all players, which we designate as jointly coherent, are precisely those strategies supporting the set of correlated equilibria. (A weaker no-arbitrage-opportunity criterion has also been shown to support the set of correlated rationalizable strategies, or a posteriori equilibria.) This result can be viewed as a natural game-theoretic extension of de Finetti's coherence theorem for the individual decision-maker, and it provides a somewhat more elementary defense of correlated equilibrium as a standard of rationality than that of Aumann [3], since the CPA is viewed as a theorem rather than an axiom. It is also somewhat weaker, since a *particular* correlated equilibrium need not always be specified.

Joint coherence is more restrictive than *independent* rationalizability in two-person games, while it may be less restrictive in games with three or more players. We have argued that a general standard of rational behavior in noncooperative games should not exclude the possibility of unobserved communication between the players, since implicit communication is needed to construct the pre-play state of joint knowledge of the structure of the game and of each other's credentials as rational players. Hence, the assumption of statistically independent choices that underlies both Nash equilibrium and independent rationalizability is viewed as excessive.

The assumption of complete information (exactly and commonly known payoffs and utilities) is certainly unrealistic, but there are several directions along which the concept of joint coherence can be extended to more general settings. One approach is to adopt Harsanyi's [19] framework in which players have "types" characterized by different payoffs and/or states of knowledge, from which nature randomly chooses at the first stage of the game, and where players may or may not engage in observable communication. In this context, it can be shown (Nau [31]) that joint coherence supports a correlated generalization of Harsanyi's Bayesian equilibrium concept; with the addition of a formal communication mechanism, it supports the concept of communication equilibrium (Myerson [28], Forges [14]). Here again, the CPA is viewed as a theorem.

Another approach would be to allow utilities to be indeterminate in a nonprobabilistic way. Operationally, the players might assert buying and selling prices (bid-ask spreads, as it were) for each of their game payoffs in terms of a common utility currency consisting of lottery tickets offering objective changes of winning a single desirable prize. Thus, a player's utility for some outcome would be represented by an interval rather than a point. More generally, the players might also assert preferences among different objectively-randomized lotteries over their possible payoffs, establishing

other kinds of inequality bounds on their utilities. (A similar characterization of jointly indeterminate probability and utility in nonstrategic settings has already been developed by Nau [30].) In this context, an arbitrage opportunity would be a system of exchanges yielding the observer a net gain in the common currency and no net transaction in the other payoffs, and joint coherence would presumably require the support of a correlated equilibrium distribution with respect to *some* consistent specification of utility functions. This would be a correlated generalization of the equilibrium concept characterized by Aumann [1] under utility theory without the completeness assumption.

Finally, we have also raised a more basic issue, namely that to adhere to the spirit of the normative theories of Savage and de Finetti, and to achieve robustness against potential violations of the completeness axiom, "Bayesian rationality in games" should be defined in terms of axioms that do not refer explicitly to probabilities and utilities, but rather to observable events, consequences, and actions. We have suggested one path along which this appears to be possible—there may be others.

REFERENCES

1. R. J. AUMANN, Utility theory without the completeness axiom, *Econometrica* **30** (1962), 445–462.
2. R. J. AUMANN, Subjectivity and correlation in randomized games, *J. Math. Econ.* **1** (1974), 67–96.
3. R. J. AUMANN, Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* **55** (1987), 1–18.
4. J. BERGER, The robust Bayesian viewpoint, in "Robustness of Bayesian Analysis" (J. B. Kadane, Ed.), pp. 63–144, North-Holland, Amsterdam, 1984.
5. D. B. BERNHEIM, Rationalizable strategic behavior, *Econometrica* **52** (1984), 1007–1028.
6. D. B. BERNHEIM, Axiomatic characterizations of rational choice in strategic environments, *Scand. J. Econ.* **88** (1986), 473–488.
7. T. F. BEWLEY, "Knightian Decision Theory: Part I," Cowles Foundation Discussion Paper, Yale University, 1986.
8. A. BRANDENBURGER AND E. DEKEL, "Bayesian Rationality in Games," Harvard University Economics Working Paper, 1986.
9. A. BRANDENBURGER AND E. DEKEL, "On an Axiomatic Approach to Refinements of Nash Equilibrium," University of California, Berkeley, Economics Working Paper, 1986.
10. A. BRANDENBURGER AND E. DEKEL, Rationalizability and correlated equilibrium, *Econometrica* **55** (1987), 1391–1402.
11. R. J. BUEHLER, Coherent preferences, *Ann. Statist.* **4** (1976), 1051–1064.
12. B. DE FINETTI, Foresight: Its logical laws, its subjective sources. 1937, translation reprinted in "Studies in Subjective Probability" (H. E. Kyburg and H. E. Smokler, Eds.), pp. 55–118, Krieger, New York, 1980.
13. B. DE FINETTI, "Theory of Probability, Vol. I," Wiley, New York, 1974.
14. F. FORGES, An approach to communication equilibria, *Econometrica* **54** (1986), 1375–1385.

