Testing the Foundations of Quantal Response Equilibrium

Mathew D. McCubbins[1], Mark Turner[2], and Nicholas Weller[3]

[1] University of Southern California, Marshall School of Business, Gould School of Law and Department of Political Science
[2] Case Western Reserve University, Department of Cognitive Science
[3] University of Southern California, Department of Political Science

Abstract. Quantal response equilibrium (QRE) has become a popular alternative to the standard Nash equilibrium concept in game theoretic applications. It is well known that human subjects do not regularly choose Nash equilibrium strategies. It has been hypothesized that subjects are limited by strategic uncertainty or that subjects have broader social preferences over the outcome of games. These two factors, among others, make subjects boundedly-rational. QRE, in essence, adds a logistic error function to the strict, knife-edge predictions of Nash equilibria. What makes QRE appealing, however, also makes it very difficult to test, because almost any observed behavior may be consistent with different parameterizations of the error function. We present the first steps of a research program designed to strip away the underlying causes of the strategic errors thought to be modeled by QRE. If these causes of strategic error are correct explanations for the deviations, then their removal should enable subjects to choose Nash equilibrium strategies. We find, however, that subjects continue to deviate from predictions even when the reasons presumed by QRE are removed. Moreover, the deviations are different for each and every game, and thus QRE would require the same subjects to have different error parameterizations. While we need more expansive testing of the various causes of strategic error, in our judgment, therefore, QRE is not useful at *predicting* human behavior, and is of limited use in *explaining* human behavior across even a small range of similar decisions.

Keywords: bounded rationality, human behavior, Nash equilibrium, behavioral game theory, strategic uncertainty, social preferences, Quantal Response Equilibrium

1 Game Theoretic Models of Behavior

A common approach to predicting strategic human behavior across many different situations utilizes the theory of non-cooperative games. At the same time, decades worth of human subject experiments now demonstrate that the basic predictions of game theory are inaccurate [1, 2, 3]. As a result, scholars have developed a variety of modifications of the standard Nash equilibrium concept that relax the strict behavioral and cognitive assumptions contained in Nash equilibrium.

One of the most well-known models that relaxes the predictions of Nash equilibrium is the Quantal Response Equilibrium (QRE) [4]. Some argue that QRE "almost always explains the direction of deviations from Nash and should replace Nash as the static benchmark to which other models are routinely compared." [5] The QRE model maintains the assumption that individuals have beliefs that are supported in equilibrium by the strategies that players choose, but that players make systematic "mistakes" or deviations in their choices. The individual deviations from Nash equilibrium strategies and outcomes can come from a variety of different sources, but two reasons seem to predominate discussions. The first reason is strategic uncertainty or bounded rationality [6]. The second reason is that individual utility functions may involve something other than the individual subject's payoffs, typically thought of as a "social preference," in which subjects appear altruistic or fair or seek to reciprocate fairness or seek to limit inequality in payoffs [7, 8].

In the QRE model the deviations between subjects' actual behavior and the Nash equilibrium of a game are represented by a parameter, λ, and as λ approaches zero, behavior is essentially random, while as λ approaches infinity, behavior is consistent with the strategies predicted by Nash equilibrium. The parameter in QRE is a free parameter that is fitted via maximum likelihood estimation and is used to fit the model to the observed behavior. The source or interpretation of the parameter, however, is not determined *a priori* but is left to the analyst. The QRE model predicts that deviations from Nash predictions will be less likely as behavior becomes more costly in terms of the losses a player suffers from non-equilibrium (non-rational) behavior.

The QRE model has been widely demonstrated to provide a better fit to observed behavior than Nash equilibrium, and has become popular in fields outside of economics, such as computer science. The most rigorous statistical examinations of QRE estimate λ from a subset of experimental data and then examine how the estimated λ "predicts" the remaining, unused or out-of-sample data (basically a form of cross-validation or training/test data) [9]. While clever and exciting, this approach falls short as a test of the QRE model for two main reasons.

First, this approach does not make specific predictions about expected behavior prior to the design of the study. That is, the QRE approach does not rule out much, if any, behavior as inconsistent with the theory, data is collected and used to estimate a free parameter that can accommodate almost any observations. Despite the empirical failure of the Nash-equilibrium model, it has the virtue that it does make clear predictions about the expected pattern of observed behavior. To put it another way, is there a distribution of observed behavior that would suggest QRE is an incorrect model of behavior?

