

Epistemic Game Theory

Lecture 13

Eric Pacuit

University of Maryland, College Park

pacuit.org

epacuit@umd.edu

May 6, 2014

What is game theory trying to accomplish?

What is game theory trying to accomplish?

What is *epistemic* game theory trying to accomplish?

R. Aumann. *What is game theory trying to accomplish?*. Frontiers of Economics, 1985.

“...in my view, scientific theories are not to be considered “true” or “false.” In constructing such a theory, we are not trying to get at the truth, or even to approximate to it: rather, we are trying to organize our thoughts and observations in a useful manner.”

“Truth, however, is not the issue. We discard a theory not because it has been “disproved,” but because it no longer works, is no longer appropriate.”

What difference does it make whether we are looking for the truth or for a workable model, as long as we are not dogmatic and are willing to consider new evidence or new ways of thinking?

What difference does it make whether we are looking for the truth or for a workable model, as long as we are not dogmatic and are willing to consider new evidence or new ways of thinking?

I think that the distinction is crucial for social science in general, and for game theory and economics in particular....People ask, since game theory offers a multiplicity of solution notions, what good can it be? Which solution notion is the right one? How do people “truly” behave?

What difference does it make whether we are looking for the truth or for a workable model, as long as we are not dogmatic and are willing to consider new evidence or new ways of thinking?

I think that the distinction is crucial for social science in general, and for game theory and economics in particular....People ask, since game theory offers a multiplicity of solution notions, what good can it be? Which solution notion is the right one? How do people “truly” behave? If one takes the point of view suggested above, this question loses much of its sharpness. None of the solution notions tells us how people truly behave....Rather, a solution notion is the scientists’ way of organizing in a single framework many disparate phenomena and many disparate ideas.”

Pitfalls of Taking it too Literally

...the validity of utility maximization does not depend on its being an accurate description of the behavior of individuals. Rather, it derives from its being the underlying postulate that pulls together most of economic theory; it is the major component of a certain way of thinking, with many important and familiar implications, which have been part of economics for decades and even centuries.

Pitfalls of Taking it too Literally

...the validity of utility maximization does not depend on its being an accurate description of the behavior of individuals. Rather, it derives from its being the underlying postulate that pulls together most of economic theory; it is the major component of a certain way of thinking, with many important and familiar implications, which have been part of economics for decades and even centuries. Alternatives such as satisficing have proved next to useless in this respect. While attractive as hypotheses, there is little theory built on them; they pull together almost nothing; they have few interesting consequences. In judging utility maximization, we must ask not “Is it plausible?” but “What does it tie together, where does it lead?”

Thus we cannot expect game and economic theory to be descriptive in the same sense that physics or astronomy are. Rationality is only one of several factors affecting human behavior; no theory based on this one factor alone can be expected to yield reliable predictions.

In fact, I find it somewhat surprising that our disciplines have any relation at all to real behavior. (I hope that most readers will agree that there is indeed such a relation, that we do gain some insight into the behavior of *Homo sapiens* by studying *Homo rationalis*.)

Descriptively speaking, then, we can expect our disciplines only sometimes to explain or provide insights into “real” phenomena.

Descriptively speaking, then, we can expect our disciplines only sometimes to explain or provide insights into “real” phenomena. We cannot expect them always to do so, because they are admittedly incomplete. We cannot even say beforehand when we expect them to do so, because we do not yet know how to integrate rational sciences like game theory and economics with non-rational sciences like psychology and sociology to yield accurate predictions.

Descriptively speaking, then, we can expect our disciplines only sometimes to explain or provide insights into “real” phenomena. We cannot expect them always to do so, because they are admittedly incomplete. We cannot even say beforehand when we expect them to do so, because we do not yet know how to integrate rational sciences like game theory and economics with non-rational sciences like psychology and sociology to yield accurate predictions. The criterion for judging our theories cannot be rigid; we cannot ask, is it right or is it wrong? Rather, we must ask, how often has it been useful? How useful has it been?

All this may sound very slippery and unsatisfactory. There are no firm predictions, no falsifiability. If our theory appears not to work, we don't lose any sleep. "Rationality is just one of the relevant factors," we say blandly; "here something else was at work."

“The point of view of this model is not normative; it is not meant to advise the players what to do. The players do whatever they do; their strategies are taken as given.

