Journal of Economic Perspectives—Volume 6, Number 4—Fall 1992—Pages 83–101

Knowledge and Equilibrium in Games

Adam Brandenburger

In recent years, economists have used game-theoretic reasoning in a variety of areas: the economics of business strategy, the economics of incentives, the theory of bargaining, the theory of bidding, trade policy, monetary policy, and others. The part of game theory most frequently applied falls under the heading of "non-cooperative" game theory. A non-cooperative game is one in which each player is faced with a specified set (or series of sets) of choices. In making choices in actual strategic situations, people must grapple with questions such as: What will the other people involved do? What are their motivations? What do they think I will do? And so on. This paper describes an approach to non-cooperative game theory that aims to capture considerations that exercise the minds of real-world strategists.

The most commonly used tool of non-cooperative game theory is the Nash equilibrium (Nash, 1951). This raises the question: Are there assumptions on what the players in a game think—including what they think other players think, and so on—that lead to consideration of Nash equilibrium? The paper provides answers to this, and related, questions.

The approach of this paper involves analyzing the decision problem facing each player in a strategic ("interactive") situation. It might be described as doing "interactive decision theory." In addition to grounding game theory in considerations that are of the essence in actual strategic situations, the approach has a number of other objectives. As conventionally formulated, non-cooperative game theory seems somewhat divorced from single-person decision theory, with its special apparatus of strategy randomization, equilibrium points,

■ *Adam Brandenburger is Associate Professor of Business Administration, Harvard Business School, Boston, Massachusetts.*

and the like. By contrast, the approach of this paper aims to make game theory more immediately accessible to people who are trained in decision theory but who are not "game theorists." A related objective is to make game theory easier to teach to students. Finally, the approach suggests new directions for research into the nature of strategic situations. Some further thoughts along these lines are offered in the closing section of the paper.

Common Knowledge of Rationality

A finite non-cooperative game in "strategic" (or "normal") form consists of a finite set of players and, for each player, a finite set of choices and a payoff function associating a payoff to that player with every specification of choices by all the players.[1] This formulation of a game might seem rather restrictive; for example, it appears not to allow for any temporal structure to the game, for any information that one player might have about previous choices made by other players and the like. Nevertheless, von Neumann and Morgenstern (1944) argued that the formulation is in fact entirely general. Their key postulate is that while the rules of a particular game may involve a player making a series of choices, all these choices can be collapsed into the single choice of a *strategy*, where a strategy is a specification of what choice the player would make in any contingency that could arise during the play of the game. Accordingly, this paper will focus on the strategic form of a game.

In the strategic form, each player must decide on a particular strategy in ignorance of the strategies chosen by the other players. This is because a player cannot condition plans on those of other players. Conditioning choices on previously observed choices of other players is, of course, allowed for in the definition of a strategy. A player's choice of strategy, then, must in general depend on a subjective assessment of the relative likelihood of different strategy choices by the other players.

Cast in this form, the problem of choosing a strategy in a game is a decision problem under uncertainty. This perspective on non-cooperative game theory is not the standard one. Traditionally, decision theory has been perceived as appropriate only when the uncertainty pertains to "exogenous" events (or to the "state of nature" as it is sometimes termed), not when the uncertainty is over the strategy choices of players in a game. For the latter, "endogenous" variety of uncertainty, it has been customary to argue that an equilibrium, or other game-theoretic, construct must be employed.[2] By contrast, the view espoused in this paper is that there is no distinction between the two types of

[1] In applications, it is often assumed that the players' choice sets are infinite, say closed intervals of the real line. However, to avoid technical complications, this paper will treat only finite games.

[2] For example, Harsanyi (1967–68, p. 167) has written: "[E]very player i ... will assign a *subjective* joint probability distribution P_i to all variables unknown to him—or at least to all unknown *independent* variables, i.e., to all variables not depending on the players' own strategy choices."

Figure 1
A Simple Two Person Game

		Player 2 L	Player 2 C	Player 2 R
Player 1	U	4, 10	3, 0	1, 3
Player 1	D	0, 0	2, 10	10, 3

uncertainty and that a player in a game should assign subjective probabilities to *all* uncertainty, including the strategy choices of other players.

Consider, for example, the game depicted in Figure 1, which shows the possible strategy choices for each player beside the rows and above the columns, and the payoff to each player inside the cells, with player 1's payoff coming first. We assume that player 1 chooses a strategy that maximizes expected payoff given some subjective assessment over player 2's choice. Does this assumption allow us to reach any conclusions concerning player 1's choice? The answer is evidently no: either U or D might be optimal for player 1, depending on that player's precise assessment. What about player 2? Simply knowing that player 2 maximizes, given some assessment over player 1's choice, is sufficient to rule out player 2 choosing R. If player 2 assesses probability greater than $1/2$ to player 1 choosing U, then player 2's best choice is L; if player 2 assesses probability less than $1/2$ to player 1 choosing U, then player 2's best choice is C; finally, either L or C is best if player 2 assesses exactly probability $1/2$ to player 1 choosing U.

Now comes an important assumption: suppose that the conclusion just reached concerning player 2 is available to player 1. What exactly is entailed by this assumption will be discussed in a moment. First, however, let us see where this line of reasoning leads. If player 1 can indeed eliminate the possibility of player 2 choosing R then player 1's subjective assessment should assign probability 0 to this event, in which case U is a better choice for player 1 than is D. If this latter conclusion is in turn available to player 2 then player 2 will choose L. Our reasoning has led to the conclusion that player 1 will choose U and player 2 will choose L.

What informational assumptions are implicit in this line of reasoning? First, player 1 is assumed to know that player 2 is *rational*, which means that the player chooses a strategy that maximizes expected payoff, given some probabilistic assessment over the other players'—in this case, player 1's—strategy choices. Second, player 1 is also assumed to know player 2's payoffs. These two assumptions together allow player 1 to conclude that player 2 will not choose R. Similarly, player 2 is assumed to know that player 1 is rational and to know player 1's payoffs, and to know that player 1 knows player 2 is rational and that player 1 knows player 2's payoffs. These assumptions allow player 2 to infer that player 1 will choose U.

The general version of these assumptions is the supposition that both the structure of the game (the set of players, the sets of strategies, and the payoff functions) and the rationality of the players are "common knowledge." To say that a fact is *common knowledge* means that everyone knows it, everyone knows that everyone knows it, and so on ad infinitum. The term "common knowledge" was first used in this connection by Lewis (1969), who attributes the basic idea to Schelling (1960). Aumann (1976) offered a precise mathematical formulation of the concept. The consequence of this supposition is that for each player any strategy that does not maximize expected payoff for some probabilistic assessment over the other players' choices should be deleted from the game. In the reduced game that results, again for each player any strategy that does not maximize expected payoff for some probabilistic assessment over opponents' choices in the reduced game should be deleted. This procedure continues until no further deletion is possible.