Second, statistical approaches to QRE take experimental data and then identify which model best "fits" (by some standard) that data. It is unspecified how to reject QRE in this approach. We could certainly find that some other model better fits the data or that combining QRE with another model leads to a better overall fit to the distribution of the data [9]; however, this approach does not reject the model.

An alternative approach to testing QRE is to identify the underlying source of the deviations from rationality (Nash equilibrium strategies) that give rise to λ, and then design experiments that remove these sources and observe whether deviations still occur. This is an ambitious research program, as there are at present about four dozen causes of bias or error identified in the literature. It is this approach that we pursue here. We discuss in the next section what is required to test a behavioral theory.

2 Testing a Behavioral Theory

We begin by identifying the basic requirements to test a theory. A simple representation is given by the following two statements, in which c = a condition or set of conditions, t = a treatment, and y = an outcome of interest.

if c + t, then y

if c + ~t, then ~y

In laboratory economic experiments, we can think of the t as the particular game being played (i.e. a prisoner's dilemma, chicken, ultimatum, matching pennies, etc.). The conditions, c, represent other parameters of the experimental design (i.e. elements of the protocol, the number of repetitions, possibilities for learning, etc.), and the outcome, y, is how subjects play a given game.

The outcome, y, is a function of t and c; therefore, to test whether the treatment causes the outcome requires that we hold constant the conditions under which the treatment is implemented; otherwise, it may be the changes in the conditions that lead to the presence or absence of a particular outcome. Overall, the framework implies that if we change some element of t (without changing anything else) and the outcome changes, then we can infer that the change in the outcome is a result of the change of t. This framework implies that we must be able to specify, in advance of the research, what are c, t and y. We will know that t is the cause of the outcome if its removal (or change) leads to a change in the outcome. On the other hand, if we remove (or change) t and y still occurs, then we know that it is not the t that is causing y to be observed.

This framework implies a straightforward way to test various theories of behavior, although it can be quite difficult to implement. In a nutshell: 1) identify the treatment and how it causes the outcome of interest; 2) identify a way to remove the aspect of the treatment that causes the behavior, and 3) determine whether the behavior persists. If the behavior persists, then we can rule out that aspect of the treatment as the cause.

3 Testing the QRE model

QRE has gained popularity in behavioral economics, political science, and computer science, with some researchers claiming it should be used as the baseline model for explaining subject behavior [5]. Despite these claims, the QRE model proves quite difficult to test – that is, it is difficult to design a test of QRE that would allow us to reject the approach [10].

One direct way to examine whether QRE can explain behavior is to create situations where the reasons for subjects to deviate from predictions are minimized, ideally to the point of non-existence. To do so requires us to be clear about the sources of the observed errors. In the QRE framework, deviations are thought to come from strategic uncertainty and/or some form of social preference. Extending our previous experiments [11, 12] we employ two different experimental approaches to study whether removing the causes of strategic error actually eliminates those deviations in experimental settings.

Our first approach is to examine subjects in situations where they do not have to make a strategic decision. In these tasks an individual subject makes a decision that by itself completely determines her payoff. Most often this is either a single-play task or it is the last stage of a sequential game of complete and perfect information, thus establishing the conditions for subjects to be sequentially rational. If strategic uncertainty leads to error, then the removal of any strategic element of the task should eliminate the deviations.

The second approach we take combines removing strategic uncertainty with reducing the likelihood that we misunderstand the subjects' utility function (that is we do not observe how social are their preferences). Although we cannot observe utility functions directly, a common reason offered for why the utility function is different than that assumed by classical game theory is that subjects have some form of social preferences in their utility function [7]. In general, these approaches explain deviations from Nash equilibrium by arguing that individuals receive some utility from the payoffs received by other actors. To eliminate this rationale for behavioral deviations, we have subjects play economic games against an algorithmic agent (computer player). We describe these experiments in more detail in the next section. While we have received yarns spun to (post-hoc) explain why subjects taking private and anonymous actions against a computer would act socially with the computer, we do not see any sequentially rational reason why subjects would be concerned with the payoffs a computer receives and would thus deviate from Nash equilibrium strategies. Together, these two approaches are good first steps to gaining leverage over the sources of deviations in game theoretic experiments.

4 Experimental Design

We describe our general experimental design in this section. One of the most important attributes of our experimental design is that we have subjects complete a battery of different, sometimes related, tasks over the course of an experimental session. This will allow us to compare behavior in one task to behavior in another task. As a result, we can examine whether the behavior we observe happens only in one game or if it happens across a variety of games. This is crucial to determining if the behavior is anomalous or reflective of a larger pattern.

