

# Against Game Theory

GALE M. LUCAS, MATHEW D. McCUBBINS, and MARK TURNER

## Abstract

People make choices. Often, the outcome depends on choices other people make. What mental steps do people go through when making such choices? Game theory, the most influential model of choice in economics and the social sciences, offers an answer, one based on games of strategy such as chess and checkers: the chooser considers the choices that others will make and makes a choice that will lead to a better outcome for the chooser, given all those choices by other people. It is universally established in the social sciences that classical game theory (even when heavily modified) is bad at predicting behavior. But instead of abandoning classical game theory, those in the social sciences have mounted a rescue operation under the name of "behavioral game theory." Its main tool is to propose systematic deviations from the predictions of game theory, deviations that arise from character type, for example. Other deviations purportedly come from cognitive overload or limitations. The fundamental idea of behavioral game theory is that, if we know the deviations, then we can correct our predictions accordingly, and so get it right. There are two problems with this rescue operation, each of them is fatal. (i) For a chooser, contemplating the range of possible deviations, as there are many dozens, actually makes it exponentially harder to figure out a path to an outcome. This makes the theoretical models useless for modeling human thought or human behavior in general. (ii) Modeling deviations are helpful only if the deviations are *consistent*, so that scientists (and indeed decision makers) can make predictions about future choices on the basis of past choices. But the deviations are not consistent. In general, deviations from classical models are not consistent for any individual from one task to the next or between individuals for the same task. In addition, people's beliefs are in general not consistent with their choices. Accordingly, all hope is hollow that we can construct a general behavioral game theory. What can replace it? We survey some of the emerging candidates.

## INTRODUCTION

Scholars employ game theory to model interdependent decision-making in bargaining, constitutional law, democratic stability, standard setting, gender roles, social movements, communication, markets, voting, coalition formation, resource allocation, war, and many other domains. For a review, with citations, see Lucas, McCubbins, and Turner (2013).

Game theory has been widely discredited: many studies have demonstrated game theory's mispredictions. People seem to be highly sensitive to frame or context and biased in their strategies. They follow heuristic decision-making. They are limited in their ability to reason and learn. Many attempts have been made to create a *behavioral* game theory by adding a correction for each of four-dozen odd mispredictions (see Lucas *et al.*, 2013). But the span of mispredictions is so great and so varied that building in corrections for them produces a model that is computationally intractable for actual human beings.

### REVIEW: THE ELEMENTS OF GAME THEORY

What is a game? A game, theoretically, is defined by identifying the players, the actions available to them, the information they have about the game, when they have such information, what they know about what others know or will know and when those others will know it, the strategies available to them that define rules for what actions they will take in making decisions, the payoffs that come with outcomes, the range of outcomes, and "equilibria." The acronym for these dimensions is PAISPOE: players, actions, information, strategies, payoffs, outcomes, and equilibria. What is an equilibrium? An equilibrium is a path of choices made by players in the game—a path that could happen, given the dispositions of all the players as defined in game theory. An equilibrium concept is a rule a player uses to pursue an equilibrium path. There are a number of proposed equilibrium concepts that players might use. They have names such as "dominant strategy," "Nash equilibrium," "Bayesian strategy," "correlated strategy," and "subgame perfect Nash equilibrium (SPNE)." The classic game theory equilibrium concepts are "dominant strategy" and "Nash equilibrium." A pure-strategy Nash equilibrium concept, for example, is a combination of actions for all the players according to which no player can benefit by unilaterally deviating from his or her combination. For a review of types of equilibrium concepts, see Maschler, Solan, and Zamir (2013). To be generalizable, all of the attempts to model interdependent choice, classical or behavioral, must assume that (i) people follow equilibrium strategies, (ii) there are specific types of people who choose a generalizable strategy over a class of tasks, or (iii) all people in performing a specific task choose a generalizable behavioral strategy. Behavioral game theory is trying to build a layer cake on top of game theory, by adding layers to the original model. Each additional layer consists of a correction to the foundation. For example, prospect theory would eliminate certain strategies associated with disfavored bets.

