

THE JOURNAL OF PHILOSOPHY

VOLUME CIII, NUMBER 5
MAY 2006

<i>page</i>		
215	<i>What Do We Want from a Theory of Justice?</i>	Amaritya Sen
239	<i>Has Game Theory Been Refuted?</i>	Francesco Guala
264	NEW BOOKS: REPRINTS	
265	NEW BOOKS: TRANSLATIONS	

HAS GAME THEORY BEEN REFUTED?*

In his famous "Rational Fools" article, Amartya Sen¹ remarked that if today you were to poll economists of different schools, you would almost certainly find the coexistence of beliefs (i) that the rational behavior theory is unfalsifiable, (ii) that it is falsifiable and so far unfalsified, and (iii) that it is falsifiable and indeed patently false (*ibid.*, p. 325).

Three decades later, Sen's poll would give pretty much the same results. But although the confusion highlighted by Sen was not unjustified in the seventies, this is not the case anymore. Or so I will argue in this paper. This paper defends the view that game theory as commonly understood is not unfalsifiable, that it suffers from important recently discovered empirical anomalies, and that it cannot accommodate the anomalies in its theoretical framework.

A few clarifications are due before we start. First, I will be concerned with game theory in its *descriptive* (or "positive") interpretation only.² Second, asking whether game theory has been refuted does not imply a commitment to some naive form of falsificationism. "Refuted theory" here is taken minimally to mean a theory that suffers from severe empirical anomalies. Finally, this paper completely sidesteps underdetermination and Duhem-Quine problems. It is not a paper about the methodological problems of empirical testing in

*The ideas in this paper have been presented at seminars at the University of Turin, the 2005 Economic Science Association meeting, and the 2005 British Society for the Philosophy of Science conference. A British Academy travel grant is gratefully acknowledged. I should thank the audiences at these meetings, Dario Castiglione, Jim Cox, John Dupré, Matthias Hild, Jim Woodward, Philippe Mongin, and an anonymous referee for their comments and suggestions. I was also inspired by several conversations with Dan Hausman and Chris Starmer. All the remaining mistakes are mine.
¹"Rational Fools: A Critique of the Behavioral Foundations of Economic Theory," *Philosophy and Public Affairs*, vi (1977): 317-44.
²On the interaction between normative and descriptive interpretations of game theory see Joseph Kadane and Patrick Larkey, "The Confusion of Is and Ought in Game Theoretic Contexts," *Management Science*, xxix (1983): 1365-79. On the theory of individual decision making (in nonstrategic situations), see, for example, Francesco Guala, "The Logic of Normative Falsification: Rationality and Experiment in Decision Theory," *Journal of Economic Methodology*, vii (2000): 59-93; and Chris Starmer, "Normative Notions in Descriptive Dialogues," *Journal of Economic Methodology*, xii (2005): 277-89.

general, but rather deals with the much more specific issue of game theory's relation with empirical evidence.

IS GAME THEORY A TAUTOLOGY?

Like other successful scientific theories, game theory can be (and is) interpreted in different ways. Thus first of all it will be necessary to clarify what we are dealing with: *What* has been refuted (if at all)? Under a standard interpretation, game theory is aimed at modeling and understanding people's behavior, in particular when the outcome of their actions depends not only on what they do, but also on other people's decisions. According to an influential textbook,³ for example,

Game theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact. The basic assumptions that underlie the theory are that decision-makers pursue well-defined exogenous objectives (they are *rational*) and take into account their knowledge or expectations of *other* decision-makers' behavior (they reason *strategically*) (*ibid.*, p. 1).

This characterization rules out, first of all, so-called "mass-action" interpretations of game theory, as for instance in evolutionary game theory models. Game theory, under the standard "epistemic" interpretation, models individual behavior as determined by players' preferences and beliefs, under strong rationality assumptions. This does not mean that mass-action interpretations are less interesting or promising than the epistemic one, but simply that the latter is by far the most developed, used, and—crucially, for our purposes—empirically tested version of game theory.⁴

The standard textbook definition also rules out interpreting game theory as a branch of pure mathematics. Although it is common to distinguish "pure" from "applied" game theory, the former is *not* a piece of uninterpreted formalism. The Osborne and Rubinstein quote significantly mentions decision makers, expectations, knowledge, objectives, and behavior, as the subject matter of game theory. Of course there is no mention of what these objectives may be, nor indeed of what the expectations are about. "Decision makers" is also left vague enough to be applicable to various real-world entities (individuals, firms, nations). But this just means that "pure" game

³ Martin Osborne and Ariel Rubinstein, *A Course in Game Theory* (Cambridge: MIT, 1994).

⁴ The distinction between these two interpretations goes back to John Nash's seminal Ph.D. dissertation.

theory is *very abstract*, rather than purely formal.⁵ The business of specifying preferences, beliefs, and the identity of decision makers is left to applied social scientists, and presumably varies depending on the intended domain application.

There are some dissenters, however. "Pure" game theorists often claim that their theory is just a *tautology*. Ken Binmore⁶ is a typical example:

Mathematical theorems are tautologies. They cannot be false because they do not say anything substantive. They merely spell out the implications of how things have been defined. The basic propositions of game theory have precisely the same character (*ibid.*, p. 96).

As with other claims of this kind, it is useful to ask first what is *meant* by "tautology," and then *what* is supposed to be a tautology.⁷ A tautology is, strictly speaking, a proposition that is true in virtue of the definition of its logical terms. This is probably not what Binmore has in mind. Most likely, he is thinking of *sets* of propositions such as:

- (1) Rational agents always choose uniquely dominant strategies.
- (2) "Defect" is the only dominant strategy in the one-shot Prisoner's Dilemma game.
- (3) Therefore, rational agents always defect in the one-shot Prisoner's Dilemma game.

This is an argument, not a proposition. It is a valid argument, in the standard logical sense: if the premises are true, the conclusion must also be true, independently of the truth of the premises. But whether each proposition in the argument is true or false is another question altogether. In order to figure out whether it is true that rational agents always defect in the one-shot Prisoner's Dilemma game, we need to check whether rational agents always choose uniquely dominant strategies, and whether "defect" is the only dominant strategy in the one-shot Prisoner's Dilemma.⁸

⁵ A similar point could be made (more elaborately, perhaps) by looking at the sort of solution concepts that are explored and considered legitimate in game theory. Rationality plays a prominent role in this activity, and rationality is not a purely formal concept. It is rationality for decision makers, and more specifically decision makers of the *human* kind.