The first step we take in our experimental design is to control the conditions of the experiment in a way that makes it unlikely that they are the cause of the deviations in behavior. Subjects were recruited using flyers and email messages distributed across a large public California university and were not compelled to participate in the experiment, although they were given $5 in cash when they showed up. The experiment lasted approximately two hours. The subjects in our experiment completed the tasks using pen and paper in a controlled classroom environment.

In all of the tasks we discuss below subjects played against an anonymous opponents (human or computer) and in all settings we used a double blind protocol to ensure their anonymity from the experimenters as well [13]. In between tasks subjects were also randomly paired with a new subject who was in another room. Subjects also only completed each task one time to ensure that they were not learning about the other subjects in the experiment, which might also cause changes in behavior. To ensure that subjects understood the tasks facing them we utilize very simply experimental tasks and we also quiz subjects before each task to ensure that they understand what they are to do. Subjects are paid for correct answers on the quizzes and if they get an answer wrong we describe the correct answer to them. Overall, subjects very rarely got wrong answers on the quizzes, which gives us confidence that their behavior is not due to a lack of understanding the task. We take great care in our control of the experimental conditions to ensure that we remove the aspects of the experiment that are typically argued to be the causes of Nash equilibrium.

We report on a portion of our battery of tasks here that involve either individual decisions or decisions by subjects acting in the role as the second player of a sequential task. These tasks do not involve strategic uncertainty about the actions of another subject. Players in all of these games know they are randomly paired with another subject in a different room. We report on the choices made by subjects when they are Player 2 in the Trust Game [14]. In this game, Player 1 and Player 2 each begin with a $5 endowment. Player 1 chooses how many dollars to send to Player 2 (ranging from $0 to $5). That amount is tripled by the experimenters and given to Player 2, who then has the sum of the tripled amount received and the initial $5 endowment. Player 2 then chooses how many of those dollars (possibly 0) to transfer to Player 1. Player 2 keeps whatever she does not transfer. In another game, called the Dictator game, The Dictator (Player 1) and the Recipient have endowments equal to exactly what Player 2 and Player 1 had in the Trust Game when the present Dictator played it as Player 2. Accordingly, the Dictator game is identical right down to the specific endowments to the second half of the Trust game: the Dictator is in exactly the position he or she was in as Player 2 in Trust. In effect, each subject replays the second half of the Trust game, but now without the reciprocity frame. The Donation game is identical to The Dictator game, except that each player begins with a $5 endowment and the amount Player 1 chooses to send is quadrupled before it is given to Player 2. In both of these games, the subject keeps all the money he does not send to the other player and there is no action required by the other player.

In what is called the Sequential Chicken Game, players maximize their payoffs if they choose the opposite of the action chosen by the other player. We implement a sequential version in which the 1st player chooses STOP or GO, and then the 2nd player observes the 1st player's choice before also choosing between STOP or GO. The payoffs in the game are shown in Table 1, where Player 1's payoffs are listed first and the grey cells represent deviations from Nash equilibrium.

		Player 2	
		GO	STOP
Player 1	GO	\$0, \$0	\$5, \$3
	STOP	\$3, \$5	\$4, \$4

Table 1. Payoffs in Sequential Chicken Game where Player 1 moves first and play is revealed to Player 2

At the end of the experiment, we present the subjects with tasks that would allow them to learn something about the choices made by subjects in the other room. They are asked to make a choice as Player 2 in the Trust game as one of the final tasks. In this last task, we have no choice but to provide subjects with information about what other subjects have done – Player 2 in Trust must know what Player 1 chose to send.

We also conducted the same experimental games reported above but had players play against computer opponents. We informed the subjects that the computers (their opponents) would always take the action that maximized the computers financial gain for the present task– that is, the computers played as perfect "Nash players." Playing a game against a computer represents a change in the treatment: it removes exactly the elements associated with playing against an anonymous person.

5 Behavior without opportunity for strategic error

The experiments we designed—including those in which people play each other—explicitly remove the opportunity for subjects to make miscalculations about the actions of others, because in all of the settings, subjects' payoffs are determined solely by their own actions. We go to great lengths to ensure subject anonymity from each other and the experimenters, and we believe that we have done as well as possible in that regard. This does not remove the possibility that subjects' utility is not related solely to their monetary payoffs, which remains a source of deviations.