“The point of view of this model is not normative; it is not meant to advise the players what to do. The players do whatever they do; their strategies are taken as given. Neither is it meant as a description of what human beings actually do in interactive situations.

“The point of view of this model is not normative; it is not meant to advise the players what to do. The players do whatever they do; their strategies are taken as given. Neither is it meant as a description of what human beings actually do in interactive situations. The most appropriate term is perhaps “analytic”; it asks, what are the *implications* of rationality in interactive situations? Where does it lead?

“The point of view of this model is not normative; it is not meant to advise the players what to do. The players do whatever they do; their strategies are taken as given. Neither is it meant as a description of what human beings actually do in interactive situations. The most appropriate term is perhaps “analytic”; it asks, what are the *implications* of rationality in interactive situations? Where does it lead? This question may be as important as, or even more important than, more direct “tests” of the relevance of the rationality hypothesis. (pg. 622)

R. Aumann. *Irrationality in Game Theory*. 1992.

I. Gilboa. *Counter-Counterfactuals*. Games and Economic Behavior, 24, pgs. 175 - 180, 1998.

“It will definitely rain this morning. However, if it doesn’t, it will definitely be very cold this afternoon.”

“Scientific theories are typically neither required nor allowed to specify what would happen if they were refuted, that is, to use counter-counterfactuals.

“Scientific theories are typically neither required nor allowed to specify what would happen if they were refuted, that is, to use counter-counterfactuals. Yet, game theory seems to be replete with counter-counterfactuals. A noncooperative “solution concept” involves a choice of a strategy for each player, and it thus specifies what would happen, as well as what would happen if what should have happened did not happen, and so forth.

“Scientific theories are typically neither required nor allowed to specify what would happen if they were refuted, that is, to use counter-counterfactuals. Yet, game theory seems to be replete with counter-counterfactuals. A noncooperative “solution concept” involves a choice of a strategy for each player, and it thus specifies what would happen, as well as what would happen if what should have happened did not happen, and so forth. It is as if game theorists are never surprised to see their predictions fail; as if we are proven wrong time and again, but we still have the audacity to keep making predictions for the future.”

“...the behaviorist stimulus-response paradigm yields predictions of people’s behavior without any reference to conscious decision making. The dominant paradigm in game theory, however, differs from behaviorism in that it attempts to model *reasoned* choice.”

W. Spohn. *How to make sense of game theory*. in Philosophy of Economics, 1982.

“...game theory...is, to put it strongly, confused about the rationality concept appropriate to it, its assumptions about its subjects (the players) are very unclear, and, as a consequence, it is unclear about the decision rules to be applied.

“Game theory is decision theory about special decision makers, namely about decision makers who theorize decision-theoretically about the other persons figuring in their decision situations.”

The basic difficulty in defining rational behavior in game situations is the fact that in general each player's strategy will depend on his expectations about the other players' strategies.

The basic difficulty in defining rational behavior in game situations is the fact that in general each player's strategy will depend on his expectations about the other players' strategies. Could we assume that his expectations were given, then his problem of strategy choice would become an ordinary maximization problem: he could simply choose a strategy maximizing his own payoff on the assumption that the other players would act in accordance with his given expectations.

The basic difficulty in defining rational behavior in game situations is the fact that in general each player's strategy will depend on his expectations about the other players' strategies. Could we assume that his expectations were given, then his problem of strategy choice would become an ordinary maximization problem: he could simply choose a strategy maximizing his own payoff on the assumption that the other players would act in accordance with his given expectations. But the point is that game theory cannot regard the players' expectations about each other's behavior as given;

The basic difficulty in defining rational behavior in game situations is the fact that in general each player's strategy will depend on his expectations about the other players' strategies. Could we assume that his expectations were given, then his problem of strategy choice would become an ordinary maximization problem: he could simply choose a strategy maximizing his own payoff on the assumption that the other players would act in accordance with his given expectations. But the point is that game theory cannot regard the players' expectations about each other's behavior as given; rather, one of the most important problems for game theory is precisely to decide what expectations intelligent players can rationally entertain about other intelligent players' behavior. This may be called the problem of mutual 'rational expectations'."

“The long and the short of all this: In the absence of more concrete results, at least a general moral may be drawn from the previous discussion. Distinguish strictly between action rationality and epistemic rationality.