⁶ *Game Theory and the Social Contract, Volume 1: Playing Fair* (Cambridge: MIT, 1994).
⁷ See, for example, Elliott Sober, *Philosophy of Biology* (Boulder: Westfield, 1993), pp. 68–72.

⁸ On this point, see also Daniel Hausman, "Testing Game Theory," *Journal of Economic Methodology*, xii (2005): 211–23.

Table 1.

	Cooperate	Defect
Cooperate	2, 2	0, 3
Defect	3, 0	1, 1

But how do we check that (1) and (2) are true? Binmore and the other supporters of the "tautology" interpretation presumably argue that (1) and (2) are *analytic*: they are true in virtue of the meaning of the nonlogical terms that are in them. So "rational agent" is defined as the kind of decision maker that always chooses uniquely dominant strategies. And the one-shot Prisoner's Dilemma game (Table 1)⁹ is by definition a game where "defect" is the only dominant strategy.

Of course there would be nothing to test, if game theory was nothing but this set of propositions. "Tautologists" like Binmore however do not conclude that game theory is empirically useless:

One reason for emphasizing the tautological nature of these tools of the game theory trade is to direct attention of the critics away from the question of how a formal game is *analyzed* to the more rewarding question of how problems of human interaction should be *modeled* as formal games (*ibid.*, p. 169).

So what can be tested presumably is the claim that certain game theory models can be used to understand (predict, explain) certain specific real-world situations.¹⁰ This seems analogous to what some versions of the semantic view of scientific theories call a "theoretical hypothesis" (but perhaps the expression "empirical hypothesis" would be less misleading).¹¹ A testable hypothesis would be, for example, that the one-shot Prisoner's Dilemma game can be used to explain (predict, and so on) the behavior of experimental subjects playing the game in Table 2.

⁹ In representing games in normal (or strategic) form, I follow the usual conventions: the first number in each cell (or consequence, or outcome) represents the payoff of the "Row" player, the second number the payoff of the "Column" player. Each row/column represents a strategy (or act) available to that particular player.

¹⁰ See also Binmore, "Why Experiment in Economics?" *Economic Journal*, cix (1999): F16-24, p. F18: "[a consolidating experimentalist] does not see as his task as asking whether economics works or not. He already knows that sometimes it does and sometimes it does not. The task is to classify economic environments into those where the theory works and those where it does not."

¹¹ See, for example, Ronald Giere, *Explaining Science* (Chicago: University Press, 1988), chapter 3.

Table 2.

	Left	Right
Up	\$2, \$2	\$0, \$3
Down	\$3, \$0	\$1, \$1

So far it looks as if the tautologists are just quibbling about what a scientific theory is. According to them, a theory is just the formalism. According to most scientists and philosophers, it is an *interpreted formalism*, or a formalism plus a set of "bridge principles" that connect it to empirical reality. But in fact some tautologists are more radical than that. Take Binmore again: "Game theorists understand the notion of a payoff in a sophisticated way that makes it tautologous that players act as though maximizing their own payoffs."¹² This strategy is defended by drawing an analogy with biology:

[In applications of game theory in evolutionary biology] the payoffs are...defined in terms of an organism's "fitness". This makes it a tautology that the organisms that survive will seem to behave as though seeking to maximize their payoffs in the game under study. The same tautological flavor is also to be found when *rational* behavior in games is studied, provided that the payoffs in the game are interpreted according to the theory of *revealed preference*... (*ibid.*, p. 104).

Revealed preference theory was originally proposed as an attempt to *reduce* the notion of preference to that of observed choice.¹³ There is remarkable disagreement nowadays about what revealed preference theory really amounts to in neoclassical economics,¹⁴ but tautologists like Binmore seem to stick close to the original understanding: "In revealed preference theory...choice behavior becomes the primitive. One observes some of the choices that a player makes and then argues that he is making choices *as though* he were equipped with a preference relation" (*ibid.*, pp. 105-06).

The first thing to notice is that Binmore's analogy with fitness in evolutionary biology is far from unproblematic. According to a common interpretation, the fitness of a trait X is not whatever

¹² *Game Theory and the Social Contract*, p. 98.

¹³ Paul Samuelson, "A Note on the Pure Theory of Consumer's Behavior," *Economica*, v (1938): 61-71.

¹⁴ See, for example, Daniel Hausman, "Revealed Preference, Belief, and Game Theory," *Economics and Philosophy*, xvi (2000): 99-115; Keith Dowding, "Revealed Preference and External Reference," *Rationality and Society*, xiv (2002): 259-84.

proportion of offsprings with X are observed in the next generation. Fitness is the *expected* reproductive success of a trait, or the probability—given the forces of natural selection—that a member of the next generation of offsprings will have trait X . Probability here can be (and usually is) interpreted as a *propensity*, rather than a frequency. This propensity may or may not be reflected approximately in the frequency of X in the next generation. It may not be because other factors besides natural selection can interfere with the heritability of X : random genetic drift, typically, or mutation, migration, recombination. It is not true that “We know who will survive because they are fittest” and “We know who are fittest because they survive” (*ibid.*, p. 98).

Interpreted this way, fitness is not a “tautology.” This of course does not, by itself, prevent one from interpreting utility in terms of revealed preferences. But there are other, independent reasons not to do so. Preferences cannot be “read off” directly from behavior because choices are determined by individual preferences *and beliefs*. This point is perhaps obscured by the convention in game theory of assuming common knowledge of the game that is being played (“I know that you know that I know... (and so forth) that such and such strategies are available, with such and such payoffs, and that all players are rational”). Of course *if* we assume common knowledge of the game, then we can derive preferences from observed choices. Whether the derivation is correct, however, depends crucially on the correctness of the common knowledge assumption.

It is not clear however why one should want to take such a strong empiricist stance regarding preferences, if one is willing to simply assume a priori a certain configuration of players’ beliefs.¹⁵ Why accord to the two key variables of game theory a different status? In fact the fathers of rational choice theory were fully aware of the need *mutually* to determine, by means of observable procedures, the structure of individual preferences *and* beliefs.

SAVAGE MEASUREMENTS

I will keep using an old-fashioned philosophical distinction, according to which a scientific theory can be analyzed in two sets of elements. On the one hand, a set of abstract concepts connected by theoretical relationships (or “internal principles”). On the other, some “bridge principles” that link the concepts with empirical reality

or measurement procedures.¹⁶ What are the bridge principles of game theory? In *The Theory of Games*, John von Neumann and Oskar Morgenstern¹⁷ famously introduce a procedure for measuring (or “operationalizing”) cardinal utility. Expected utility theory is the “bridge” connecting game theory with the real world.