	Amount passed by subjects in each experimental task										
	$0	$1	$2	$3	$4	$5	$6	$7	$8	$9	$10
Donation	87	32	23	11	6	21	n/a	n/a	n/a	n/a	n/a
Dictator	132	11	18	8	1	7	0	1	0	1	1
Trust Player 2	118	10	20	8	8	8	0	1	0	1	6

Table 2. Choices in three different experimental tasks when playing with other humans. Trust results include the 80 players who received $0 from Player 1. If those are excluded the number who sent $0 drops to 38**.**

In Table 2 we report the results for the different tasks in which subjects had to decide how much, if any, money to pass to another player. Recall that in all of these settings the choices of players determine directly and solely what is kept or passed there was no need for strategic calculation. Therefore, the possibility of deviations caused by the inability of players to understand and predict the strategic interaction is eliminated by the experimental design. Yet, subjects send money to another anonymous player, which contrasts with the Nash equilibrium. In all three tasks, many of the players (the number of subjects is 180 in these tasks) send money to the other player. Furthermore, the assumption within QRE that costly mistakes will be less likely is inconsistent with these data. The incidence of deviations does not decline in a consistent pattern as we move away from the choice to send $0.

We turn now to the sequential chicken game and in particular to examining the behavior of Player 2, who already knows the move of Player 1 before making her own choice in this task. There are 17 subjects in the role of player 2 that choose GO even after the 1st player already chose "GO." These players knowingly made a choice that leads both subjects to get a payoff of $0 instead of the second player's choosing STOP and guaranteeing $3 for herself and $5 for the first player. In addition, there are 36 players in the role of Player 2 who choose STOP even after Player 1 also chose STOP. This choice gives Player 2 a lower payoff than if she chooses GO in this situation. Player 2 faces zero strategic uncertainty, because, before choosing, Player 2 is informed of Player 1's choice. Yet, overall, 53 subjects (29.4%) chose the action that leads to the lower payoff.

Across a wide number of tasks, we observe that subjects take actions that are inconsistent with standard game theoretic expectations and that cannot be explained by strategic uncertainty. However, it may be that these deviations come not from strategic uncertainty but from the fact that the subjects' utility functions differ from our assumptions about them, which is another source of behavioral deviations in the QRE approach. We turn now to a set of experiments designed to address that question.

6 Misunderstood utility functions

The results in the prior section demonstrate that subjects still deviate from basic Nash equilibrium even when we remove the strategic elements; however, it is still possible that the behavior we observe is a result of the fact that we misunderstand their utility functions. In this setting, that will usually mean that individual's utility functions include something other than their own monetary payoffs. This could include a concern for reducing inequality, maximizing social welfare, punishing others for the choices they made, or anything else that makes them take actions in which they knowingly earn less money because their utility depends in some way on what others earn as well. This is not a very specific argument, and the unobservability of utility functions means that if we allow ourselves to utilize post-hoc rationalizations, we can justify nearly any behavior via utility function modification.

We turn now to examining the actions our human subjects take when matched to a computer that they know will take actions to maximize its payoffs – a perfect Nash player. In Table 3, we display the humans' behavior (for 40 subjects) across a range of different experimental tasks. This time we focus on the decisions humans make as the first player in Trust, choosing how much of $5 to pass the computer, rather than the second player, because computer players never pass money to the other player and the computer will never return money as Player 2 in the Trust game. Despite the fact that there is no strategic uncertainty and players know they are playing a computer that will maximize its own payoffs, we still observe that players choose to pass money to the computer player in all three situations. In both Donation and Trust, at least 25% of subjects in the role of Player 1 pass money to a computer opponent.

In the sequential chicken game, the 2nd player is always faced with a computer who chose "GO," which is the Nash equilibrium. In this setting, the human players chose "STOP" 39 out of the 40 times.

The overall pattern of behavior when subjects are playing against computer players is much closer to Nash behavior. By combining tasks where our subjects' actions directly determine their payoffs with computer opponents who will always take the payoff maximizing option, we remove all of the strategic uncertainty that often faces players in game theoretic experiments. In addition, play against a computer eliminates much (though perhaps not all) of the ways that a subject's utility function can depend on the payoffs to another subject. Despite removing the explanations for deviations that underpin QRE, we still observe behavior that is inconsistent with Nash equilibrium. In the Donation and Trust games, 25% of subjects deviate from the Nash predictions, whereas in the sequential chicken game, only 2.5% of subjects deviate.

	Amount passed by subjects in each experimental task					
	$0	$1	$2	$3	$4	$5
Donation	30	3	5	2	0	0
Dictator	38	0	1	0	0	1
Trust Player 1	28	3	5	1	1	2

Table 3: Choices when subjects play against computers who play as perfect Nash players.