This iterated deletion procedure can be characterized in another useful way. Recall the definition of a strongly dominated strategy: a strategy of a player is strongly dominated if that player has some other (possibly mixed) strategy that, for every specification of choices by the other players, yields a strictly higher expected payoff. A well-known result states that a strategy is strongly dominated if and only if there is no probability distribution over the choices of the other players for which the strategy in question maximizes the player's expected payoff.[3] Thus in the game depicted in Figure 1, the strategy R of player 2 is strongly dominated by a mixed strategy that gives weight $1/2$ to each of L and C. A strategy will be called *iteratively undominated* if it remains after all strongly dominated strategies have been deleted from the game, all strategies that become strongly dominated in the reduced game are then eliminated, and so on. The conclusion of the preceding paragraph can now be restated as the following proposition, versions of which are proved, with differing degrees of formality, in Bernheim (1984, 1986), Pearce (1984), Brandenburger and Dekel (1987), and Tan and Werlang (1988).

Proposition 1: Suppose that the structure of the game and the rationality of the players are common knowledge. Then each player chooses an iteratively undominated strategy.

A question arises as to whether the full force of common knowledge is really needed here. For example, in the game depicted in Figure 1, player 1's knowledge about player 2, player 2's knowledge about player 1, and player 2's knowledge about player 1's knowledge about player 2 sufficed. In other words, knowledge to (at most) two levels was needed. More generally, for a given finite game there is an obvious upper bound on the number of levels of

[3]The necessity direction of this result is straightforward; sufficiency is a deeper result that relies on the supporting hyperplane theorem. For a proof see, for example, Pearce (1984, Appendix B).

knowledge needed to conclude that players choose iteratively undominated strategies. However, if one is to state a sufficient condition that applies to all finite games, common knowledge is needed.

A point worth emphasizing is that the development so far assumes that each player makes a definite choice of strategy, with no attempt to randomize. In terms of the usual parlance, players choose "pure" and not "mixed" strategies.[4] There is, of course, no difficulty if a player should happen to want to choose a mixed strategy. A commitment to roll a die, for example, can simply be viewed as an extra pure strategy. However, at no point in what follows do we require players to choose mixed strategies.

At this juncture, the reader familiar with this sort of analysis may be asking what is really new here. After all, arguments based on dominance and iterated dominance go back a long way in game theory (for example, Luce and Raiffa, 1957). Nevertheless, it is only rather recently that the decision-theoretic foundations of iterated dominance have been spelled out in detail, as in the papers cited immediately before Proposition 1. Nor is the delay a mystery: a proper understanding could not develop until the crucial notion of common knowledge was in place.

In fact, Bernheim and Pearce are led to consider "rationalizable" strategies, which differ somewhat from iteratively undominated strategies, because they make a further common knowledge assumption: they require it to be common knowledge that each player's probability distribution on the strategy choices of the other players exhibit no dependence between those choices. In the language of probability theory, the distributions are required to be "stochastically independent."

In two-person games rationalizable strategies are clearly the same as iteratively undominated strategies, since a player faces only one other player and hence the stochastic independence requirement has no bite.

For games with more than two players, any strategy which is rationalizable is iteratively undominated. The three-person game depicted in Figure 2 demonstrates that the converse is false. Player 1 chooses the row, player 2 the column, and player 3 the matrix. No strategy of any player is strongly dominated, hence all strategies are iteratively undominated. In particular, notice that B is an optimal strategy for player 3 if that player assesses probability 1/2 to player 1 choosing U and player 2 choosing L, and probability 1/2 to player 1 choosing D and player 2 choosing R. This distribution is stochastically dependent: player 3 believes that if player 1 chooses U, player 2 chooses L; if player 1 chooses D, player 2 chooses R.

On the other hand, if player 3 is restricted to an independent distribution over player 1's and player 2's choices, a simple calculation shows that B can never be optimal. Following the deletion of B, next D and R, and then finally

[4]The reader should not be misled into thinking that play of mixed strategies is somehow involved by virtue of the mention of mixed strategies in the definition of a strongly dominated (pure) strategy. Mixed strategies appear in this context as a formal device only: they are used to test whether a given (pure) strategy should be deleted from the game.

Figure 2
A Simple Three Person Game

	L	R
U	1, 1, 1	1, 0, 1
D	0, 1, 0	0, 0, 0

A

	L	R
U	2, 2, .7	0, 0, 0
D	0, 0, 0	2, 2, .7

B

	L	R
U	1, 1, 0	1, 0, 0
D	0, 1, 1	0, 0, 1

C

C, can be deleted leaving only *U*, *L*, and *A* as rationalizable strategies in the sense of Bernheim and Pearce.

Non-cooperative game theory has traditionally taken the route of assuming stochastic independence, on the grounds that this captures the idea that strategies are chosen "independently." Thus in the example just described, player 3's dependent probability distribution would be ruled out for failing to reflect the "independence" of player 1's and player 2's choices. But this argument is far from watertight, as Aumann (1987, p. 16) explains:

> 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.

An analogy (for which I am grateful to Bob Aumann) may be helpful here. In elementary probability theory a sequence of coin tosses is often represented by an independently and identically distributed (i.i.d.) sequence of random variables. But this is appropriate only when the parameter of the coin is known with certainty before the first toss. In any other case one would surely not assess the same probability to heads on the first toss, and to heads on the 101st toss following 90 heads in the first 100 tosses; yet that is what independence demands.[5] Likewise, in the game-theoretic context there would in general be information about a second player's choice of strategy to be gleaned from observation of a first player's choice; dependence in probability assessments

[5]The answer in probability theory is to assume instead that the sequence of random variables is "exchangeable" (in the sense of de Finetti). This captures the "physical independence" of successive tosses, yet allows that successive outcomes contain information that can be used to modify probability estimates.

should therefore be permitted. The assumption of independence, which is normally accepted without discussion in non-cooperative game theory, in fact becomes quite questionable once the decision-theoretic viewpoint is adopted.

Pure-Strategy Nash Equilibrium

In the game depicted in Figure 1, the assumption of common knowledge of the structure of the game and the rationality of the players was sufficient to pin down a unique choice for each player. In many games, however, the assumption has little bite by itself, and places few or no restrictions on the strategies the players can choose. In the game depicted in Figure 3 (based on a game in Bernheim, 1984), no strategy is strongly dominated, and so the common knowledge assumption does not exclude any choice of either player from consideration.

However, consider the case where player 1 decides to play T after assessing probability one to player 2 choosing L, and player 2 chooses L after assessing probability one to 1 choosing B. While this situation is not ruled out by anything in our discussion so far, there is nevertheless an "inconsistency:" player 2's choice is based on an incorrect assessment of player 1's choice. Of course, this "inconsistency" may simply be an accurate reflection of the uncertainty facing player 2 but, while admitting this as a realistic possibility, we are nevertheless going to pursue the consequences of ruling out such situations.