But the number of adjustments needed to build a behavioral game theory is so vast that it cannot yield generalizable models. The literature has proposed many adjustments in terms of bounds, biases, heuristics, and context dependencies. This presents two problems for behavioral game theory. First, experimental economists working in the laboratory and knowing that subjects make one or more of these adjustments are no longer in the position of knowing subjects' true payoffs: we know only the experimental economists' view of the experimental earnings, typically thought to be the subjects' earnings. What experimentalists do not consider is "context," that is, factors such as unobserved experimenter demand and framing. Second, and most important, under such adjustments, we do not know what the subjects believe. Most games, especially those simple enough to be tested in the lab and simple enough that we can strip out most (if not all) of the effect of framing and context, assume that subjects share *common knowledge*. Classical game theory requires players to have correct and *consistent* beliefs. To have "correct beliefs" is to regard other players as following classical game theory and to predict that they follow classical equilibrium strategies. Indeed, it is also required that players know that other players know that they themselves are following classical strategies and so on, ad infinitum. As Lupia, Levine, and Zharinova (2010) note, the condition is even stronger, in that a classical equilibrium concept "requires shared conjectures ... . Common Nash refinements ... continue to require that actors share identical conjectures of other players' strategies" (p. 106). This is part of what economists assume when they accept that the players in a game share "common knowledge." As Smith (2000, p. 9) writes, citing two other winners of the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel:

"The common knowledge assumption underlies all of game theory and much of economic theory. Whatever be the model under discussion ... the model itself must be assumed common knowledge; otherwise the model is insufficiently specified and the analysis incoherent" (Aumann 1987, p. 473). Without such common knowledge people would fail to reason their way to the solution arrived at cognitively by the theorist. This is echoed by Arrow (1987) when he notes that a "monopolist, even ... where there is just one in the entire economy, has to understand all these [general equilibrium] repercussions ... has to have a full general equilibrium model of the economy." (p. 207)<sup>1</sup>

Nash showed that any two-player zero-sum game has an equilibrium and it was later proved that finding this equilibrium is computationally tractable. (A zero-sum game has fixed payoffs, where higher payoffs to one player result in corresponding lower payments to other players). This was the

---

1. Smith here quotes a reprint of an original article by Arrow (1986).

appeal of noncooperative game theory: we could find equilibria that in turn predict outcomes of interest. But, even in the case where there are only two players and two actions, we cannot expect humans to solve non-zero-sum games, as they are computationally intractable or infeasible, falling into the class of computational problems identified as PPAD-complete<sup>2</sup> (for history, definition, and analysis, see Daskalakis, Goldberg, & Papadimitriou, 2009; see also Daskalakis, Goldberg, & Papadimitriou, 2005; Chen & Deng, 2006).

The core problem for building a behavioral game theory is that, as we add the biases, heuristics, and context dependencies suggested by decades of research, we are implicitly increasing the dimensionality of the computational problem of finding an equilibrium. It may seem that building a range of possible deviations into the model would help us build a better model. Doing so requires the bold assumption that a given subject, when faced, for example, with a series of non-zero sum games, deviates from classical game theory in a way that is *consistent* from one game to the next, and even from one choice to the next inside the same game. The result is indeed a generalizable theory—perhaps a false theory, but one that is at least generalizable. But for that theory to give us purchase on modeling human thought or predicting human behavior, the deviations in beliefs and therefore strategies, must—at a minimum—be consistent. They also must be common knowledge: subjects will have no way to compute consistent strategies if the way in which subjects deviate from strategies is private information belonging to only the capricious subject and unknowable to other subjects. That is the bedrock on which the rescue operation for failed game theory must be built. Our evidence, however, like all other evidence of which we are aware that touches on this point, shows that this presumed bedrock does not in fact exist.

## NEW FINDINGS

### RESEARCH DESIGN

There are two important new features of our experimental design. First, we created a battery of tasks. Our experiments use a battery of up to 17

---

2. PPAD stands for “polynomial parity argument on directed graphs.” It is a complexity class regarded as exceptionally difficult. In computational complexity theory, a problem  $X$  is called *hard* for a complexity class  $C$  if any problem in  $C$  can be reduced to  $X$ , which implies that no problem in  $C$  is harder than  $X$ , as a solution to  $X$  provides a solution to any problem in  $C$ . If  $X$  is both hard for  $C$  and in  $C$ , then  $X$  is called *C-complete*. A problem that is  $C$ -complete is the hardest problem, computationally, in  $C$ , or rather, there is no harder problem in  $C$ . A PPAD-complete problem is in principle “so hard to calculate that all the computers in the world couldn’t find it in the lifetime of the universe” (Hardesty, 2009). Accordingly, it is difficult to imagine that human beings playing a game would seek solutions by trying to perform such a computation. “By showing that some common game-theoretical problems are so hard that they’d take the lifetime of the universe to solve, Daskalakis is suggesting that they can’t accurately represent what happens in the real world” (Hardesty, 2009).

games, several of which we constructed by modifying the standard form of well-known games such as “Prisoner’s Dilemma,” “Public Goods,” “Stag Hunt,” “Ultimatum,” “Trust,” “Chicken,” and “Dictator.” The purpose of our modifications is to minimize or eliminate the framing of the game and to present the games, to the extent possible, as starkly as they are defined in prominent textbooks in game theory. For details, see Lucas *et al.* (2013). We emphasize that, unlike the typical method of running experiments in psychology and economics, where subjects face the same task repeatedly, our method presents subjects with a battery of tasks. In addition, unlike psychology experiments, where subjects are typically paid in the form of satisfying a course requirement, and unlike economics experiments, where subjects are typically paid at random for only one of the two to three dozen repetitions of the task, our subjects know that they are paid in cash according to *every* action they take in *every* task.