“The long and the short of all this: In the absence of more concrete results, at least a general moral may be drawn from the previous discussion. Distinguish strictly between action rationality and epistemic rationality. If your concern is action rationality, then design full decision models for your subjects and determine rational action by the rule of maximizing expected utility;

“The long and the short of all this: In the absence of more concrete results, at least a general moral may be drawn from the previous discussion. Distinguish strictly between action rationality and epistemic rationality. If your concern is action rationality, then design full decision models for your subjects and determine rational action by the rule of maximizing expected utility; and if this alone does not satisfy you, if you search for some account for the epistemic assumptions written into the decision models, then keep strictly to some rules of epistemic rationality as basic and as widely acceptable as possible. Otherwise, theoretical and foundational confusion threatens.” (pg. 267)

J. Kadane and P. Larkey. *The Confusion of Is and Ought in Game Theoretic Contexts*. Management Sciences, 29:12, pgs. 1365 - 1379.

- ▶ Descriptive theories are concerned with empirical phenomena, but stop with a description
- ▶ Explanatory theories go further by addressing “why questions”.
- ▶ Predictive theories discuss what behavior will be.
- ▶ Normative theories can be divided into two types of statements:
 - Speculative statements are nonoperational usually consisting of a goal or criterion with no precise instructions on how one might accomplish the goal
 - Prescriptions are operational in that they give both a goal and a feasible algorithm for achieving that goal.

“ . . . Experiments have been performed which show that individuals do not reason about uncertainty in the way described (by Bruno de Finetti in Theory of Probability). The experiments provide a descriptive view of man's attitudes: de Finetti's approach is normative. To spend too much time on description is unwise when a normative approach exists, for it is like asking people's opinion of $2 + 2$, obtaining an average of 4.31 and announcing this to be the sum. It would be better to teach them arithmetic. I hope that (de Finetti's) book will divert psychologists' attentions away from descriptions to the important problem ... of how to teach people to assess probabilities. ” (Lindley)

Objectives of Game Theory

“... it is crucial that the social scientist recognize that game theory is not descriptive, but rather (conditionally) normative. It states neither how people do behave nor how they should behave in an absolute sense, but how they should behave if they wish to achieve certain ends. It prescribes for given assumptions courses of action for the attainment of outcomes having certain formal “optimum” properties. These properties may or may not be deemed pertinent in any given real world conflict of interest. If they are, the theory prescribes the choices which must be made to get that optimum.”

(Luce and Raiffa)

If you are a player (or advisor to a player) you need two theories.

If you are a player (or advisor to a player) you need two theories. First, you need a prescriptive theory to guide your decisions. You want a decision procedure that results in better outcomes (on average) than alternative procedures.

If you are a player (or advisor to a player) you need two theories. First, you need a prescriptive theory to guide your decisions. You want a decision procedure that results in better outcomes (on average) than alternative procedures. Second, you need a predictive theory of the opponent's behavior—a theory to predict his behavior, to formulate premises for your own decisions.

If you are a player (or advisor to a player) you need two theories. First, you need a prescriptive theory to guide your decisions. You want a decision procedure that results in better outcomes (on average) than alternative procedures. Second, you need a predictive theory of the opponent's behavior—a theory to predict his behavior, to formulate premises for your own decisions. The theory of your opponent's behavior, the decision procedure you believe that the opponents will use, may or may not be like the decision procedure you are using.

If you are a player (or advisor to a player) you need two theories. First, you need a prescriptive theory to guide your decisions. You want a decision procedure that results in better outcomes (on average) than alternative procedures. Second, you need a predictive theory of the opponent's behavior—a theory to predict his behavior, to formulate premises for your own decisions. The theory of your opponent's behavior, the decision procedure you believe that the opponents will use, may or may not be like the decision procedure you are using. If you have good information about the problem your opponent perceives and reason to believe that he will solve it as you would solve it, your own (prescriptive) decision procedure applied to your understanding of your opponent's decision problem may be a useful first approximation to predicting your opponent's behavior.

		Bob	
		L	R
Ann	U	6,3	6,7
	D	10,3	10,7

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

		Bob	
		L	R
Ann	U	2,1	1,0
	D	1,0	0,1

Game 2

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

		Bob	
		L	R
Ann	U	2,1	1,0
	D	1,0	0,1

Game 2

Game 1: U strictly dominates D and L strictly dominates R .

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

		Bob	
		L	R
Ann	U	2,1	1,0
	D	1,0	0,1

Game 2

Game 1: U strictly dominates D and L strictly dominates R .