Call a fact *mutual knowledge* if everyone knows it. This is to be distinguished from common knowledge, which also requires higher levels of knowledge—everyone knows that everyone knows, and so on. The "inconsistency" just discussed is eliminated by supposing the players' strategy choices to be mutual knowledge. To see how this works, refer again to the game in Figure 3. Can the pair of strategies (T, L) be mutual knowledge? The answer is "no," provided

Figure 3
Illustrating Mutual Knowledge

		Player 2		
		L	C	R
Player 1	T	7, 0	0, 5	0, 7
	M	5, 0	2, 2	5, 0
	B	0, 7	0, 5	7, 0

player 2 is rational, for then player 2 will choose R rather than L. In fact, the only pair of strategies that can be mutual knowledge, given that both players are rational, is (M, C). This strategy pair also constitutes a *pure-strategy Nash equilibrium*: a profile of (pure) strategy choices such that each player's choice is optimal given the choices of the other players. The general connection is summarized in the following proposition (Aumann and Brandenburger, 1991).

Proposition 2: Suppose that each player is rational and that the strategy choices of the players are mutual knowledge. Then the choices constitute a pure-strategy Nash equilibrium.

The proof is immediate: Since a player is assumed to know the strategy choices of the other players, rationality of the player means that the player's own choice must be optimal against the other choices. But this is exactly the definition of a Nash equilibrium.

While the proposition may seem transparent, it does contain a surprise: common knowledge is not needed. This is despite the folklore that has grown up connecting the notions of Nash equilibrium and common knowledge. The basis of the folklore goes something like this: player 1 plays his part of a Nash equilibrium because he thinks that player 2 is playing her part, a belief which is held because he thinks she thinks he is playing his part, and so on. The circularity of Nash equilibrium does indeed seem related to the infinite hierarchy of beliefs inherent in common knowledge, yet Proposition 2 suggests that the relationship is illusory. This is perhaps the more unexpected given that common knowledge plays a role in Proposition 1. In fact, we shall see later that common knowledge does crop up in the context of Nash equilibrium, but in a manner more subtle than that suggested by conventional wisdom.

The hypothesis of Proposition 2 leads, in the case of the game depicted in Figure 3, to a definite specification of the players' choices (for player 1 the choice M and for player 2 the choice C), but in general such a sharp conclusion is not possible. If a game possesses several pure-strategy Nash equilibria, then play of any of the equilibria is consistent with the hypothesis of Proposition 2; which equilibrium will obtain depends on the precise mutual knowledge that the players possess.

Mixed-Strategy Nash Equilibrium

While some games possess multiple Nash equilibria, others possess none at all, at least if players choose only pure strategies. A simple example is depicted in Figure 4. This game possesses no pure-strategy Nash equilibria, from which it follows (by Proposition 2) that no pure-strategy pair can be mutual knowledge, given that both players are rational.

The traditional way to "solve" a game such as that depicted in Figure 4 is to augment the players' choices to include mixed strategies, that is, randomized choices over the sets of pure strategies. In the example of Figure 4, there is a

Figure 4
A Game with No Pure Stategy Equilibria

		Player 2 L	Player 2 R
Player 1	U	2, 0	0, 1
Player 1	D	0, 1	1, 0

(unique) Nash equilibrium in mixed strategies, in which player 1 chooses *U* with probability 1/2 and *D* with probability 1/2, while player 2 chooses *L* with probability 1/3 and *R* with probability 2/3. Each player's mixed strategy maximizes expected payoff given the other player's mixed strategy.

A fundamental result of non-cooperative game theory (Nash, 1951) states that if mixed strategies are introduced, then every finite game possesses an equilibrium. But the use of mixed strategies, while mathematically convenient, has always been conceptually troubling. To begin with, there is never a strict incentive for a player to randomize. Take player 1 in the example: given player 2's mixed strategy, player 1 is indifferent between carrying out the 50 : 50 randomization described above and any other randomization between *U* and *D* (including choosing either *U* or *D* for sure). This would not matter if no such "deviation" by player 1 upset the equilibrium in question, but this is not the case. It is only when player 1 performs the 50 : 50 randomization that player 2 is prepared to play her part of the equilibrium. Further clouding the picture is the feeling that people simply do not randomize when making decisions. Given these difficulties, it is perhaps not surprising that many economists have tended to reject use of equilibria that involve mixed strategies.

Does the decision-theoretic approach to non-cooperative game theory shed some light on the matter? In particular, can a different interpretation be placed on mixed strategies? And, if so, what assumptions on the players' state of knowledge lead to mixed-strategy equilibrium?

In answering, it will be helpful to begin with a definition: the *conjecture* of a player is the player's probability assessment over the strategy choices of the other players. Consider now a mixed-strategy equilibrium. Allow that players do not actually randomize; rather, each player chooses some definite pure strategy. Allow also that the other players need not know which one; or, in other words, that choices need not be mutual knowledge. Each player's mixed strategy can then be thought of as representing not a conscious randomization by that player, but rather the (common) conjecture held by the other players about that player's choice.

This interpretation of mixed-strategy equilibrium, which originated in Harsanyi (1973) and was further developed by Armbruster and Böge (1979), Böge and Eisele (1979), Aumann (1987), Tan and Werlang (1988), and

Brandenburger and Dekel (1989), among others, offers an appealing solution to the interpretational difficulties described above. Wholesale rejection of mixed-strategy equilibria is unnecessary; instead, a mixed-strategy equilibrium can be viewed as an "equilibrium in conjectures."

To take an example, the mixed-strategy equilibrium of the game in Figure 4 can be interpreted as an "equilibrium in conjectures," in which player 1 assesses probability 1/3 to player 2 choosing L and probability 2/3 to player 2 choosing R, and player 2 assesses probability 1/2 to player 1 choosing U and probability 1/2 to player 1 choosing D.

The question, however, remains: how does such an "equilibrium in conjectures" come about? In particular, what assumptions on the players' state of knowledge lead to their conjectures being in equilibrium? It turns out that the answer is significantly different in the two-player and many-player situations. The following result provides the answer for two-player games (Aumann and Brandenburger, 1991).

Proposition 3: Consider a two-person game. Suppose that the structure of the game, the rationality of the players, and their conjectures are mutual knowledge. Then the conjectures constitute a mixed-strategy Nash equilibrium.

To see why Proposition 3 is true, it will help to have in mind the following characterization of Nash equilibrium. The mixed-strategy pair (σ_1, σ_2), where σ_1 is a mixed strategy of player 1 and σ_2 is a mixed strategy of player 2, is a Nash equilibrium if and only if every pure strategy of player 1 assigned positive weight by σ_1 maximizes player 1's expected payoff given that player 2 is playing σ_2 and every pure strategy of player 2 assigned positive weight by σ_2 maximizes player 2's expected payoff given that player 1 is playing σ_1. Turning to the proof of Proposition 3, let us, in a suggestive notation, write σ_1 for the conjecture that player 1 knows player 2 to hold concerning player 1's choice of strategy. Similarly, let σ_2 denote the conjecture that player 2 knows player 1 to hold concerning player 2's choice of strategy. Now, since player 1 knows player 2 to be rational and also knows player 2's payoff function, any strategy of player 2 to which player 1 assigns positive probability, that is, any strategy of player 2 assigned positive weight by σ_2, must maximize player 2's expected payoff given the conjecture σ_1, that player 1 knows player 2 to hold. Similarly, since player 2 knows player 1 to be rational and also knows player 1's payoff function, any strategy of player 1 to which player 2 assigns positive probability, that is, any strategy of player 1 assigned positive weight by σ_1, must maximize player 1's expected payoff given the conjecture σ_2, that player 2 knows player 1 to hold. The preceding two sentences establish that (σ_1, σ_2) satisfies the condition characterizing Nash equilibrium.