Second, we added prediction markets, based on Plott and Roust (2005) and Wolfers and Zitzewitz (2004). In these prediction markets, subjects could earn additional money by placing bets on the choices that were made by the players with whom they were matched and in many cases by placing bets on the collective choices of all subjects who had played the game during previous runs of the experiment. We create a market, specifically a betting market, where we invite subjects to bet on other players’ choices. We quiz the subjects on the betting procedure so that we could both motivate them to work hard to understand what the bets entailed and also to measure each subject’s compliance with respect to the betting tasks. The subjects were paid if their bets regarding the other player’s actions were accurate and they were given a chance to also double down on their bets if they felt confident about their predictions. Subjects understood that they can earn this extra money from betting, and how much for each bet. For example, in the Trust Game, there are two players randomly paired—Player 1 and Player 2; both start with \$5; Player 1 can select any number of dollars from 0 to 5 to give away; those dollars are taken from Player 1, tripled by the experimenters, and given to Player 2; after receiving them, Player 2 has the opportunity to transfer back any number of the dollars that Player 2 has in total, including Player 2’s original \$5. At this point, Player 2 may have any number of dollars between 5 and 20. We ask Player 1 to guess how many dollars Player 2 will return. Later, but before Player 2 learns Player 1’s choice, we ask Player 2 to guess how many dollars Player 1 selected to give away. We also ask Player 2 to guess how much Player 1 predicted Player 2 would transfer back to Player 1. After Player 2 learns Player 1’s choice, we then ask Player 2 to guess how much Player 1 predicted that Player 2 would return. All players know that a player earns \$3 for each correct guess and nothing for a guess that is wrong. The questions we ask vary slightly for each task, but, as an example, here

is an exact question we ask Player 1 in the Trust Game: “How much money do you guess the other person transferred back to you? If you guess correctly, you will earn \$3. If not, you will neither earn nor lose money.” Players know that, with one exception, they can *never* learn whether their predictions (bets) were right or wrong and that subjects *never* have any information about other subjects’ guesses (bets). The exception is the rare case in which Player 2 in a sequential game, such as Trust or Ultimatum, must know Player 1’s choice in order for Player 2 to understand Player 2’s situation and payoffs. For example, Player 2 in Trust must be informed of how many dollars Player 1 chose to give away, if Player 2 is going to know how many dollars Player 2 has and therefore what the possible actions and payoffs are. Even then, the delivery of this information to Player 2 and Player 2’s subsequent choice are postponed to the last set of choices a Player makes, so as to have no effect on previous choices. Payments to subjects are made in a lump sum, without accounting or explanation, individually, anonymously, and privately, when the experiment is completed.<sup>3</sup> We do this to eliminate any opportunity for subjects to make inferences about other players’ choices (either individually or collectively), and subjects know this before they make any of their many choices in the extensive battery. For details on these prediction markets in a range of games, see Lucas *et al.* (2013).

## RESULTS

ACROSS our battery of tasks, our results verify decades of research demonstrating that subjects do not follow game-theoretic predictions. Consider a traditional Ultimatum Game. According to Andreoni and Blanchard (2006, p. 307), this game “has come to symbolize the power” of classical game theory and “its utter failure in practice.” In this “bargaining game,” a “proposer” makes a take-it-or-leave-it offer to a “responder,” who then accepts or rejects the offer. The classical equilibrium prediction is that if players care only about their own monetary payoffs, then the responder will accept any positive offer the “proposer” makes. (More technically, not counting the “endowment” of dollars with which players begin—in our experiment, for example, the proposer begins with \$10 and the responder with \$0—if we assume that players care only about the money they earn, then the Nash Equilibrium prediction for the responder is that the responder will accept any positive offer the “proposer” makes; and the Nash equilibrium prediction for the proposer is that, knowing that the responder will accept any positive offer, the proposer will reason by backward induction to choose the smallest possible positive offer allowed by the game.) Despite the stark framing of our experimental tasks,

---

3. Except for on time show up fees that were paid to all subjects and except for payments for the first quiz relating to general experimental instructions.

our results generally replicate what others have found about the poverty of classical Nash equilibrium predictions (see Camerer, 2003). For example, only slightly more than 6% of our subjects chose the classical strategy when they play the Ultimatum Game.