These results suggest that there may be something about the games themselves that makes them more amenable to non-Nash behavior even after we remove the standard rationales that underpin the QRE model.

7 Discussion

Prior research has established that the quantal response equilibrium provides a better fit to experimental behavior than Nash equilibrium, but it is also clear from the results we report that subjects' behavioral deviations do not occur simply from the standard explanations that underpin QRE. We show via a variety of experiments that even when we design the experimental conditions and treatments to remove the reasons why people might deviate from Nash equilibrium we still observe such deviations. The deviations also differ across games, but we lack an explanation for why that occurs. We could still estimate the lambda parameter of the QRE model for that data we discuss in this paper, but the parameter would have no substantive meaning given that the causes of the deviations do not exist in these experiments.

Additionally, as we report elsewhere, the decision to send money in one of these tasks is not predictive of the decision to send money in one of the other tasks [11]. That is, players may send money in one setting but then not send money in another very similar setting. In the context of the QRE model, this implies that the lambda parameter is not consistent within a subject across very similar, albeit not identical, tasks. The QRE model does not assume that the parameter will be consistent, but it makes it difficult to use the QRE approach to predict behavior in new tasks if individual deviations are inconsistent across similar tasks.

The results in this paper suggest that even though QRE improves on Nash equilibrium in its ability to fit observed behavior, we cannot be sure that the observed deviations come from the sources assumed in QRE, because deviations persist even when subjects are placed in environments that remove the sources of the deviations. Therefore, these results suggest that we need to revise the basic cognitive and behavioral assumptions of game theory in order to build models that more accurately reflect human capabilities and actions.

Acknowledgments. McCubbins acknowledges the support of the National Science Foundation under Grant Number 0905645. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Turner acknowledges the support of the Centre for Advanced Study at the Norwegian Academy of Science and Letters.

References

1. Smith, Vernon. 2000. Bargaining and Market Behavior: Essays in Experimental Economics. Cambridge University Press.
2. Camerer, Colin. 2003. Behavioral Game Theory. Cambridge University Press.
3. Gigerenzer, Gerd. 2008. Rationality for Mortals: How People Cope with Uncertainty. New York: Oxford University Press.
4. McKelvey, R. and Palfrey, T. 1995. Quantial Response Equilibria for Normal Form Games," Games and Economic Behavior. 10, 6-38
5. Camerer, C., Ho, T-H, and Chong, J-K. 2001. "Behavioral Game Theory: Thinking, Learning and Teaching," Paper Presented at the Nobel Prize Symposium.
6. Simon, Herbert. 1956. "Rational Choice and the Structure of Environments." Psych. Rev. 63: 129-138.
7. Rabin, Matthew. 1993. "Incorporating Fairness Into Game Theory and Economics." The American Economic Review.83, 1281-1302.
8. Souten, J., DeCremer, D., & van Dijk, Erik. 2006. "Violating Equality in Social Dilemmas: Emotional and Retributive Reactions as a Function of Trust, Attribution, and Honesty." Personality & Social Psychology Bulletin. 32, 894-906P
9. Wright, J and K. Leyton-Brown. 2012. "Beyond Equilibrium: Predicting Human Behavior in Normal Form Games" Conference of the Association for the Advancement of Artificial Intelligence (AAAI-10).
10. Haile, Philip A. Ali Hortaçsu and Grigory Kosenok. 2008. "On the empirical content of quantal response equilibrium." August. Cowles Foundation Discussion Paper no. 1227
11. McCubbins, Mathew D., Mark Turner, and Nicholas Weller. 2012a. "The Mythology of Game Theory." Proceedings of the International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction. Springer Lecture Notes in Computer Science. Eds. Shanchieh Jay Yang, Ariel Greenberg and Mica Endsley.
12. McCubbins, Mathew D., Mark Turner, and Nicholas Weller. 2012b. "The Challenge of Flexible Intelligence for Models of Human Behavior." Association for Advancement of Artificial Intelligence Spring Symposium on Game Theory for Security, Sustainability and Health. AAI Technical Report.
13. Hoffman, E., McCabe, K., Shachat, K., Smith, V., 1994. "Preferences, Property Rights and Anonymity in Bargaining Games." Games and Economic Behavior. 7, 346-380
14. Berg, J., Dickhaut, J., McCabe, K. 1995. "Trust, Reciprocity and Social History." Games and Economic Behavior. 10, 122-142.