The story has become part of the folklore of game theory. Von Neumann was happy to build the theory on the generic notion of “payoff.” Indeed, he conceived of payoffs basically as *objective* (material, possibly monetary) properties of outcomes or consequences. Morgenstern persuaded von Neumann to recast the theory in utility jargon, in order to appeal to the average neoclassical economist. Von Neumann quickly scribbled on the back of an envelope the measurement procedure that was to appear in the second (1947) edition of *The Theory of Games*.

Von Neumann and Morgenstern’s measurement procedure of utility is well known: starting with three outcomes x , y , and z such that $x > y > z$, the utility scale is normalized by assigning $U(x) = 1$ and $U(z) = 0$. A subject is then presented with the choice between the prospect y for sure and the prospect $[p, x; (1 - p), z]$; if she prefers either the risky prospect or the sure outcome, the probabilities in the former are varied until she is indifferent. At this stage, the utility of y can be computed by $U(y) = p$, and the procedure iterated for any value between x and z .

This clearly assumes that we already know subjects’ probability assignments or beliefs. In practice, since von Neumann and Morgenstern restrict their discussion to lotteries with objective probabilities, we just assume that subjects’ assignments mirror the objective probability measures (their beliefs are accurate representations of the world). When this is not the case, subjective probability assignments must be worked out from people’s choices, in a way that is symmetric to the one above. Start with people’s preferences about sets of consequences, and then infer which states of the world are considered more probable from the observation of individuals’ acts (or choices of strategy). In practice, it is usually assumed that more money

¹⁶ We do not have to assume that these two elements are precisely separable. The distinction is useful only for heuristic purposes, and should not be read as implying any deep philosophical commitment on the theory-observation distinction, or on the nature of scientific theories. In fact, it should be broadly compatible both with old-fashioned “syntactic” approaches (like the “standard view” of theories) and with more fashionable “semantic” approaches (where the “bridge principles” are sometimes called “theoretical hypotheses,” see above).

¹⁷ *The Theory of Games and Economic Behavior* (Princeton: University Press, 1944).

is preferred to less, and then individuals are offered money prizes associated with certain types of event. But again one has to take a firm fulcrum (preferences for the measurement of beliefs, beliefs for the measurement of preferences) for the lever of expected utility theory to lift any weight.

These procedures were first sketched by Frank Ramsey and then refined by Leonard Savage.¹⁸ Since the latter's systematization has become standard in contemporary game theory, I shall refer to it as the *Savage measurement procedure*. According to "kosher" game theory, the application of the theory should start with a Savage measurement, that is, with the identification of players' utilities and beliefs. Once this has been done, one can begin to figure out what kind of game they are playing, and apply the solution concepts of game theory.¹⁹ Most experimental game theory is not kosher from this respect, however. This causes problems, because the results of experiments can be (and are) criticized for not testing the theory at all.

WHAT KIND OF GAME ARE WE PLAYING?

Take the one-shot Prisoner's Dilemma game again (Table 1 above). Here the free riding strategy is Down for Row and Right for Column, leading to a payoff of 1 unit each. These strategies are uniquely dominant, and determine the unique rational solution (the Nash solution or Nash equilibrium) of the game. Informally, whatever Column will do, Row will be better off by playing Down; and similarly Column will always be better off by playing Right, given what the opponent does. But both would be better off if they played Up-Left. The Nash equilibrium in a one-shot Prisoner's Dilemma game is not a Pareto-optimal solution.

¹⁸ Ramsey, "Truth and Probability" in Richard Braithwaite, ed., *The Foundations of Mathematics and Other Logical Essays* (New York: Routledge, 1931), pp. 156-98; Savage, *The Foundations of Statistics* (New York: Wiley, 1954).

¹⁹ Sometimes Binmore writes (and speaks, in conversation) as if his revealed preference theory was nothing but an application of Savage's measurement procedure. He says for example ("hawk" here stands for "defect," "dove" for "cooperate"): "the preference (dove, hawk) \succ (hawk, hawk) is deduced from the prior information that, if Adam knew he had to choose between only (dove, hawk) and (hawk, hawk), then he actually would choose (hawk, hawk)" (*Game Theory and the Social Contract*, p. 106). Notice that according to this formulation one does *not* observe behavior in a Prisoner's Dilemma game first, and *on the basis of this*, define the preference scales. Rather, one elicits the preference scales by pair wise comparison of the various outcomes (as in a simplified Savage measurement—more about this below), and *then* uses such information to predict behavior in the Prisoner's Dilemma game. This does not make a "tautology" out of game theory, because of course there may well be a mismatch between the preferences as "revealed" in the two situations. And it does not reduce the concept of preference to that of choice, because in order to identify preferences in the simplified Savage procedure you need to postulate a given belief structure.

Table 3.

	Left	Right
Up	2, 2	0, 0
Down	0, 0	1, 1

As is well known, in a "standard" Prisoner's Dilemma experiment a considerable number of subjects play cooperatively (Up-Left).²⁰ The game played in such experiments, however, is not necessarily the one in Table 1. It is the one already presented in Table 2 above. The "naive" interpretation of these experimental results is that many human beings (fortunately) do not behave as predicted by the theory. But a number of game theorists reject this interpretation. They say, for example, that the preference rankings of the subjects who play cooperatively in these experiments are inadequately represented by the numbers in the classic Prisoner's Dilemma matrix. Or, in other words, that Table 1 and Table 2 are not necessarily isomorphic for all experimental subjects.

According to the orthodox interpretation of game theory, the numbers in Table 1 represent the (ordinal) structure of agents' preferences. Thus, the argument goes, if we observe anomalous behavior in the experiment, then the initial conditions postulated in the model were probably not instantiated in the experiment. Subjects were not *really* playing the Prisoner's Dilemma, but another game of their choice. A game like the one in Table 3, for example.

This game has *two* Nash equilibria. The transformation of Row's payoff from 3 to 0 in the Down-Left cell can be explained, for example, by Row's anticipated feelings of guilt in letting Column down, given that the latter was willing to cooperate (and vice-versa for Column in Up-Right). The Prisoner's Dilemma is then transformed in a coordination game, where Up-Left is a legitimate equilibrium.

Binmore has been a vocal supporter of this interpretation.²¹

[those who are anxious to deny that people seek only their own narrowly conceived selfish ends] argue that the players may care about the welfare of their opponents, or that they may actively want to keep their promises out of feelings of group solidarity or because they would otherwise suffer the pangs of a bad conscience. Such players will *not* be playing the

²⁰ See, for example, Anatol Rapoport and Albert Chammanah, *Prisoner's Dilemma: A Study in Conflict and Cooperation* (Ann Arbor: Michigan UP, 1965).