Game 2: U strictly dominates D

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

		Bob	
		L	R
Ann	U	2,1	1,0
	D	1,0	0,1

Game 2

Game 1: U strictly dominates D and L strictly dominates R .

Game 2: U strictly dominates D , and *after removing* D , L strictly dominates R .

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

		Bob	
		L	R
Ann	U	2,1	1,0
	D	1,0	0,1

Game 2

Game 1: U strictly dominates D and L strictly dominates R .

Game 2: U strictly dominates D , and *after removing* D , L strictly dominates R .

Theorem. In all models where the players are *rational* and there is *common belief of rationality*, the players choose strategies that survive iterative removal of strictly dominated strategies (and, conversely...).

		Bob	
		L	R
Ann	U	3,-3	-1,1
	D	-9,9	3,-3

What is your advice to Ann for optimal play? Should she play U or D on the next move and how should he decide?

What is your advice to Ann for optimal play? Should she play U or D on the next move and how should he decide? Do you have a general strategy to recommend to Ann for the next 100 plays? If so, is your recommended strategy independent of or conditional on Bob's behavior during the 100-play sequence?

What is your advice to Ann for optimal play? Should she play U or D on the next move and how should he decide? Do you have a general strategy to recommend to Ann for the next 100 plays? If so, is your recommended strategy independent of or conditional on Bob's behavior during the 100-play sequence? What literatures would you draw upon for your advice to Ann in this situation?

What is your advice to Ann for optimal play? Should she play U or D on the next move and how should he decide? Do you have a general strategy to recommend to Ann for the next 100 plays? If so, is your recommended strategy independent of or conditional on Bob's behavior during the 100-play sequence? What literatures would you draw upon for your advice to Ann in this situation? Game theory? Classical statistical decision theory? Bayesian decision theory? Psychology?

You may, as most game theorists do, postulate some form of rational behavior for both players as your explanatory/predictive theories. But, as a rational person, yourself, you should not, in light of a substantial body of experimental and natural evidence, expect these theories to work well predictively except in the most simple, contrived settings.

The Is-Ought Confusion

1. Introspective Theorizing
2. Objective and Subjective Rationality
3. The Third Party Perspective
4. Mathematical Tractability

Introspective Theorizing

“With people, in contrast to light beams and water, we usually believe we are dealing with conscious decisions or adaptations in the pursuit of goals, immediate or remote, within the limits of their information and their comprehension of how to navigate through their environment toward whatever their objectives are.

Introspective Theorizing

“With people, in contrast to light beams and water, we usually believe we are dealing with conscious decisions or adaptations in the pursuit of goals, immediate or remote, within the limits of their information and their comprehension of how to navigate through their environment toward whatever their objectives are. In fact we can often ascribe to people some capacity to solve problems—to calculate or to perceive intuitively how to get from here to there.

Introspective Theorizing

“With people, in contrast to light beams and water, we usually believe we are dealing with conscious decisions or adaptations in the pursuit of goals, immediate or remote, within the limits of their information and their comprehension of how to navigate through their environment toward whatever their objectives are. In fact we can often ascribe to people some capacity to solve problems—to calculate or to perceive intuitively how to get from here to there. And if we know what problem a person is trying to solve, and if we think he actually can solve it, and if we can solve it too, we can anticipate what our subject will do by putting ourself in his place and solving his problem as we think he sees it. This is the method of “vicarious problem solving” that underlies *most* of the microeconomics.” (Schelling)

Objective and Subjective Rationality

Even in many of the relatively simple conflict situations created in the laboratory, subjects depart from optimal behavior as best it can be deduced in a particular game using particular assumptions about the rationality of subjects. These results need not be interpreted as evidence of irrational or nonrational behavior because subjects may still be behaving rationally in a weak personalistic sense (i.e., they are solving the problem they perceive in the best way they can devise).

Third Party Perspective

“When we depart from (the) extreme case in which there are two players with strictly opposing interests, game theory has very little advice to offer us.

Third Party Perspective

“When we depart from (the) extreme case in which there are two players with strictly opposing interests, game theory has very little advice to offer us. True, it does establish a relevant vocabulary and a pattern of thinking, but is silent when it comes to telling us precisely how we, as one of the players of the game, should go about analyzing our problem.