But the question we focus on with our battery of experimental games is, when players “deviate” from the classical strategies, do they do so consistently? The answer is no, and if there is no consistency, then the modifications to classical game theory cannot provide a generalizable model of behavior. For example, the first half of our Trust Game is exactly like a game that we call “Donation”: a given subject as Player 1 in Trust is in exactly the same payoff and action situation as he or she is when he or she is the Donor in Donation. Accordingly, we can measure whether a given subject is consistent across these two situations. In addition, the second half of our Trust Game is exactly like our Dictator Game: a given subject as Player 2 in our Trust Game faces the same incentives and action possibilities as he or she faces when he or she is the Dictator in our Dictator Game.<sup>4</sup> We thus can measure whether a given subject is consistent across these two situations. In addition, all subjects completed a Trust Game where they made choices as Player 1 and also completed a Trust Game where they made choices as Player 2. We therefore can examine the consistency of a given subject’s behavior *within a game*, namely Trust (see Lucas *et al.*, 2013, for details).

Our findings show that in Trust, there is large variance in behavior across subjects in the role of Player 1 and also in the role of Player 2. (We discuss findings in more detail in the Appendix.)

They also show that there is large variance in behavior by the same subject in those two roles. We expected to find that subjects deviate from predicted behavior for each and every task, but, for example, our results show much more than that for Trust, Dictator, and Donation: fewer than 15% of our subjects deviate *consistently* from classical equilibrium concepts across the four tasks involved in those games.<sup>5</sup> For these four tasks, by even the most minimal definition of consistency, only 42% of our subjects either consistently follow classical equilibrium concepts *or* consistently deviate from them. That is, a minority of subjects have even the most minimal consistency from task to task. We report on similar results in Lucas *et al.* (2013).

This suggests that simple amendments to game theory, such as adding social preferences of one sort or another, risk preference, or discovering each subject’s individual level of experience or ability to undertake the reasoning necessary to choose an optimal strategy, or the level of randomness in each

4. Of course, one difference remains, the Trust task is interactive, whereas the Dictator task is not.

5. That is, in making choices as Player 1 or Player 2 in Trust, in making choices in Dictator or Ultimatum, subjects sometimes deviate and sometimes do not deviate from classical Nash equilibrium predictions. Those that deviate do not always do so, and few subjects deviate on every choice and few never deviate.

individual's choices, cannot by themselves explain the inconsistent pattern of choices across the battery of games in our experiment. For example, sometimes the subject may "appear" altruistic, in one choice, and then "appear" greedy in another choice, even though the incentives are identical and are stated identically for the two tasks. Some subjects are altruistic some of the time, other subjects are altruistic most of the time, and some subjects are never altruistic. At best, we need amendments to game theory that differ for each and every game and for each individual. This would make it difficult, or impossible, to provide a general explanation of behavior that is built on game theory.

Game theoretic predictions about behavior depend on the assumption that beliefs and choices are aligned. Researchers rarely possess knowledge of the actual beliefs of subjects. But our within-subject experiments allow us to test beliefs, and we find that for individual subjects, there is routine and ubiquitous inconsistency across choice and beliefs (see Appendix and McCubbins, Turner, & Weller, 2012). We demonstrate that subjects' beliefs are often inconsistent with equilibrium predictions, which has not been widely appreciated. Our findings also show that these deviations are not consistent; they depend on the specific setting and task. These deviations are so pervasive and so various even within a single subject that it seems unwarranted to refer to them as *deviations*. On the contrary, consistent "Nash behavior and beliefs" appear to be remarkable deviations from human cognitive patterns and human behavior.

#### ADVANCE DIRECTIVE FOR BEHAVIORAL GAME THEORY: DO NOT RESUSCITATE

Our results show, as has often been shown, that subjects deviate from the predictions made by classical equilibrium analysis in game theory. We emphasize that (i) subjects often deviate from these predictions, but that (ii) for the vast majority of subjects, their deviations are themselves not consistent even across similar tasks, and (iii) there is large variance in how different subjects choose for any individual task. Moreover, we show (iv) that individuals' beliefs about other subject's choices or beliefs do not support classical Nash equilibrium strategies, and (v) that there is large variance between subjects and among a single subject's beliefs from task to task. In addition, we show that individuals do not hold common beliefs about game strategy or deviations from equilibrium. Individuals' beliefs seem to be specific to particular settings and not generalizable from one setting to the next. Indeed, it may be misleading to refer to these patterns of action and belief as "deviations" at all. There are no consistent deviations from classical equilibrium concepts, and thus there is no general behavioral fix. There are about four