²¹ *Fun and Games: A Text on Game Theory* (Lexington, MA.: D.C. Heath, 1992).

Prisoners' Dilemma. They will be playing some other game with different payoffs (*ibid.*, pp. 313-14).

In other words: it is a common fallacy to interpret the payoffs in a game theory model as representing physical properties or (typically) monetary outcomes. They represent utilities. And people may derive utility from (or, more precisely, may have preferences about) whatever features of the consequences they happen to care about—including other players' payoffs, one's own anticipated guilt, regret, and so forth.

This is in line with the liberal tradition that informs much of economic theory. According to this tradition, it is important that social scientists remain neutral about the contents of people's tastes and beliefs. From a normative viewpoint (and economics is a theory of rationality as well as a theory of actual behavior) the requirements imposed on preferences and beliefs should be purely formal, of the *consistency* type. (Remember Hume: it is perfectly rational to prefer the end of the world to a scratch on my finger.) So, even though economists often assume for simplicity that outcomes are defined strictly in monetary terms, this ought not to be always the case: outcomes can (and should, whenever required) be redescribed or "refined" to include whatever properties the agents find relevant.

This idea (the possibility of "refining" the consequences) has provoked controversy since the early days of expected utility theory.²² An explicit attempt to legitimize it is due to Peter Hammond, who puts the "refining" strategy right at the core of his "consequentialism." Consequentialism, as usually understood, is the idea that rationality is forward-looking and independent of procedural considerations.

²² A (probably incomplete) list of contributions to this debate includes Paul Samuelson, "Utility, Preference and Probability," in Joseph Stiglitz, ed., *The Collected Papers of Paul Samuelson, Volume I* (Cambridge: MIT, 1952), pp. 127-36; Amos Tversky, "A Critique of Expected Utility Theory: Descriptive and Normative Issues," *Erkenntnis*, ix (1975): 163-73; Sen, "Rationality and Uncertainty," *Theory and Decision*, xxii (1985): 109-27, and "Internal Consistency of Choice," *Econometrica*, lxi (1993): 495-521; Mark Machina, "Dynamic Consistency and Non-Expected Utility Models of Choice Under Risk," *Journal of Economic Literature*, xxvii (1989): 1622-68; Michael Bacharach, "Preferenze Razionali e descrizioni," in Maria Carla Galavotti and Guido Gambetta, eds., *Epistemologia ed economia* (Bologna: CLUEB, 1988), pp. 43-68; Peter Hammond, "Consequentialist Foundations for Expected Utility Theory," *Theory and Decision*, xxv (1988): 25-78; Robert Sugden, "Rational Choice: A Survey of Contributions from Economics and Philosophy," *Economic Journal*, ci (1991): 751-85; John Broome, *Weighing Goods* (Cambridge: Blackwell, 1991); Paul Anand, *Foundations of Rational Choice Under Risk* (New York: Oxford, 1992); and Bertrand Munier, "Comment," in Kenneth Arrow, Ennio Colombaro, Mark Perlman, and Christian Schmidt, eds., *The Rational Foundations of Economic Behavior* (New York: St. Martin's, 1996), pp. 43-47.

Consequentialism requires everything which should be allowed to affect decisions to count as a relevant consequence—behavior is evaluated by its consequences, and nothing else (*op. cit.*, p. 26).

Hammond's "consequentialism" tries to transform genuine consequentialism into an entirely general but admittedly trivial proposition:

If regrets, sunk costs, even the structure of the decision tree itself, are relevant to...behavior, they are therefore already in the consequence domain. Indeed the content of a...theory of behavior is then largely a matter of what counts in practice as a relevant consequence, rather than whether consequentialism and other abstract axioms are satisfied (*op. cit.*, p. 26).

Many people find this problematic. It is often argued that if you are allowed to redefine consequences (or the argument of the utility functions) as you like, then economic theory becomes empty. This critique goes back a long time,²³ but is essentially misguided. It is a false dilemma to stipulate that *either* the payoffs are to be interpreted as money, *or* the theory is empty. People have preferences, and they are preferences *about* something. There is no reason a priori to assume that all individuals care about the same things. One difficult empirical task in testing game theory is precisely to figure out what they care about. Once we know what they care about, we can check whether they are rational—that is, whether the models of rational choice provide a correct account of human behavior. The theory is not empty. It can be used to say something substantial about the real world.

To illustrate, consider an analogy with classical mechanics. Newton's law of universal gravitation ($F=ma$) can be interpreted (and usually is interpreted) as a description of the relation between mass, acceleration, and force once *all* the forces at work in a given system have been taken into account. As such, it seems hard to falsify. Imagine a fictional classical (pre-relativity) physicist observing a prima facie violation of the law; before coming to any conclusion about the validity of the universal law of gravitation, she would begin to search for some unobserved forces at work on the system under study. Such forces could be gravitational, but also of some other kind, not (yet) described by Newtonian mechanics (electromagnetic forces, for example). But the physicist's behavior would not make the law empty:

²³ See, for example, Terence Hutchison, *The Significance and Basic Postulates of Economic Theory* (New York: Kelley, 1938), pp. 114-15; for recent versions, see Machina, "Dynamic Consistency," pp. 83-87; Alexander Rosenberg, *Economics: Mathematical Politics or Science of Diminishing Returns?* (Chicago: University Press, 1992), chapter 5; and Anand, *Foundations of Rational Choice Under Risk*, pp. 105-06.

The message is: one cannot proceed to make any inference about the empirical status of game theory without figuring out first what individuals care about. If they care about fairness, for example, it is still entirely possible that they are rational maximizers of utility—once the latter has been defined appropriately to capture fairness considerations. The definition of the utility function as ranging exclusively over one's own monetary outcomes is unnecessarily restrictive and in principle can be relaxed to allow for more varied (and interesting) individual preferences.

PROBLEMS WITH REFINEMENTS

The possibility of refining the options does not make the theory empty. There are still two ways in which game theory can be refuted. First, suppose that cooperators in the one-shot Prisoner's Dilemma really did care about their own monetary outcomes only (and preferred more money to less). Perhaps they do not understand the concept of dominance; perhaps they do not pay enough attention to the structure of the game. If either of these hypotheses were correct, the theory would turn out to be false.

This is just a speculation. As a matter of fact most experimental game theorists (including Cox and Weibull) agree that to assume selfish preferences over money only is too restrictive. There is more in the argument of our utility functions, and this extra content should be accommodated in the framework of game theory.