As is evident from the proof, a corollary of Proposition 3 is that the players' expected payoffs are also determined: they are equal to the expected payoffs in the mixed-strategy Nash equilibrium to which the players' conjectures correspond. Notice that it is *not* claimed that the players' (pure) strategy choices are

Figure 5
An Interactive Belief System

	L_1	L_2	L_3
U_1	A_1 (2/5)	B_1 (1/5)	
U_2		A_2 (1/5)	B_2 (1/10)
U_3			A_3 (1/10)

in equilibrium. Indeed this may be impossible, as in the case of the game of Figure 4.

The conditions of Proposition 3 are insufficient in games with more than two players. To see this, consider a three-person game in which players 1 and 2 each have just one choice—say U for player 1 and L for player 2—and player 3 has two choices, A and B.[6] We are going to describe a situation in which there is mutual knowledge of conjectures, yet player 1's and player 2's conjectures about player 3 disagree, so that no well-defined mixed-strategy profile emerges. The payoffs are unimportant for our purposes and so will be left unspecified, but they could easily be chosen so as to make both the structure of the game and the players' rationality mutual knowledge. Hence the example will suffice to establish the claim made at the beginning of this paragraph.

To construct the situation, we introduce a little more formal apparatus than has been employed so far. Accordingly, consider the "interactive belief system" depicted in Figure 5. Here it is assumed that there are three different possible "types" of player 1, corresponding to the three rows and labelled U_1, U_2, and U_3, three different possible "types" of player 2, corresponding to the three columns and labelled L_1, L_2, and L_3, and five different possible "types" of player 3, corresponding to the five boxes and labelled A_1, B_1, A_2, B_2, and A_3, where a *type* of a player describes the player's strategy choice and the probabilities the player assesses to the various types of the other players.[7]

The labelling of types indicates each type's strategy choice: U for all types of player 1, L for all types of player 2, and alternately A and B for the various types of player 3.

[6]The following example is taken from Aumann and Brandenburger (1991), and is based on an idea of John Geanakoplos.
[7]More generally, a type of a player also describes the player's payoff function, but we avoid this extra complexity here. See Aumann and Brandenburger (1991) for a full treatment.

As for probabilities, it is assumed in Figure 5 that the types' assessments are consistent with the existence of a "common prior." To understand this idea, it will help to have the notion of a "state of the world," which is simply a specification of a type for each player. Thus in the example there are $3 \times 3 \times 5 = 45$ different possible states of the world. A *common prior* is then a probability distribution on the states of the world such that for each player, the probability assessment of a given type of that player is equal to the conditional distribution obtained from the common prior by conditioning on the event that the player's type is in fact the given type.

The numbers in parentheses in Figure 5 give the common prior probabilities of the five states indicated. (All other states are assumed to have zero prior probability.) What, then, is the probability assessment of type U_1 of player 1? The answer is found by conditioning on the row in question, namely the first one. Thus type U_1 assesses probability 2/3 to player 2 being of type L_1 and player 3 being of type A_1, and probability 1/3 to player 2 being of type L_2 and player 3 being of type B_1. To take another example, the probability assessment of type L_2 of player 2 is found by conditioning on the second column: type L_2 assesses probability 1/2 to player 1 being of type U_1 and player 3 being of type B_1, and probability 1/2 to player 1 being of type U_2 and player 3 being of type A_2. Finally, player 3 always assesses probability one to the true state of the world: if, for example, the true state is (U_2, L_2, A_2), player 3 assesses probability one to player 1 being of type U_2 and player 2 being of type L_2.

An interactive belief system is a formal description of the players' choices, of their conjectures about other players' choices, of their beliefs about other players' conjectures, and so on. It should be clear how a particular state of the world determines the players' choices and also their conjectures. Take the state (U_1, L_1, A_1). The players' choices are evidently U, L, and A, respectively. Player 1's conjecture derives immediately from type U_1's assessment: player 1 assesses probability 2/3 to player 2 choosing L and player 3 choosing A, and probability 1/3 to player 2 choosing L and player 3 choosing B. The other players' conjectures are found in similar fashion.

But what is player 1's belief concerning player 2's conjecture, or player 1's belief concerning player 2's belief concerning player 3's conjecture, and so on? Are these higher-order beliefs already determined in some fashion, or must they too be specified? The beauty of the interactive belief system model is that, appearances to the contrary, all these higher-order beliefs are already determined. To see this, simply note that, at the state (U_1, L_1, A_1) say, player 1 assesses probability 2/3 to player 2 being of type L_1 and probability 1/3 to player 2 being of type L_2. Hence player 1 believes that player 2 either: a) assesses probability one to player 1 choosing U and player 3 choosing A; or b) assesses probability 1/2 to player 1 choosing U and player 3 choosing B, and probability 1/2 to player 1 choosing U and player 3 choosing A. Indeed, player 1 assesses probability 2/3 to the first situation and probability 1/3 to the second. It should be apparent that Figure 5 can be used in like fashion to calculate all higher-order beliefs.

Notice an implicit assumption here: the interactive belief system is in some sense "transparent" to the players themselves. This is not, in fact, a formal assumption; rather, it is a tautology, reflecting the fact that the interactive belief system already describes any ignorance on the part of the players.[8]

Before proceeding with using the interactive belief system of Figure 5 to describe a situation of mutual knowledge of conjectures, let us delay a little longer and explain why it was important to examine the interactive belief system model in such detail. This paper began by promising to be much more explicit than is customary in non-cooperative game theory about the decision problem facing each player in a game. The interactive belief system model gives rigorous shape to this approach, and allows one to state and prove all the results described in this paper as formal theorems. This applies not only to Propositions 4 and 5 below, which rely on the interactive belief system model in an obvious way, but also to Propositions 1–3. Although the argument for the latter proceeded at an intuitive level, the reader should now be in a position to see that they can also be given formal content; they are not simply "stories."

Consider again Figure 5. We assert that if the true state is (U_2, L_2, A_2), then the players' conjectures are mutual knowledge. At this state: a) player 1 assesses probability 2/3 to player 2 choosing L and player 3 choosing A, and probability 1/3 to player 2 choosing L and player 3 choosing B; b) player 2 assesses probability 1/2 to player 1 choosing U and player 3 choosing B, and probability 1/2 to player 1 choosing U and player 3 choosing A; and, finally, c) player 3 assesses probability one to player 1 choosing U and player 2 choosing L. Player 1 knows player 2's and player 3's conjectures, as they are the same at the only other state, (U_2, L_3, B_2), that player 1 considers possible. Similarly, player 2 knows player 1's and player 3's conjectures, as they are the same at the only other state, (U_1, L_2, B_1), that player 2 considers possible. Finally, player 3 knows player 1's and player 2's conjectures, since player 3 knows the true state. So the conjectures are mutual knowledge. Yet player 1's and player 2's conjectures about player 3 disagree: player 1 thinks it twice as likely that player 3 will choose A than B, while player 2 considers the two choices equally likely. Hence no well-defined mixed-strategy profile emerges. The conclusion is that the conditions of Proposition 3 do not suffice once there are more than two players.