dozen deviations from classical game theoretic predictions identified in the literature. We find that the same individual subject will be deviating from game theoretic predictions in as many ways as there are tasks in our experiment and that across all subjects in our experiment we see a great variability from one subject to the next in the pattern of deviations. It is unsurprising, for example, that some subjects are altruistic in some settings. It is unsurprising that we find that some fraction of the subjects are altruistic for one or another task. It is more surprising that subjects are altruistic in one task and then not altruistic when offered the identical incentives again later on. All these different variations of course interact with each other, giving a complex and unpredictable landscape of complex variation running over individuals and groups. Thus, adding behavioral assumptions to the general model of cognition within game theory cannot make these general models more suitable for predicting behavior.

We have shown further that the protected core of game theory—the unrecognized cognitive model, or Theory of Mind (McCubbins *et al.*, 2012), of noncooperative game theory—fails repeatedly in hypothesis testing. The assumptions about human cognition that are part of game theory, including the predictions of classical game theory and its refinements, are at odds with what we know about actual human cognition. This is no surprise, because the equilibrium concepts were not constructed based on how actual humans think, reason, or make decisions. We do not yet see a way forward to creating a behavioral game theory that offers meaningfully generalizable predictions. Accordingly, our advance directive would say: Do Not Resuscitate.

## THE NEXT STEP

What are the alternatives to game theory, both classical and behavioral? What possibilities are there for forming testable hypotheses about interdependent decision-making?<sup>6</sup> Within cognitive science, there are a number of lines of research about how actual human beings make actual decisions, several of them reviewed by Turner (2001). These lines of research in cognitive science have had virtually no consideration inside economics. Here are a few of them:

*Variation Across Domains and Situations.* Entire subfields of cognitive science are dedicated to the ways in which human thought varies across different domains and situations. Given basic considerations of evolutionary development and fitness, there is no reason to assume a priori that the way a human

---

6. Another way in which behavioral game theory has sought to account for ubiquitous deviations from game-theoretic predictions is to add a random factor to human decision-making. One such prominent approach is “quantal response equilibrium” (see McKelvey & Palfrey, 1995). But it is largely impossible to put quantal response equilibrium to the test as almost any pattern of behavior is consistent with its predictions (see McCubbins *et al.*, 2013), especially when it is expanded to “heterogenous quantal response equilibrium” (Rogers, Palfrey, & Camerer, 2009).

being thinks and acts with respect to food, mating, entertainment, and so on would follow the same patterns or principles of reasoning and decision. All models of game theory, classical and behavioral, assume the opposite, namely, that although people might have different preferences with respect to these different domains, their patterns of reasoning and choice must be uniform across them. Game theory models two situations as identical if they have the same game structure, regardless of the content. It does not matter, for example, that the “Stag Hunt” game might concern growing vegetables or killing soldiers. To a cognitive scientist, or maybe just to anybody other than a game theorist, that approach looks like a nonstarter. It certainly should be discarded as an assumption.

*Learning.* Entire subfields of cognitive science are dedicated to the various mental operations involved in learning, and to the study of how human thought and action depend on the highly flexible and powerful learning for which human beings are equipped. To give one example, cognitive science routinely considers the power of analogy or blending: one remembers a previous specific situation, perhaps in childhood, or from a biography, or even from a science fiction novel, and remembers, too, its outcomes; one then uses those concepts and knowledge to inform one’s understanding and decisions in the present specific situation. These traditions in cognitive science have no status inside game theory, even though they provide a basis for hypotheses and tests about decision-making. Indeed, the validity of measuring such powers of analogy and blending is so unquestioned in Western culture that assessments of this ability are used as part of the process of deciding who has what IQ and which applicants to college should be admitted.

*Complexity and Nonlinearity.* The world is rich, and in the typical situation, actors are engaged in simultaneous games that overlap. In life, any action is usually a move in many different games. Strategies to maximize expected utility over all these games are typically nonlinear. In principle, the output of any subgame of any game can be input to any subgame of any other game. Game theory by contrast assumes a partitioning of thought and action to tiny scripts of activity that are pretty much separate from all others. This assumption could be discarded.

*Adaptive Behavior.* In the typical situation, people are adaptive: their first and strongest disposition is often not to play the game but to reinvent it, change it. Their decisions can be driven by attempts to change the game from the outside. Game theory leaves no room for this normal and routine thought and behavior.