Third Party Perspective

“When we depart from (the) extreme case in which there are two players with strictly opposing interests, game theory has very little advice to offer us. True, it does establish a relevant vocabulary and a pattern of thinking, but is silent when it comes to telling us precisely how we, as one of the players of the game, should go about analyzing our problem. One difficulty is that the theory attempts to be neutrally prescriptive, to give advice simultaneously to each player of the game, and it cannot accomplish this except in a small subset of strategic conflict situations.”
(Raiffa)

Mathematical Tractability

“Now it may be that the problems are intrinsically interesting, that the results are elegant and aesthetically pleasing, that the mathematics proceed are real contributions to the discipline of mathematics, and that proving theorems in game theory keeps a large number of academics and journal editors employed. It does not follow that the work is contributing to our knowledge of how individuals and organizations do and should behave in conflict situations.” (Kadane and Larkey, pg. 1375)

The Bayesian Perspective

“ . . . the achievement of determinate solutions for two person, non-zero-sum games through the estimation of subjective probabilities requires the introduction of an assumption to the effect that the individual employs some specified rules of thumb in assigning probabilities to the choices of the other player.

The Bayesian Perspective

“ . . . the achievement of determinate solutions for two person, non-zero-sum games through the estimation of subjective probabilities requires the introduction of an assumption to the effect that the individual employs some specified rules of thumb in assigning probabilities to the choices of the other player. But this is not a very satisfactory position to adopt within the framework of the theory of games. Logically speaking, there is an infinite variety of rules of thumb that could be used in assigning subjective probabilities, the game theory offers no persuasive reason to select anyone of these rules over the others.

The Bayesian Perspective

“... the achievement of determinate solutions for two person, non-zero-sum games through the estimation of subjective probabilities requires the introduction of an assumption to the effect that the individual employs some specified rules of thumb in assigning probabilities to the choices of the other player. But this is not a very satisfactory position to adopt within the framework of the theory of games. Logically speaking, there is an infinite variety of rules of thumb that could be used in assigning subjective probabilities, the game theory offers no persuasive reason to select anyone of these rules over the others. This problem can be handled by introducing new assumptions (or empirical premises) about such things as the personality traits of the players. But such a course would carry the analyst far outside the basic structure of the theory of games, requiring a fundamental revision of the basic perspective of game theory.” (Young)

Negative Conclusions

“Most of game theory is speculative-normative. The only standards that can be applied to this type of theory are the standards that apply to mathematics research: Is the logic correct (internally consistent)? Are the results new? Do the results suggest further work?

Negative Conclusions

“Most of game theory is speculative-normative. The only standards that can be applied to this type of theory are the standards that apply to mathematics research: Is the logic correct (internally consistent)? Are the results new? Do the results suggest further work? Even if the answer to all three questions is “yes” and the results are published, researchers interested in prescription and prediction must then ask, “So what?” There is nothing in these standards to make us sanguine about the utility of accumulated game theoretic results in the short or long run. (Kadane and Larkey, pg. 1377)

Taking the Bayesian norm as **prescriptively** compelling for my play leads me to want the best **description** I can find of my partner/opponent's play.

“Joseph B. Kadane and Patrick D. Larkey (1982) expressed a similar view, but unlike Luce and Raiffa (1957), they eschewed the “strategy aspects”. This ignores *the* fundamental insight of game theory, an insight that is captured by the idea of rational expectations introduced here: that a rational player must take into account that the players reason about each other in deciding how to play. ” (Aumann and Dreze, pg. 81)

R. Aumann and J. Dreze. *Rational Expectations in Games*. American Economic Review, 98:1, pgs. 72 - 86, 2008.

“...a rational player may or may not model his counterpart as rational. He does not violate the axioms of Bayesian rationality if he models his counterpart as not completely rational....players can have whatever models they may have of the other player, which however many uncertain parameters, again, on the marginal distribution of the other player's move affects the optimal decisions.” (Kadane, pg. 407)

J. Kadane. *Principles of Uncertainty*. CRC Press, 2011.

“As the theories rise in strength, they require more and ore restrictive assumptions about what the players believe about each other. Thus more and more is placed into phrases like “common knowledge”, “common knowledge of rationality”, and “common priors.” The usefulness of these special assumptions ahas to be determined case-by-case in applications.” (Kadane, pg. 407)

J. Kadane. *Principles of Uncertainty*. CRC Press, 2011.