15. D. GALE, "The Theory of Linear Economic Models," McGraw-Hill, New York, 1960.
16. F. J. GIRON AND S. RIOS, Quasi-Bayesian behavior: A more realistic approach to decision making? in "Bayesian Statistics" (J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, Eds.), University Press, Valencia, Spain, 1980.
17. I. J. GOOD, Subjective probability as the measure of a non-measurable set, 1962, reprinted in "Studies in Subjective Probability" (H. E. Kyburg and H. E. Smokler, Eds.), Krieger, New York, 1980.
18. J. M. HARRISON, Independence and calibration in decision analysis, *Manage. Sci.* **24** (1977), 320-328.
19. J. C. HARSANYI, Games with incomplete information played by 'Bayesian' players: I-III, *Manage. Sci.* **14** (1967, 1968), 159-182, 320-334, 486-502.
20. J. C. HARSANYI, Subjective probability and the theory of games: Comments on Kadane and Larkey's paper, and Rejoinder to reply, *Manage. Sci.* **28** (1982), 120-124, 124-125.
21. S. HART AND D. SCHMEIDLER, Existence of correlated equilibria, *Math. Oper. Res.* **14** (1989), 18-25.
22. J. B. KADANE AND P. D. LARKEY, Subjective probability and the theory of games, and Reply, *Manage. Sci.* **28** (1982), 113-120, 124.
23. J. B. KADANE AND P. D. LARKEY, The confusion of is and ought in game theoretic contexts, *Manage. Sci.* **29** (1983), 1365-1379.
24. B. O. KOOPMAN, The bases of probability, 1940, reprinted in "Studies in Subjective Probability" (H. E. Kyburg and H. E. Smokler, Eds.), pp. 121-131, Krieger, New York, 1980.
25. E. A. LEAMER, Bid-ask spreads for subjective probabilities, in "Bayesian Inference and Decision Techniques" (P. Goel and A. Zellner, Eds.), pp. 217-231, North-Holland, Amsterdam, 1986.
26. H. MOULIN AND J. P. VIAL, Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved on, *Int. J. Game Theory* **7** (1978), 201-221.
27. R. B. MYERSON, Acceptable and predominant correlated equilibria, *Int. J. Game Theory* **15** (1986), 133-154.
28. R. B. MYERSON, Multistage games with communication, *Econometrica* **54** (1986), 323-358.
29. J. NASH, Non-cooperative games, *Ann. of Math.* **54** (1951), 286-295.
30. R. F. NAU, "Indeterminate Probabilities and Utilities on Finite Sets," Duke University, Fuqua School of Business Working Paper 8929, 1989.
31. R. F. NAU, "Joint Coherence in Games of Incomplete Information," Duke University, Fuqua School of Business Working Paper 8916, 1989.
32. D. G. PEARCE, Rationalizable strategic behavior and the problem of perfection, *Econometrica* **54** (1984), 1029-1050.
33. L. J. SAVAGE, "The Foundations of Statistics," Wiley, New York, 1954.
34. C. A. B. SMITH, Consistency in statistical inference and decision, *J. Roy. Statist. Soc. B* **23** (1961), 1-25.
35. P. SUPPES, The measurement of belief, *J. Roy. Statist. Soc. B* **36** (1974), 160-175.
36. T. TAN AND S. WERLANG, The Bayesian foundations of solution concepts of games, *J. Econ. Theory* **45** (1988), 370-391.
37. H. VARIAN, The arbitrage principle in financial economics, *J. Econ. Perspect.* **1** (1987), 55-72.
38. P. WALLEY, "The Elicitation and Aggregation of Beliefs," Technical Report, Department of Statistics, University of Warwick, 1982.
39. P. M. WILLIAMS, Indeterminate Probabilities, in "Formal Methods in the Methodology of Empirical Sciences" (M. Przełęczki, K. Szaniawski, and R. Wójcicki, Eds.), Ossolineum and D. Reidel, Dordrecht, Holland, 1976.