This leads to the second way in which game theory could turn out to be refuted. Suppose that subjects cared about something else besides their own monetary gain, but that game theory was not flexible enough to accommodate this "something else" within its framework. Were the bridge principles unable to account for the observations, in particular, the only option left would be to modify the theory of rational play (the "hard core" of game theory).

This possibility is rarely discussed, perhaps because of a taken-for-granted analogy between folk psychology and formal rational choice theory. According to our everyday understanding, *any* observed pattern of behavior can in principle be rationalized by imputing an appropriate—perhaps highly sophisticated—structure of beliefs and desires. The intertemporal coherence of beliefs and desires can also always be restored by simply modifying their contents—no matter what the individual does or is going to do next. Alexander Rosenberg is the philosopher who has most forcefully argued that "the fundamental explanatory strategy of economic theory is of a piece with that of our ordinary explanations of human action" (*op. cit.*, p. 118). And the view that "no intentional action can be internally irrational" (in

the law says that *once* all the forces have been taken into account, they obey the universal law of gravitation—and in some cases we might well be wrong in computing the total force in the system.²⁴

So, there is a division of labor in classical mechanics: on the one hand general principles like the law of universal gravitation are supposed to hold for all bodies and all forces. On the other, more specific laws connect specific kinds of forces to other quantities of interest. For example: gravitational forces are described by equations like $F = G(m_1 m_2 / r^2)$; electrostatic forces by Coulomb's law ($F = q_1 q_2 / 4\pi\epsilon_0 r^2$), and so on. Similarly in economics one can envisage a division of labor between general principles of rational play like those of game theory, and the more specific task of connecting payoffs (utilities) with empirical reality, by specifying the argument of the utility function.

Game theorists like James Cox and Jorgen Weibull have expressed this viewpoint clearly in recent articles.²⁵ Here is Cox, for example:

In their seminal work on game theory, von Neumann and Morgenstern (1944, 1947) thought it necessary to simultaneously develop a theory of utility and a theory of play for strategic games. In contrast, much subsequent development of game theory has focused on analyzing the play of games to the exclusion of utility theory. In the absence of a focus by game theorists on utility theory, it is understandable that experimentalists testing the theory's predictions have typically assumed that agents' utilities are affine transformations of (only) their own monetary payoffs in the games. This interpretation of game theory incorporates the assumptions that agents do not care about others' (relative or absolute) material payoffs or about their intentions. There is a large experimental literature based on this special-case interpretation of the theory, which I shall subsequently refer to as the model of "self-regarding preferences." The part of the literature concerned with public goods experiments and trust and reciprocity experiments has produced replicable patterns of inconsistency with predictions of the model of self-regarding preferences. For example, the patterns of behavior that have been observed in one-shot trust and reciprocity games are inconsistent with the subgame perfect equilibria of that model. But this does *not* imply that the observed behavior is inconsistent with game theory, which is a point that has not generally been recognized in the literature (*op. cit.*, pp. 260–61).

²⁴ This fictional example is not far from historical reality, as shown by Imre Lakatos, "Falsificationism and the Methodology of Scientific Research Programmes," in *Philosophical Papers, Volume 1* (New York: Cambridge, 1970), pp. 8–101, and others.

²⁵ Cox, "How to Identify Trust and Reciprocity," *Games and Economic Behavior*, xlvii (2004): 260–81; see also Weibull, "Testing Game Theory," in Stephan Huck, ed., *Advances in Understanding Strategic Behavior* (New York: Palgrave, 2004), pp. 85–104, especially pp. 85–86.




Figure 1

the above sense), has been dubbed the *Plato principle* by Donald Davidson.²⁶

The Plato principle, however, is false when applied to game theory, because the formal theory of rational choice is less flexible than folk psychology, and as a consequence has more limited explanatory resources. This applies to the theory both in its normative and in its descriptive version, but given the focus of this paper, let us phrase it in a descriptive idiom. There is a tension between certain refinements of the consequences and the Savage measurement procedure. Consider the sequential Prisoner's Dilemma game in Figure 1.²⁷ The first monetary prize between brackets at each end-node is A's monetary payoff, the second one is B's monetary payoff; the letters *f*, *g*, and so on, are arbitrary labels assigned to the consequences. Experimenters report data where 40% of first-movers (A) play a *trust* move (Up) in a game of this kind, and second-movers (B) respond Left 35% of the time.²⁸

Can we "refine" A and B's preferences in order to account for this behavior? It depends. Suppose A and B are motivated by two sets of

considerations. They prefer more money to less, but also have a taste for fairness. In equation (1) deviations from fairness are defined with respect to a standard "fair payoff" π^f :

$$(1) V_i = \pi_i - \alpha_i(|\pi_i^f - \pi_i| - \sum_{j \neq i} \beta_j |\pi_j^f - \pi_j|).$$

Player *i*'s material payoffs are denoted by π_i , and the parameter α introduces *i*'s concern for the pattern of distribution of the payoffs (as opposed to his own payoffs). When $\alpha=0$, the player is purely selfish (his utility V_i is a direct function of his material payoff only). Everything on the right-hand side of α represents fairness considerations. When the standard of fairness is equality of payoffs ($\pi_i^f = \pi_j^f$), this becomes a model of inequality aversion.²⁹ The parameter β weighs the impact of the first set of fairness considerations (inequality in one's payoffs) against that of the second one (inequality in other players' payoffs).

This is just one example, representative of a wider class of models of fairness preferences. It suffices to raise two points. First, these are not alternatives to the theory of rational play at the core of game theory. They are attempts to pin down the bridge principles connecting game theory with observed behavior (or better: with observable features of the outcomes, like the distribution of monetary gains π_i). Such models play in the hands of the refiners, and their empirical successes in predicting experimental data are not failures of game theory.

Second, these models add to the refiners' weaponry because they are *genuinely* consequentialist in spirit. The extensive form of the game in Figure 1 is quite redundant. The equation in (1) can be applied simply by looking at the normal form of the game (Table 2 above). In principle, this is fully compatible with the orthodox Savage measurement procedure. Indeed, equation (1) could be read as a generalization of some class of preference structures observed by means of the Savage procedure.

How would this measurement work, in practice? Suppose we want to measure B's utilities. All we have to do is to present player B with the choices in Figure 2. B's preferences, as revealed in these binary-choice tasks, will suffice to predict her behavior in the sequential

²⁶ See, for example, his "Paradoxes of Irrationality," in Richard Wollheim and James Hopkins, eds., *Philosophical Essays on Freud* (New York: Cambridge, 1982), pp. 289–305, see p. 294.

²⁷ Borrowed from Hausman, "Revealed Preference," p. 107.