*Construal.* Cognitive science routinely investigates our rich capacities for differing *construals* of the same given material. “The mountain range runs from Canada to Mexico” and “The mountain range runs from Mexico to Canada” deal with the same stuff but call for quite different emphases

and viewpoints. They also both call for conceptualization by using the idea of *motion* (“runs from ... to”), even though in some sense we think that no motion is involved. Construal is a crucial part of interdependent decision-making, because actors try to reconstrue history and to get other actors to do the same. In the typical situation, actors work at conceptual reinterpretation of the history of play, so as to persuade other actors that the value and status of a past action must be changed, and further, to persuade them that the action led to nodes different from those to which it was once thought to lead. Conceptually, the history of the game is not fixed. Game theory assumes the opposite.

*What's Up?* Actors must operate in general without knowing what game they are in, and the question always arises, who has the authority to recognize and establish the game being played? Actors attempt to influence other actors' thoughts about the game being played.

*Identity.* Cognitive science routinely considers the work people do to construct an identity for themselves, and to carry it and vary it appropriately from situation to situation. It may often be that the principle payoff in any scripted activity is not the local payoffs but the actor's concept of a personal identity. When we enter different situations, different rooms, different moments, what does the present offer by way of allowing us to construct an identity? What looks like fatal inconsistency from the point of view of game theory might look like fruitful experimentation, learning, and fluidity from the point of view of actual human beings.

Within the social sciences, it has been shown that institutions (laws, constitutions, auction mechanisms, common agency, families, friendships, societal structures, and so on) serve to create not only incentives for choice but also a set of shared mental models about players, actions, payoffs, outcomes, and perhaps most important, information. This would imply that the study of institutions might supply some of the cognitive grounding that game theory is currently lacking. Yet another assumption of game theory is that from knowledge about players, actions, payoffs, outcomes, and information, one can derive strategies and equilibria. It is unknown whether this is true, but what is clear is that institutions have already influenced how subjects derive strategies and equilibria. People, developed within institutions, are thus, when they enter our experiments, far from a *tabula rasa*. There is little reason to imagine that the narratives that we can give them in an experiment are strong enough to offer much hope of overcoming that training within institutions. There is little reason to think that these narratives in experiments would substitute for the purported “10,000 h” of learning and practice needed to be successful at a given task. Accordingly, the cognitive study of decision-making must include the study of learning within institutions.

In short, we think that the future for game theory, if there is one, would come from a grounding in cognitive science, or more generally, in the analysis of how the cognitively modern human mind works, what its basic mental operations are, and how they are deployed in situations. We know that strategic games such as chess arose very late in human evolution and even in human culture and that people are very poor at such games in general and must undergo extensive training in order to play them well. Such games of strategy are perhaps the last place one should look for a model of human decision-making generally. Game theory has placed itself into that cul-de-sac, but there is no reason for that sterility and isolation. One could instead begin with how actual human beings think.

#### APPENDIX ON EXPERIMENTAL DATA

This appendix provides details of subjects' behavior in Trust, Dictator, and Donation. First, on examining the play within the Trust Game, we find that 56% of the subjects as Player 1 sent money to Player 2. On average, they sent \$1.43, with a standard deviation of \$1.70. On average, as Player 2, they return \$1.23, with a standard deviation of \$2.29. Our emphasis is not on the well-established deviance from classical predictions in Trust (indeed, the standard deviation from our experiments is somewhat smaller than that usually reported) but rather on the large variance in behavior both across subjects in both tasks and by the same subject across different tasks. Only 1 of the 80 subjects who as Player 2 received \$0 from Player 1 returned any money to Player 1. Of the 100 subjects who did receive money as Player 2, 62 returned something. The average returned for this subset is \$2.22, again with a large variance (the standard deviation is \$2.71). Second, for the 62 who as Player 2 returned money to Player 1 after receiving money, only 40 sent money when they were the Dictator; and of those 40, only 29 sent money when they were in the role of Donor. Was a subject's pattern of deviation from classical equilibrium consistent? There are 42 subjects who deviate from classical equilibrium as *both* Player 1 and Player 2 in Trust. Of these 42, 33 also deviated as Donor and, of these 33, 26 deviated as Dictator. We see that fewer than 15% of our subjects *consistently* deviate from classical equilibrium concepts across these four tasks. In sum, by even the most minimal definition of consistency, only 42% of our subjects either consistently follow classical equilibrium concepts (specifically, "SPNE") *or* consistently deviate from them, in these four tasks. That is, a minority of subjects have even the most minimal consistency from task to task. We report on similar results in the study by Lucas *et al.* (2013).