Notice that at (U_2, L_2, A_2) player 2 does not know that player 1 knows player 2's conjecture. The reason is that player 2 thinks it possible that player 1 is of type U_1, in which case, as we saw above, player 1 entertains two different

[8]Further discussion of this and other aspects of the interactive belief system model can be found in Aumann (1987) and Aumann and Brandenburger (1991). Relevant foundational papers include Armbruster and Böge (1979), Mertens and Zamir (1985), Tan and Werlang (1985), and Brandenburger and Dekel (1992).

possibilities for the conjecture that player 2 holds. Certainly the players' conjectures are not common knowledge at (U_2, L_2, A_2). It turns out that exactly this extra condition is sufficient for mixed-strategy Nash equilibrium in many-person games, as the following result indicates (Aumann and Brandenburger, 1991).

Proposition 4: Suppose that there is a common prior on the set of states of the world. Suppose further that at some state the structure of the game and the rationality of the players are mutual knowledge, and that the conjectures of the players are common knowledge. Then for each player, all the other players hold the same conjecture about that player, and the resulting profile of conjectures constitutes a mixed-strategy Nash equilibrium.

We discussed earlier the folklore connecting Nash equilibrium and common knowledge and saw that the precise nature of the connection, if indeed any existed, was not obvious. With Proposition 4, common knowledge finally enters the picture. But the connection that emerges is unexpected: it is common knowledge of conjectures, not of the structure of the game nor of the players' rationality, that matters for Nash equilibrium—and then only in games with more than two players.[9] Once on the scene, however, common knowledge is needed in full force. It can be shown, using a modification of the interactive belief system of Figure 5, that Proposition 4 is false if common knowledge of conjectures is replaced by iterated knowledge of conjectures to any finite number of levels. Aumann and Brandenburger (1991) offer details.

Proposition 4 bears a relation to Aumann's (1976) well-known theorem on the impossibility of "agreeing to disagree:" if two individuals start with a common prior and their beliefs concerning some event are common knowledge, then their beliefs must be equal.[10] Likewise, in the game-theoretic context, if two players' conjectures are common knowledge, then their conjectures about a third player's choice of strategy must agree. Hence common knowledge of conjectures rules out the situation earlier in which player 1's and

[9]Ben Polak has pointed out the following interesting result. Suppose that at some state the structure of the game and the players' conjectures are common knowledge. Then if rationality is mutual knowledge it is also common knowledge. (For a proof, see Aumann and Brandenburger, 1991.) In this sense, common knowledge of rationality does play a part in equilibrium, but it should be viewed not as a condition for equilibrium but rather as an "epiphenomenon" of (somewhat tightened) conditions for equilibrium.

[10]To gain some intuition for this result, consider two individuals 1 and 2 whose beliefs concerning some event E are common knowledge. Since 1's belief is common knowledge, it must be that the conditional probability of E is the same, equal to $\frac{1}{4}$ say, whether conditioned on: a) the information 1 actually possesses; or b) any piece of information that 2 thinks 1 might possess (otherwise 2 would not know 1's belief); or c) any piece of information that 1 thinks 2 thinks 1 might possess (otherwise 1 would not know that 2 knows 1's belief); and so on. It follows that the conditional probability of E is again $\frac{1}{4}$ if conditioned only on the union of all these pieces of information. But exactly the same argument can be made for individual 2 and, since the unions are the same for both people (this needs some thought!), the conditional probability of E calculated in this latter fashion must also be $\frac{1}{4}$. This will be so only if the conditional probability of E, conditioned on 2's information, is $\frac{1}{4}$.

player 2's conjectures concerning player 3's choice differed. (Recall that player 1's and player 2's conjectures were mutual knowledge, but not common knowledge.) Proposition 4 goes further than this, however, in that common knowledge of conjectures also implies that each player's conjecture concerning the choices of the other players is stochastically independent. It is this, together with agreement between the players' conjectures, that leads to Nash equilibrium.

We end this section with a brief remark on the assumption of a common prior. This assumption, which is discussed at length in Harsanyi (1967–68) and Aumann (1987), always generates much controversy. Here we content ourselves with observing that the assumption of a common prior is implicit throughout almost all of economic theory. It serves to focus attention on the role of informational differences in interactive situations; without it, analysis is often inconclusive. No doubt, the last word remains to be written on this subject.

Correlated Equilibrium

The conditions of Propositions 1–4 refer to the state of knowledge of the players themselves. This "inside" perspective is natural, given the decision-theoretic approach of the paper.

However, this section alters perspective and asks the following question from the point of view of an outside observer. Suppose we are given an interactive belief system with a common prior, in which each player is rational at every state of the world. What probability distribution will we, as outside observers, assign to the players' strategy choices? At first sight it would seem that very little can be said about this distribution. After all, the answer would appear to depend on all the details of the interactive belief system, which could be very complicated. Nevertheless, a definite answer exists, as the following result indicates (Aumann, 1987).

Proposition 5: Suppose that there is a common prior on the set of states of the world and that each player is rational at every state. Then the distribution of the players' strategy choices constitutes a correlated equilibrium distribution.

To see what a correlated equilibrium distribution is, and to understand why Proposition 5 is true, consider the game, sometimes called "chicken," depicted in Figure 6*a*.[11] This game has three Nash equilibria: a) one in which player 1 chooses D and player 2 chooses L; b) one in which player 1 chooses U and player 2 chooses R; and c) one in which (in the usual terminology) player 1 chooses U with probability $2/3$ and D with probability $1/3$ while player 2 chooses L with probability $2/3$ and R with probability $1/3$.

[11]Although this is a two-person game, there is no loss of generality since the theory of correlated equilibrium is the same for two-person and many-person games.

Figure 6
A Game of Chicken

Player 2

(a)

Player 1	L	R
U	6, 6	2, 7
D	7, 2	0, 0

(b)

	L	R
U	p	q
D	r	s

(c)

	L	R
U	$\frac{1}{3}$	$\frac{1}{3}$
D	$\frac{1}{3}$	0

A correlated equilibrium distribution, on the other hand, is a probability distribution (p, q, r, s) on the four possible pairs of choices in Figure 6*a*, as illustrated in Figure 6*b*, which satisfies the following system of linear inequalities:

$$6p + 2q \geq 7p,$$
$$7r \geq 6r + 2s,$$
$$6p + 2r \geq 7p,$$
$$7q \geq 6q + 2s.$$

What conditions do these inequalities embody? The first says that if the probability of player 1 choosing U is positive, then U is optimal for player 1 given the distribution on player 2's choices calculated by conditioning on player 1 playing U. (To see this, divide both sides by $(p + q)$.) The remaining inequalities refer to D, L, and R, respectively.