²⁸ See, for example, Kenneth Clark and Martin Sefton, "The Sequential Prisoner's Dilemma: Evidence on Reciprocation," *Economic Journal*, cxi (2001): 51–68.

²⁹ See for example, David Levine, "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics*, 1 (1998): 593–622; Ernst Fehr and Klaus Schmidt, "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, cxiv (1999): 817–68; Cary Bolton and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity and Cooperation," *American Economic Review*, xc (2000): 166–93.




Figure 2

Prisoner's Dilemma game, if they conform to the preference structure of (1). In this scenario, $f=f'$, $g=g'$, $h=h'$, and $k=k'$.

But let us suppose instead that B is not primarily motivated by distributional concerns. The reason she plays the equal payoff (f) in Figure 1 could be that she is *reciprocating* A 's first move. She chooses Left in response to a "trust" move because she wants to be kind towards A , but has no reason to do so when A has not played at all. In this scenario, looking at B 's choices in Figure 2 will be useless.³⁰

MODELING RECIPROCITY

Reciprocity is a tricky notion. Consider the *basic reciprocity formula*:

$$(2) V_i = \pi_i + \rho[\kappa(\cdot)\sigma(\cdot)]$$

Here ρ is a parameter weighing the impact of reciprocity considerations against the usual self-interested preference for one's own monetary gains (π_i). κ is a *kindness function*, representing how "nice" or "offensive" j 's behavior is towards i . The other function, σ , captures i 's reaction to j 's behavior. A *negative* kindness combined with a *positive* reaction (or vice-versa) results in a reduction of overall utility (V_i). Reacting negatively to negative kindness, and positively to positive kindness, leads to an increase of overall utility.

At first sight, it may look as if the arguments of these two functions could be filled in consistently with the spirit of genuine consequentialism. But this is unlikely. Consider the reaction function first: a negative reaction is one that *hurts* j , typically by reducing her payoff

with respect to a given standard or reference point. The standard is most naturally interpreted as what j *would have* achieved, had i played a different strategy. We shall represent this possible payoff with π_i^j . Similarly, j 's kindness towards i depends not only on i 's relative payoff with respect to some absolute standard (say, equality, as in most models of fair preferences), but on what i *could* have achieved, had j played differently. Again, we will use π_i^j to represent such counterfactual payoff, where the counterfactual mark (') refers to the fact that it is a *possible* payoff—an outcome that could in principle be achieved in *this* game.³¹

$$(3) V_i = \pi_i + \rho[\kappa(\pi_i - \pi_i^j)\sigma(\pi_j - \pi_j^i)]$$

The most advanced models of reciprocity in this spirit make use of the tools provided by *psychological game theory*.³² In psychological game theory the payoffs depend not just on the material structure of the consequences, but also on players' *beliefs*. Following Matthew Rabin's seminal work,³³ reciprocity can be represented as involving a calculation of what player i believes that j believes that i will play. Player j 's intention is perceived as "nice" when i believes that j believes that i will make a "nice" move, and yet j chooses a strategy that reduces i 's monetary payoffs with respect to the reference point.³⁴

A discussion of psychological game theory is besides the scope of this paper. How counterfactual considerations are modeled is not crucial: perhaps intentions are not even relevant *all* the time; perhaps in some cases B reacts to A 's "nasty" moves only to signal that she is making a mistake, or in order to push the outcome towards a Pareto-superior solution (B might be driven by selfish motives, that is, and might have a crude stimulus-response model of A 's behavior).³⁵

³¹I am simplifying a lot here. The counterfactual "reference point" in $\kappa(\pi_i - \pi_i^j)$ is in many occasions better represented as π_i^j , where j stands for the "fairest" payoff that could have been achieved in the game. It is likely that π_i^j should be identified in each game by comparing i and j 's various possible payoffs.

³²First proposed by John Geanakoplos, David Pearce, and Ennio Stacchetti, "Psychological Games and Sequential Rationality," *Games and Economic Behavior*, 1 (1989): 60-79.

³³"Incorporating Fairness into Game Theory and Economics," *American Economic Review*, lxxxiii (1993): 1281-302.

³⁴The two prominent models in this tradition are Armin Falk and Urs Fischbacher, "A Theory of Reciprocity," Working Paper 6 (2000), Institute for Empirical Research in Economics, University of Zurich; and Martin Dufwenberg and Georg Kirchsteiger, "A Theory of Sequential Reciprocity," *Games and Economic Behavior*, xlvii (2003): 268-98.

³⁵As a matter of fact, there is some puzzling evidence suggesting that in dilemma experiments some subjects "reciprocate" even when they are playing against a computer; see Daniel Houser and Robert Kurzban, "Revisiting Kindness and Confusion in Public Goods Experiments," *American Economic Review*, xcii (2002): 1062-69.

³⁰ See also Hausman, "Revealed Preference," p. 108, for this point.

What really matters is the role of counterfactual considerations themselves. The identification of the preferences associated with each consequence requires an examination not just of what will happen, but also of *what might have happened*. To put it another way, the utilities are context-dependent: they do not just depend on the outcomes taken in isolation, but on the *whole structure of the game*.

RECIPROCIITY CANNOT BE "REFINED"

Remember Hammond: "Consequentialism requires everything which should be allowed to affect decisions to count as a relevant consequence"; if reciprocity affects decisions, it must be accounted for at the moment of describing the consequences, according to consequentialism.

This turns out to be problematic in the Savage framework, however. The problem has been around for a long time. In a letter from Robert Aumann to Savage, in 1971,³⁶ the issue is first raised of how to distinguish *acts* (say, taking an umbrella) from *consequences* (for example, getting wet). Aumann points out that the distinction is fuzzy. Savage agrees but says in a very pragmatic vein that this may not matter as long as "I seem to be able to couch my decision problems in terms of them"; for "...a consequence is in the last analysis an experience" (*ibid.*, p. 79).

But this is overoptimistic. Savage defines *acts* as functions from *states of the world to consequences*. (For example: the act "taking the umbrella" assigns the consequence "staying dry" both to the state "it rains" and to the state "it does not rain"; "not taking the umbrella" assigns the consequence "getting wet" to the state "it rains" and the consequence "staying dry" to the state "it does not rain.") All the existing versions of expected utility theory prove the existence of the utility function (representation theorem) from the assumption that decision makers form preferences over *every* act that can possibly be constructed by combining consequences with states of the world. John Broome³⁷ calls this feature the *rectangular field assumption*.