We turn now to reporting on the often hidden and never-tested parts of "equilibrium concepts" in game theory, that is, the assumptions regarding

subjects' beliefs and knowledge. In the Trust Game, for example, the classical SPNE prediction for both players is that they will send \$0 for all tasks. Thus, if our subjects hold beliefs that support SPNE, both Player 1 and Player 2 should expect the other person in each task to send \$0. Further, they should expect that the other player expects that they will send \$0. (And so on ad infinitum, they should expect that the other expects that they expect that the other player will send \$0, and so on.) When acting as Trust Player 1, however, 88 of 180 subjects bet, and thus can be thought to believe, that Player 2 will return some money to them. Likewise, when in the role of Player 2 in Trust, the majority of participants (112 of 180) believe that Player 1 will send them *more* than the equilibrium amount of \$0. Indeed, only 21% (38 of 180) hold the "correct" SPNE beliefs and, as both Player 1 and Player 2, bet that the other person will send nothing. (More elaborately, as the "other person" is always that other subject with whom the subject has been randomly paired for that specific task, only 21% of the players expect both when they are Player 1 and later when they are Player 2 that the other subject with whom they have been randomly paired for that particular game will send nothing.) In total, 58 of 180 participants bet when they are in both roles—that is, Player 1 and Player 2—that the other person will send more than \$0. In general, participants do not hold beliefs that support a SPNE.

Examining consistency of belief in more depth, we can ask, for example, how many of the 180 subjects in our analysis consistently held beliefs that support an SPNE? As Player 2, participants made guesses about Player 1's prediction of how much Player 2 would return, and only 92 of 180 made guesses that support an SPNE. Of those 92, only 41 also held SPNE-consistent beliefs as Player 2 when guessing how much Player 1 predicted that Player 2 guessed Player 1 would transfer. Of those 41, 33 were also SPNE-consistent as Player 2 when guessing how much Player 1 would transfer. Of those 33, 29 were also SPNE-consistent as Player 1 when guessing how much Player 2 predicted that Player 1 guessed that Player 2 would return. Of those 29, 27 were also SPNE-consistent as Player 1 when guessing both (i) how much Player 2 would return and (ii) how much Player 2 predicted that Player 1 would send. In sum, only 15% of our subjects *consistently* adhere to beliefs that support an SPNE as Player 1 and as Player 2 in a single game, Trust.

## REFERENCES

- Andreoni, J., & Blanchard, E. (2006). Testing subgame perfection apart from fairness in ultimatum games. *Experimental Economics*, 9, 307–321.
- Arrow, K. J. (1986). Rationality of self and others in an economic system. *The Journal of Business*, 59(4), Part 2: The Behavioral Foundations of Economic Theory. pp. S385–S399.

- Aumann, R. (1987). Game theory. In *New palgrave dictionary of economics* (Vol. 2, p. 473). New York, NY: Macmillan.
- Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- Chen, X., and Deng, X. (2006). Settling the complexity of 2-player Nash-equilibrium. *Proceedings of the 47th Annual IEEE symposium on foundations of computer science* (FOCS'06; pp. 261–272). doi:10.1109/FOCS.2006.69
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2005). The complexity of computing a Nash equilibrium (TR05–115). *Electronic Colloquium on Computational Complexity*. Retrieved from <http://eccc.hpi-web.de/eccc-reports/2005/TR05-115>
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *Communications of the ACM*, 52(2), 89–97. doi:10.1145/1461928.1461951
- Hardesty, L. (2009, November 9). What computer science can teach economics. *MIT News*. Retrieved from <http://web.mit.edu/newsoffice/2009/game-theory.html>
- Lucas, G., McCubbins, M., & Turner, M. (2013). Can we build behavioral game theory? Available at SSRN: <http://ssrn.com/abstract=2278029> or 10.2139/ssrn.2278029
- Lupia, A., Levine, A. S., & Zharinova, N. (2010). Should political scientists use the self confirming equilibrium concept? Benefits, costs and an application to the Jury Theorem. *Political Analysis*, 18, 103–123.
- Maschler, M., Solan, E., & Zamir, S. (2013). *Game theory*. Cambridge, England: Cambridge University Press.
- McCubbins, M. D., Turner, M., & Weller, N. (2012). The theory of minds within the theory of games. In H. R. Arabnia, D. de la Fuente, E. G. Kozerenko, P. M. Lamonica, R. A. Liuzzi, J. A. Olivas, A. M. G. Solo & T. Waskiewica (Eds.), *Proceedings of the 2012 international conference on artificial intelligence* (Vol. I, pp. 515–521). Las Vegas, NV: CSREA Press.
- McCubbins, M. D., Turner, M., & Weller, N. (2013). Testing the foundations of quantal response equilibrium. In A. M. Greenberg, W. G. Kennedy & N. D. Bos (Eds.), *Social computing, behavioral-cultural modeling and prediction. Lecture notes in computer science* (Vol. 7812, pp. 144–153). Berlin: Springer.
- McKelvey, R., & Palfrey, T. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10, 6–38.
- Plott, C. & Roust, K. (2005). The design and testing of information aggregation mechanisms: a two-stage parimutuel IAM. Social Science Working Paper # 1245. Pasadena, California Institute of Technology. Retrieved from <http://www.hss.caltech.edu/SSPapers/sswp1245.pdf>
- Rogers, B. W., Palfrey, T. R., & Camerer, C. F. (2009). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144(4), 1440–1467.
- Smith, V. (2000). “Rational choice: The contrast between economics and psychology” in *Bargaining and market behavior: Essays in experimental economics* (pp. 7–24). New York, NY: Cambridge University Press.