The distribution of strategy choices that arises in a Nash equilibrium satisfies the defining inequalities for a correlated equilibrium distribution, with the added property that the distribution is stochastically independent. Hence, all Nash equilibrium distributions are correlated equilibrium distributions. However, Figure 6*c* displays a correlated equilibrium distribution of "chicken" that is not a Nash equilibrium distribution.

Why should the distribution of the players' strategy choices, under the conditions of Proposition 5, constitute a correlated equilibrium distribution? The argument goes as follows. Fix an interactive belief system and suppose that (p, q, r, s) is the resulting distribution of the players' strategy choices. Player 1 is rational at every state of the world, hence U must be optimal for player 1 whenever he chooses it. Consider the set of states in which player 1 chooses U. If player 1's only information was that the true state was contained in this set, U would still be the optimal choice for player 1. (Notice the similarity between the argument here and in note 10.) But the conditional probabilities that player 1 would, in this situation, assign to player 2 choosing L versus R are exactly $p/(p + q)$ versus $q/(p + q)$. This gives the first inequality above. The argument for the remaining inequalities in similar.

Concluding Remarks

This paper has discussed sufficient conditions for various non-cooperative solution concepts: iterated dominance, rationalizability, pure-strategy Nash equilibrium, mixed-strategy Nash equilibrium in two-person games, mixed-strategy Nash equilibrium in many-person games, and correlated equilibrium. Table 1 summarizes conditions on the players' knowledge that lead to each of these concepts.

Table 1
A Summary of Solution Concepts

	Conditions		
Solution Concept	*Structure of the Game*	*Rationality*	*Choices or Conjectures*
Iterated dominance	Common knowledge	Common knowledge	—
Rationalizability	Common knowledge	Common knowledge	Common knowledge of stochastic independence of conjectures
Pure-strategy Nash equilibrium	—	Fact of rationality	Mutual knowledge of choices
Mixed-strategy Nash equilibrium in two-person games	Mutual knowledge	Mutual knowledge	Mutual knowledge of conjectures
Mixed-strategy Nash equilibrium in many-person games	Mutual knowledge	Mutual knowledge	Common prior plus common knowledge of conjectures
Correlated equilibrium	—	Rationality at every state	Common prior

There are also converse results of the following kind: for a given solution, say a profile of iteratively undominated strategies, there is an interactive belief system in which the appropriate conditions are satisfied.[12] But typically there are also interactive belief systems in which the conditions do *not* hold. Thus the conditions described in this paper are sufficient but not necessary. This is very natural. For example, there is no reason why players should not blunder into a Nash equilibrium "by accident," so to speak, without anybody knowing much of anything.

Conspicuous by its absence has been any discussion in this paper of the large number of so-called "refinements" of Nash equilibrium that have been

[12]For proofs, see Aumann (1987), Brandenburger and Dekel (1987), and Aumann and Brandenburger (1991).

developed in recent years. Nor has this paper discussed the related issue of elimination of weakly, rather than strongly, dominated strategies. This omission was deliberate. Work on providing decision-theoretic foundations for refinements of Nash equilibrium has started, but it would be premature to attempt a survey at this point.

The results summarized in Table 1 suggest many further questions. What do other combinations of assumptions give? What happens if, instead of common knowledge, there is iterated knowledge to a large, but finite, number of levels? What if everyone is rational, but assigns some small positive probability to the event that everyone else is not? And so on.[13] Questions such as these, which reflect very nicely the concerns of the real-world strategist, can be explored in a systematic fashion using the decision-theoretic approach to non-cooperative game theory.

■ *I would like to acknowledge the considerable intellectual debt that I owe my friend, Eddie Dekel. Many hours spent working together have done much to shape the views expressed here. I am deeply indebted to Bob Aumann for correspondence and conversations over a long period of time concerning game theory. Ben Polak, John Pratt, and Gus Stuart provided detailed and insightful comments on an earlier draft. I also wish to thank Ken Binmore, Elon Kohlberg, and Vijay Krishna for much help. Financial support from the Harvard Business School Division of Research is gratefully acknowledged.*

[13]For some answers in the context of specific games, see, for example, Stuart (1991) and Aumann (1992).

References

Armbruster, W. and W. Böge, "Bayesian Game Theory." In Moeschlin, O. and D. Pallaschke, eds., *Game Theory and Related Topics.* Amsterdam: North-Holland, 1979, 17–28.

Aumann, R., "Agreeing to Disagree," *Annals of Statistics,* 1976, *4,* 1236–39.

Aumann, R., "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica,* January 1987, *55*:1, 1–18.

Aumann, R., "Irrationality in Game Theory." In Dasgupta, P., *et al.,* eds., *Economic Analysis of Markets and Games.* Cambridge: MIT Press, 1992, 214–27.

Aumann, R. and A. Brandenburger, "Epistemic Conditions for Nash Equilibrium," Working Paper 91-042, Harvard Business School, 1991.

Bernheim, D., "Rationalizable Strategic Behavior," *Econometrica,* July 1984, *52,* 1007–28.

Bernheim, D., "Axiomatic Characterizations of Rational Choice in Strategic Environments," *Scandinavian Journal of Economics,* 1986, *88,* 473–88.

Böge, W. and Th. Eisele, "On Solutions of Bayesian Games," *International Journal of Game Theory,* 1979, *8*:4, 193–215.

Brandenburger, A. and E. Dekel, "Rationalizability and Correlated Equilibria," *Econometrica,* November 1987, *55*:6, 1391–1402.

Brandenburger, A. and E. Dekel, "The Role of Common Knowledge Assumptions in Game Theory." In Hahn, F., ed., *The Economics of Missing Markets, Information, and Games.* Oxford: Oxford University Press, 1989, 46–61.

Brandenburger, A. and E. Dekel, "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory,* forthcoming 1992.

Harsanyi, J., "Games with Incomplete Information Played by 'Bayesian' Players," Parts I–III, *Management Science,* November 1967, *14*:3, 159–82; January 1968, *14*:5, 320–34; March 1968, *14*:7, 486–502.

Harsanyi, J., "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory,* 1973, *2*:1, 1–23.

Lewis, D., *Conventions: A Philosophical Study.* Cambridge: Harvard University Press, 1969.

Luce, R. and H. Raiffa, *Games and Decisions.* New York: Wiley, 1957.

Mertens, J-F. and S. Zamir, "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory,* 1985, *14*:1, 1–29.

Nash, J., "Non-Cooperative Games," *Annals of Mathematics,* 1951, *54,* 286–95.

Pearce, D., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica,* July 1984, *52*:4, 1029–50.

Schelling, T., *The Strategy of Conflict.* Cambridge: Harvard University Press, 1960.

Stuart, H. W., "Irrationality in the Finitely Repeated Prisoner's Dilemma," unpublished, Harvard Business School, 1991.