The rectangular field assumption is not *implied* by expected utility theory, from a logical point of view. But the existence of the expected utility function is derived from the rectangular field assumption and other fundamental principles.³⁸ So a violation of the rectangular field

³⁶ Reprinted in Jacques Drèze, *Essays on Economic Decisions Under Uncertainty* (New York: Cambridge, 1987), pp. 76–81.

³⁷ Broome, *Weighing Goods*, pp. 115–17; see also Sugden, "Rational Choice," pp. 761–63.

³⁸ Broome, *Weighing Goods*, p. 117.

assumption does not logically contradict expected utility theory (although we still lack an alternative representation theorem that does without that assumption). It does, however, seriously invalidate the Savage measurement procedure.

For measurement purposes it is crucial that acts can be constructed arbitrarily. Take any state of the world, and combine it with any consequence in our theoretical universe: whatever we obtain by these arbitrary combinations must make sense, and must be something the agent can form a preference about. Now suppose that the consequence *g* of player *B* in the sequential Prisoner's Dilemma game of Figure 1 is defined as "getting \$3 but being sorry for player *A*" (who had offered to cooperate). This is surely meaningful in the context of a game like the one above. But does it make sense in a different context, where that consequence does *not* follow from the act of choosing Right in the above game (after *A* has chosen Up)? We must distinguish two cases (*ibid.*, p. 116):

- (1) The (refined) consequence is *causally implausible* outside the specific game; for example, why should *B* feel sorry for *A*, if they are not playing the game of Figure 1?
- (2) The (refined) consequence is *logically impossible* outside the specific game; for example, imagine we redefined *g* as "getting \$3 after player *A* has chosen Up": clearly this is inconsistent with the game that *B* is playing in, say, Figure 2.

To sum up: when too much is included in the description of outcomes, the consequence itself remains tied to the specific game and cannot be used to construct arbitrarily other acts (or functions from states of the world to consequences).³⁹ Refining options to include reciprocity considerations clearly raises problems of both kinds. If players care about "what might have been," then this must be included in the definition of a consequence. Indeed the consequences

³⁹ See also Sugden, "Rational Choice"; Munier, "Comment"; and Bruno Verbeek, "Consequentialism, Rationality, and the Relevant Description of Outcomes," *Economics and Philosophy*, xvii (2001): 181–205. These and related problems have been addressed in the literature by introducing so-called "state-dependent" utilities. At the price of introducing some extra axioms, it has been shown that individual choices can be described as maximizing a state-dependent utility function over acts. This leads to some problems, however, in the estimation of subjective probabilities; see, for example, Mark Schervish, Teddy Seidenfeld, and Joseph Kadane, "State-Dependent Utilities," *Journal of the American Statistical Association*, LXXXV (1990): 840–47. Such problems can be overcome if one is ready to abandon the Savage measurement procedure and especially the strong operationalism associated with it; see, for example, Edi Karni and Philippe Mongin, "On the Determination of Subjective Probability by Choices," *Management Science*, XLVI (2000): 233–48.

must be refined so as to include entire branches of the game. But then one cannot use such consequences to create arbitrarily other acts, if such acts are to make sense at all. The Savage measurement procedure is not flexible enough to neutralize reciprocity counter-examples.⁴⁰

EVIDENCE OF RECIPROcity

Reciprocity has been discussed by experimental game theorists for more than two decades now. Most of this discussion however has taken place *before* the existence of reciprocal motives had been established empirically. Consider social dilemma games again. As already mentioned, many subjects seem willing to cooperate in a one-shot sequential Prisoner's Dilemma game. But this behavior of course does not constitute decisive evidence in favor of reciprocity, because the mutual cooperation outcome provides an *equal* distribution of the material payoffs (see Figure 1 above) and therefore might be preferred by players that are averse to inequality.

A crucial test for reciprocity must implement some variant of the Savage measurement procedure. Players' preferences over a set of outcomes must be observed *first* "in isolation," and *then* in the context of a game that is likely to trigger reciprocity considerations (or more generally a concern for "what might have been"). The comparison between choices in these two contexts will provide direct evidence for reciprocity as well as a proper test of game theory. Several experimenters have done precisely this in recent years, although experimental game theorists themselves do not seem to have noticed the general implications of these results.

The experiments usually start with so-called Dictator's games. A Dictator's game⁴¹ is not even a strategic game, strictly speaking, for it involves a straight choice between different allocations. The outcome is not determined by an interaction between the players: one player chooses, and the other takes whatever has been chosen. A Dictator's game then, like a stripped-down version of the Savage measurement procedure, simply elicits the (ordinal) structure of preferences of the experimental subjects. Since the ordering, rather than the cardi-

⁴⁰ Regret is another motivation that can lead to problems with the rectangular field assumption not only in the realm of game theory, but also in individual (nonstrategic) decision making. See, for example, Sugden, "Regret, Recrimination and Rationality," *Theory and Decision*, xix (1986): 77-99. Unlike reciprocity, however, we do not have strong evidence about regret at the moment.

⁴¹ See, for example, Robert Forsythe, Joel Horowitz, N.E. Savin, and Martin Sefton, "Fairness in Simple Bargaining Experiments," *Games and Economic Behavior*, vi (1994): 347-69.




Figure 3

nality, of preferences is what matters in many simple experimental games, the full-blown Savage method is unnecessary.

Once preferences have been elicited in the Dictator's game, it is then possible to check whether players behave rationally in more complex strategic situations. Take the sequential Prisoner's Dilemma. Cox (*op. cit.*) uses a slightly more complex version of this setting, called the "Investment game." Player A here chooses how much of a given endowment to pass over to B. Whatever A chooses, the sum will be tripled by the experimenter. Player B then decides how much of this investment to return to A.

After the game is played, Cox offers some subjects exactly the same sums which have been selected by players of type A in the first experiment—except that type-B players this time do not know that it is an investment from another player (the sum is presented as a "pie from the sky"). The experimenter can therefore check whether the behavior of B players in this Dictator's game is radically different from their behavior in the Investment game. As a matter of fact it is: in the Investment task the average amount returned by B players is more than twice the amount allocated to A in the Dictator's task (\$4.94 versus \$2.06).

Gary Charness and Matthew Rabin⁴² use even simpler settings (like, for example, Figure 3). The first part of the experiment (on the left) is again a Dictator's game. The first payoff between brackets is A's monetary gain and the second one is B's, as usual. The figures under each outcome represent the observed distribution of player B's

⁴² "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, cxvii (2002): 817-69.

that leads to each consequence. They care about what might have been and how they got where they are now.