- Turner, M. (2001). "Choice," chapter 3 of *Cognitive dimensions of social science*. New York, NY: Oxford University Press.
- Wolfers, J., & Zitzewitz, E. (2004). *Prediction markets*. *Journal of Economic Perspectives*.

#### GALE M. LUCAS SHORT BIOGRAPHY

**Gale M. Lucas** is currently a postdoctoral research associate at University of Southern California's Institute for Creative Technologies. She is also affiliated with USC's Marshall School of Business. Dr. Lucas earned her PhD in Social Psychology from Northwestern University, where she was named a National Science Foundation Graduate Research Fellow. She has taught at Northwestern University, Willamette University, and Western Oregon University.

#### MATHEW D. McCUBBINS SHORT BIOGRAPHY

**Mathew D. McCubbins** is Professor of Political Science and Law at Duke University, and Director of the Center for Democracy and the Rule of Law, Duke School of Law. Web page with link to CV: [www.mccubbins.us](http://www.mccubbins.us)

#### MARK TURNER SHORT BIOGRAPHY

**Mark Turner** is Institute Professor and Professor of Cognitive Science at Case Western Reserve University. He is the Founding Director of the Cognitive Science Network; Co-Director of the Red Hen Lab; Founding President of the Myrfield Institute for Cognition and the Arts; Fellow of the Institute for Advanced Study, the Center for Advanced Study in the Behavioral Sciences, the National Humanities Center, the John Simon Guggenheim Memorial Foundation, the Institute of Advanced Study at Durham University, the Centre for Advanced Study at the Norwegian Academy of Science and Letters, the New England Institute for Cognitive Science and Evolutionary Psychology, the National Endowment for the Humanities, and the Institute for the Science of Origins; Extraordinary Member of the Humanwissenschaftliches Zentrum der Ludwig-Maximilians-Universität; and External Research Professor of the Krasnow Institute for Advanced Study. His research is presented at <http://markturner.org>

#### RELATED ESSAYS

Models of Revealed Preference (*Economics*), Abi Adams and Ian Crawford  
Deterrence (*Sociology*), Robert Apel and Daniel S. Nagin  
History and Epistemology of Anthropology (*Anthropology*), Arjun Appadurai

Agency as an Explanatory Key: Theoretical Issues (*Sociology*), Richard Biernacki and Tad Skotnicki

Choice Architecture (*Psychology*), Adrian R. Camilleri and Rick P. Larrick

Heuristic Decision Making (*Political Science*), Edward G. Carmines and Nicholas J. D'Amico

Neoliberalism (*Sociology*), Miguel A. Centeno and Joseph N. Cohen

Culture and Cognition (*Sociology*), Karen A. Cerulo

The Inherence Heuristic: Generating Everyday Explanations (*Psychology*), Andrei Cimpian

Behavioral Economics (*Sociology*), Guy Hochman and Dan Ariely

Emotion and Decision Making (*Psychology*), Jeff R. Huntsinger and Cara Ray

Genetic Foundations of Attitude Formation (*Political Science*), Christian Kandler *et al.*

Search and Learning in Markets (*Economics*), Philipp Kircher

Complexity: An Emerging Trend in Social Sciences (*Anthropology*), J. Stephen Lansing

Empirical Models of Bilateral Contracting (*Economics*), Robin S. Lee

From Individual Rationality to Socially Embedded Self-Regulation (*Sociology*), Siegwart Lindenber

Understanding Risk-Taking Behavior: Insights from Evolutionary Psychology (*Psychology*), Karin Machluf and David F. Bjorklund

Heuristics: Tools for an Uncertain World (*Psychology*), Hansjörg Neth and Gerd Gigerenzer

Event Processing as an Executive Enterprise (*Psychology*), Robbie A. Ross and Dare A. Baldwin