Tan, T. and S. Werlang, "On Aumann's Notion of Common Knowledge—An Alternative Approach," unpublished, Department of Economics, Princeton University, 1985.

Tan, T. and S. Werlang, "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory,* August 1988, *45,* 370–91.

von Neumann, J. and O. Morgenstern, *Theory of Games and Economic Behavior.* Princeton: Princeton University Press, 1944.

This article has been cited by:

1. Christian W. Bach, Jérémie Cabessa. 2023. Lexicographic agreeing to disagree and perfect equilibrium. *Journal of Mathematical Economics* **109**, 102908. [Crossref]
2. Áron Tóbiás. 2021. Meet meets join: the interaction between pooled and common knowledge. *International Journal of Game Theory* **50**:4, 989-1019. [Crossref]
3. Davide Marchiori, Sibilla Di Guida, Luca Polonio. 2021. Plasticity of strategic sophistication in interactive decision-making. *Journal of Economic Theory* **196**, 105291. [Crossref]
4. Roy Chen, Yan Chen, Yohanes E. Riyanto. 2021. Best practices in replication: a case study of common information in coordination games. *Experimental Economics* **24**:1, 2-30. [Crossref]
5. Edith Elkind, Umberto Grandi, Francesca Rossi, Arkadii Slinko. 2020. Cognitive hierarchy and voting manipulation in k -approval voting. *Mathematical Social Sciences* **108**, 193-205. [Crossref]
6. Vincent P. Crawford. 2019. Experiments on Cognition, Communication, Coordination, and Cooperation in Relationships. *Annual Review of Economics* **11**:1, 167-191. [Crossref]
7. . Newcomb's Problem **70**, . [Crossref]
8. Shinji Teraji. Society and Knowledge 263-308. [Crossref]
9. Thomas A. Weber. 2016. A robust resolution of Newcomb's paradox. *Theory and Decision* **81**:3, 339-356. [Crossref]
10. Shinji Teraji. Norms, Coordination, and Order 175-205. [Crossref]
11. Yuan Xiaoyong, Tang Hongyan, Li Ying, Jia Tong, Liu Tiancheng, Wu Zhonghai. A Competitive Penalty Model for Availability Based Cloud SLA 964-970. [Crossref]
12. V. P. Crawford. 2014. Boundedly rational versus optimization-based models of strategic thinking and learning in games. *Voprosy Ekonomiki* :5, 27-44. [Crossref]
13. Vincent P. Crawford. 2013. Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games. *Journal of Economic Literature* **51**:2, 512-527. [Abstract] [View PDF article] [PDF with links]
14. Vincent P. Crawford,, Miguel A. Costa-Gomes,, Nagore Iriberri. 2013. Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature* **51**:1, 5-62. [Abstract] [View PDF article] [PDF with links]
15. Thomas A. Weber. 2013. A Robust Resolution of Newcomb's Paradox. *SSRN Electronic Journal* . [Crossref]
16. Steven A. Lippman, Kevin F. McCardle. 2012. Embedded Nash Bargaining: Risk Aversion and Impatience. *Decision Analysis* **9**:1, 31-40. [Crossref]
17. Till Grüne-Yanoff, Aki Lehtinen. Philosophy of Game Theory 531-576. [Crossref]
18. Jimmy Teng. 2012. Refining Nash Equilibrium by Bayesian Iterative Conjectures Approach. *SSRN Electronic Journal* . [Crossref]
19. Emiliano Lorini, Frédéric Moisan. 2011. An Epistemic Logic of Extensive Games. *Electronic Notes in Theoretical Computer Science* **278**, 245-260. [Crossref]
20. Paul Weirich. 2011. Introduction: Interactive Epistemology. *Episteme* **8**:3, 201-208. [Crossref]
21. Mihnea C. Moldoveanu, Joel A. C. Baum. 2011. "I Think You Think I Think You're Lying": The Interactive Epistemology of Trust in Social Networks. *Management Science* **57**:2, 393-412. [Crossref]
22. Emiliano Lorini, François Schwarzentruber. 2010. A Modal Logic of Epistemic Games. *Games* **1**:4, 478-526. [Crossref]

23. Boudewijn de Bruin. 2009. Overmathematisation in game theory: pitting the Nash Equilibrium Refinement Programme against the Epistemic Programme. *Studies in History and Philosophy of Science Part A* **40**:3, 290-300. [Crossref]
24. M. Kibria, Abbas Jamalipour. 2009. Game theoretic outage compensation in next generation mobile networks. *IEEE Transactions on Wireless Communications* **8**:5, 2602-2608. [Crossref]
25. Emiliano Lorini, François Schwarzentruber, Andreas Herzig. Epistemic Games in Modal Logic: Joint Actions, Knowledge and Preferences All Together 212-226. [Crossref]
26. Ivan Moscati. 2009. Interactive and Common Knowledge in the State-Space Model. *SSRN Electronic Journal* . [Crossref]
27. Troels Bjerre Søensen. 2008. Solution to exchanges 7.2 puzzle. *ACM SIGecom Exchanges* **7**:3, 1-4. [Crossref]
28. Eduardo Zambrano. 2008. Epistemic conditions for rationalizability. *Games and Economic Behavior* **63**:1, 395-405. [Crossref]
29. M. Rubaiyat Kibria, Abbas Jamalipour. Lossy utility based outage compensation in Next Generation Networks 389-396. [Crossref]
30. Mihnea C. Moldoveanu, Joel A. C. Baum. 2008. The Epistemic Structure and Dynamics of Social Networks. *SSRN Electronic Journal* **4**. . [Crossref]
31. . BRANCHES AND SCHOOLS OF ECONOMICS AND THEIR METHODOLOGICAL PROBLEMS 295-296. [Crossref]
32. Vernon Smith. Economics in the Laboratory 334-355. [Crossref]
33. Nicolai J. Foss. 2007. Strategic belief management. *Strategic Organization* **5**:3, 249-258. [Crossref]
34. David J. Cooper. 2007. An introduction to the symposium on behavioral game theory. *Economic Theory* **33**:1, 1-10. [Crossref]
35. S. Abu Turab Rizvi. 2007. Aumann's and Schelling's Game Theory: The Nobel Prize in Economic Science, 2005. *Review of Political Economy* **19**:3, 297-316. [Crossref]
36. Nicolai J. Foss. 2007. Strategic Belief Management. *SSRN Electronic Journal* . [Crossref]
37. Esther-Mirjam Sent. 2006. The Tricks of the (No-)Trade (Theorem). *History of Political Economy* **38**:Suppl_1, 305-321. [Crossref]
38. Justin Lenzo, Todd Sarver. 2006. Correlated equilibrium in evolutionary models with subpopulations. *Games and Economic Behavior* **56**:2, 271-284. [Crossref]
39. Ka-Man Lam, Ho-Fung Leung. 2006. A Trust/Honesty Model with Adaptive Strategy for Multiagent Semi-Competitive Environments. *Autonomous Agents and Multi-Agent Systems* **12**:3, 293-359. [Crossref]
40. Yohsuke Ohtsubo, Amnon Rapoport. 2006. Depth of reasoning in strategic form games. *The Journal of Socio-Economics* **35**:1, 31-47. [Crossref]
41. Amnon Rapoport. Games: Centipede . [Crossref]
42. Mihnea C. Moldoveanu. 2006. Thinking Strategically about Thinking Strategically. *SSRN Electronic Journal* . [Crossref]
43. Ivan Moscati. 2006. Epistemic Virtues and Theory Choice in Economics. *SSRN Electronic Journal* . [Crossref]
44. Steven A. Lippman, Kevin F. McCardle. 2004. Sex, Lies, and the Hillblom Estate: A Decision Analysis. *Decision Analysis* **1**:3, 149-166. [Crossref]
45. Larry Samuelson. 2004. Modeling Knowledge in Economic Analysis. *Journal of Economic Literature* **42**:2, 367-403. [Abstract] [View PDF article] [PDF with links]