CONCLUDING REMARKS

Game theory is the piece of contemporary economics that has had the most dramatic impact outside the boundaries of the discipline. Sixty years from its inception, game theory is widely used in fields as different as political science, sociology, biology, psychology, law, military strategy, and, last but not least, philosophy. Its empirical status, then, is an issue of wide scientific and philosophical relevance.

For many years a considerable degree of confusion about the status of game theory was fostered by the proliferation of seemingly anomalous but indecisive evidence. Now that experiments of the right kind have been done—experiments that really test game theory—it has become impossible to argue that the anomalies can be digested as usual. In order to see that, it is necessary to appreciate the crucial distinction between the theory of rational play and the bridge principles that connect the former with measurable properties of the real world. These two elements are sometimes conflated, sometimes treated as if they were completely independent. This probably leads to the (mistaken) belief that game theory is as flexible an explanatory tool as the folk-psychological theory of action of which it constitutes a formal refinement.

The Savage procedure (and the rectangular field assumption upon which it is based) imposes some limits on the amount of information that can be incorporated in the description of the outcomes of a game. These technical limitations make sure that the spirit of genuine consequentialism (players are forward looking, subscribe to sub-game perfect rationality, and so on) is preserved.

It is important to understand the precise implications of what was just said: nothing prevents us from applying the tools of game and utility theory across a limited range of experimental games, assuming the subjects have consistent preferences over a limited number of highly specific (or highly refined) options, once the outcomes have been (re-)described so as to include the overall structure of the game.⁴⁵ Such preference information might turn out to be useless

⁴⁵ The Savage procedure prescribes *first* to measure the utility of the outcomes of a game in isolation and *then* to see how the subjects behave in more complex strategic situations (games). The rectangular field assumption is required if one wants to elicit utility measures for a *wide* range of consequences, or in other words if one is trying to specify the *fine-grained* structure of an individual's utility function. But the application of game theory in specific instances of course does not require this full-blown measurement to be carried out in practice. In fact, as we have seen, the simple games that are customarily tested in the laboratory often require only a simplified or stripped-down version of the Savage procedure—because the mere ordinal ranking of a handful of outcomes is what really matters.




Figure 4

choices. A majority of subjects goes for the Pareto-superior outcome, but there is significant evidence of inequality aversion too.

The same (sub-)game is then transferred in a different context, where *B* is asked to make the same choice *after* player *A* has already made a move (Figure 3, on the right). As we can see, choices in this second setting differ substantially from those in the simple Dictator's game. Concerns for equality in this specific case seem to be largely swept away by *B*'s desire to respond to *A*'s kindness (*A* could have chosen unilaterally to give *B* absolutely nothing, after all).

Figure 4 presents another example of the importance of "what might have been." In the Dictator's game (left), *B* players display a considerable willingness to sacrifice a small part of their material payoff to increase *A*'s income. This willingness disappears in the enhanced game (right), where *B* moves knowing that *A* has refused the (550, 550) split (a phenomenon that Charness and Rabin dub "concern withdrawal").

There are many more experiments confirming the importance of reciprocity.⁴⁶ Much current work in experimental game theory is now focusing on the robustness of these patterns.⁴⁴ It is already possible to say, however, that in a number of settings *the structure of the game matters*. Subjects do not just care about the material consequences but also about the alternative moves that are available and the process

⁴³ See, for example, Armin Falk, Ernst Fehr, and Urs Fischbacher, "On the Nature of Fair Behavior," *Economic Inquiry*, xli (2003): 20–26; Kevin McCabe, Mary Riggdon, and Vernon Smith, "Positive Reciprocity and Intentions in Trust Games," *Journal of Economic Behavior and Organization*, li (2003): 267–75.

⁴⁴ Reciprocity seems to be a fairly erratic phenomenon. Cristina Bicchieri argues, convincingly in my view, that social norms (such as reciprocity) are highly context-sensitive and usually must be "triggered" or "cued" by some factor signalling that "this is a context where the norm applies," that is, where it is generally expected that people behave according to the norm. See, for example, Bicchieri, *The Grammar of Society* (New York: Cambridge, 2006).

outside the specific game situation, of course, but could still be used to predict and summarize choices in a limited set of games. What cannot be done, if consequences become overly sophisticated, is to have a unique utility function across the set of *all* the games that can be constructed using those (and other) consequences.

From this respect, the "naive" and often ridiculed convention of using monetary payoffs as initial conditions or proxies of players' utilities, has much to recommend. It is well known among historians and philosophers of science that the success of scientific theories is often crucially dependent on the creation of instruments of observation that allow adequate measurements of their key variables. Consider Galileo's telescope, but also the microscope for the atomic research program, the steam engine for thermodynamics, and so forth. Monetary payoffs are objective and observable, and provided we have the right theory at hand, can be used readily to generalize to other situations that are similar (in the monetary payoffs structure) to those that were tested in the laboratory. If this route is followed of course the "core" of game theory must be modified to account for the existing evidence. The partnership between the standard theory of rational play and the Savage procedure has proven to be unsustainable (given the evidence). We can keep the core of game theory and reject the bridge principles (the Savage procedure), but face difficult problems of applicability. The other viable alternative is to keep the bridge and change the theory of play.

Why did it take such a long time to design and run the right experiments to test game theory? One reason is that, surprisingly, the Savage procedure was never part of the official toolkit of the experimental game theorist. This is the result of complicated historical factors, in particular the fact that the methodology of experimental economics (especially so-called "induced value theory"⁴⁶) was developed mainly in the context of market experiments and was aimed at the active *inducement*, rather than the passive observation of preferences. But a proper discussion of these issues would require a long detour into the history and methodology of experimental game theory, which cannot be attempted here.⁴⁷

Interestingly, the refutation of game theory took place almost simultaneously with the theory's most remarkable successes in policy-

making. The reform of inefficient (usually centralized) methods of allocation and the design of new markets (for electricity, telecommunication, health services, and so on) is now increasingly informed by the principles of game theory. These principles moreover are often subjected to preliminary testing in controlled conditions before being applied to real-world situations. These tests, however, are usually aimed at spotting and avoiding precisely the sort of problematic situations that the theory is unable to deal with. Other-regarding preferences for instance can simply be ruled out by design, by making sure that the available information and the structure of the game prevent such concerns from arising in the first place. This is economic engineering, as opposed to theory-testing: the world is modified so as to "fit" the conditions that allow the theory to work. But since engineering successes often *do* drive science, we should not rush to any conclusion about the future of game theory. Its refutation and its remarkable success can (and will) coexist in the minds of many social scientists for whom the theory is, and is likely to remain, the best game in town.

FRANCESCO GUALA

University of Exeter