46. Christian Montet, Daniel Serra. Non-Cooperative Games with Complete and Perfect Information 62-140. [Crossref]
47. P.J. Gmytrasiewicz. Issues in rational planning in multi-agent settings 8 pp.. [Crossref]
48. Frank Heinemann, Gerhard Illing. 2002. Speculative attacks: unique equilibrium and transparency. *Journal of International Economics* **58**:2, 429-450. [Crossref]
49. Craig R. Fox, Martin Weber. 2002. Ambiguity Aversion, Comparative Ignorance, and Decision Context. *Organizational Behavior and Human Decision Processes* **88**:1, 476-498. [Crossref]
50. Robert Grafstein. 2002. What Rational Political Actors Can Expect. *Journal of Theoretical Politics* **14**:2, 139-165. [Crossref]
51. Barry Nalebuff. Prisoners' Dilemma 1489-1495. [Crossref]
52. Piotr J. Gmytrasiewicz, Sanguk Noh. Implementing a Decision-Theoretic Approach to Game Theory for Socially Competent Agents 97-118. [Crossref]
53. John Hillas, Elon Kohlberg. Chapter 42 Foundations of strategic equilibrium 1597-1663. [Crossref]
54. Mihnea Moldoveanu, Howard Stevenson. 2001. The self as a problem: the intra-personal coordination of conflicting desires. *The Journal of Socio-Economics* **30**:4, 295-330. [Crossref]
55. Shyam Sunder. 2001. Knowing What Others Know: Common Knowledge, Accounting, and Capital Markets. *SSRN Electronic Journal* . [Crossref]
56. Amnon Rapoport, Wilfred Amaldoss. 2000. Mixed strategies and iterative elimination of strongly dominated strategies: an experimental investigation of states of knowledge. *Journal of Economic Behavior & Organization* **42**:4, 483-521. [Crossref]
57. Kin Chung Lo. 2000. Epistemic conditions for agreement and stochastic independence of ϵ-contaminated beliefs. *Mathematical Social Sciences* **39**:2, 207-234. [Crossref]
58. Hervé Defalvard. 2000. Croyances individuelles et coordination sociale. *L'Actualité économique* **76**:3, 341-364. [Crossref]
59. Pierre Allan, Cédric Dupont. 1999. International Relations Theory and Game Theory: Baroque Modeling Choices and Empirical Robustness. *International Political Science Review* **20**:1, 23-47. [Crossref]
60. Miguel Costa-Gomes, Vincent P. Crawford, Bruno Broseta. 1999. Cognition and Behavior in Normal-Form Games: An Experimental Study. *SSRN Electronic Journal* . [Crossref]
61. J. Berg, A. Engel. 1998. Matrix Games, Mixed Strategies, and Statistical Mechanics. *Physical Review Letters* **81**:22, 4999-5002. [Crossref]
62. Peter Newman. P 1-198. [Crossref]
63. John B. Van Huyck, Raymond C. Battalio, Frederick W. Rankin. 1997. ON THE ORIGIN OF CONVENTION: EVIDENCE FROM COORDINATION GAMES*. *The Economic Journal* **107**:442, 576-596. [Crossref]
64. Larry G. Epstein. 1997. Preference, Rationalizability and Equilibrium. *Journal of Economic Theory* **73**:1, 1-29. [Crossref]
65. Luc Lismont, Philippe Mongin. On the Logic of Common Belief and Common Knowledge 3-34. [Crossref]
66. H. S. Shin, T. Williamson. Representing the Knowledge of Turing Machines 169-190. [Crossref]
67. Hans Jørgen Jacobsen. 1996. On the Foundations of Nash Equilibrium. *Economics and Philosophy* **12**:1, 67-88. [Crossref]
68. Giacomo Bonanno. 1995. Rationality and Coordination, Bicchieri Cristina. Cambridge University Press, 1994, xiii + 270 pages. *Economics and Philosophy* **11**:2, 359-366. [Crossref]

69. Raymond Dacey. 1995. A general model of international interactions and peace economics. *International Interactions* **21**:2, 155-180. [Crossref]
70. Dale O. Stahl. 1995. Lexicographic rationalizability and iterated admissibility. *Economics Letters* **47**:2, 155-159. [Crossref]
71. E.A. Billard, J.C. Pasquale. 1995. Probabilistic coalition formation in distributed knowledge environments. *IEEE Transactions on Systems, Man, and Cybernetics* **25**:2, 277-286. [Crossref]
72. Cristina Bicchieri. The Epistemic Foundations of Nash Equilibrium 91-146. [Crossref]
73. Peter Vanderschraaf. 1995. Convention as correlated equilibrium. *Erkenntnis* **42**:1, 65-87. [Crossref]
74. Luc Lismont, Philippe Mongin. 1994. On the logic of common belief and common knowledge. *Theory and Decision* **37**:1, 75-106. [Crossref]
75. Hyun Song Shin, Timothy Williamson. 1994. Representing the knowledge of turing machines. *Theory and Decision* **37**:1, 125-146. [Crossref]
76. Xiaohua Lu, Dale O. Stahl. 1994. Responsive and strong responsive evolutionary dynamics. *International Journal of Game Theory* **23**:2, 101-117. [Crossref]
77. Vernon L. Smith. 1994. Economics in the Laboratory. *Journal of Economic Perspectives* **8**:1, 113-131. [Abstract] [View PDF article] [PDF with links]
78. Barry O'Neill. Sources in Game Theory for International Relations Specialists 9-30. [Crossref]
79. Carsten Herrmann-Pillath. 1994. Evolutionary rationality, "Homo economicus," and the foundations of social order. *Journal of Social and Evolutionary Systems* **17**:1, 41-69. [Crossref]
80. Eric Pacuit, Rohit Parikh. Social Interaction, Knowledge, and Social Software 441-461. [Crossref]
81. Mihnea C Moldoveanu, Joel A.C Baum, Tim J Rowley. INFORMATION REGIMES, INFORMATION STRATEGIES AND THE EVOLUTION OF INTERFIRM NETWORK TOPOLOGIES 221-264. [Crossref]