Journal of Theoretical Politics 18(4): 498–511
DOI: 10.1177/0951629806067455
http://jtp.sagepub.com

STRONG AND WRONG

THE USE OF RATIONAL CHOICE THEORY IN EXPERIMENTAL ECONOMICS

Andrew Schotter

ABSTRACT

The main purpose of rational choice theory is to lay out in clear and transparent terms what conditions are necessary and/or sufficient for the validity of statements about consistent human behavior. Strong criteria for rationality are 'wrong' if understood as a positive description. However, their very strength provides a sharp guide for experimental social science's project of mapping the properties of individual behavior.

KEY WORDS • behavioral economics • experiments • rational choice

1. Introduction

Many people have summarized the theme of George W. Bush's presidential campaign as saying it is better to be 'strong and wrong' than to be wavering and inconclusive (a flip-flopper). Ironically (or perhaps unfortunately) in this article I am going to argue something similar about economic theory specifically and rational choice theory in general. It is useful to be strong and wrong. The difference between myself and our President, however, is that I consider the benefits of being strong and wrong to come mostly from what we can learn from being wrong rather than the thrill of pretending to be strong. Before entering into the main argument of this article, I begin with a brief outline of what I see as the epistemological benefits of rational choice theory.

2. The Rational Choice Debate

I must admit that I am always confused when I hear the debates over rational choice theory that take place in other social sciences. To be more precise, I consider rational choice theory and the modelling that accompanies it, to be simply an effort to delineate under what circumstances certain statements are valid. Put differently, it is an effort to separate cheap talk from logically coherent statements that follow from transparent assumptions. For example, say that you are speaking to a friend who says, '*As we all know*, every time you raise the price of a good people consume less of it'. Now I would suspect that if I were to take a person on the street and say that to them, they would certainly agree. The statement obviously has a lot of intuitive appeal and is actually convincing to most people. It seems to be a no-brainer. In fact your friend actually stated a theorem that, like all theorems, needs to be proven. The theorem is that all demand functions are downward sloping. Of course we know this is false. If the price of a good is raised then, whether more or less of it is consumed, depends on the relative magnitudes of the income and substitution effect for that good as well as its income elasticity of demand, that is, whether the good is inferior or superior. My claim is that the main purpose of rational choice theory is to lay out in clear and transparent terms what conditions are necessary and/or sufficient for these types of statements to be valid. Using rational choice techniques allows us to separate the logical wheat from the intuitive sounding but wrong chaff.

This is a rather innocuous agenda and should be unobjectionable. It simply follows from asking the question of what it means to know something in the social sciences. To my mind there are only two answers to this question. One is I know something, that is, a statement is valid logically, if it is derived from a set of assumptions that are clearly laid out using accepted rules of deduction. Alternatively, I can state that I know something if I can be convinced by the tools of empirical analysis that it is corroborated by data.

To accomplish this agenda it is often necessary to build models or prove theorems using an axiomatic approach. Let's take model building. In our demand-theory example above, in order to establish that the original statement was false we economists had to make a variety of assumptions about the preferences and calculating abilities of economic agents. For example, economic agents were assumed to have a complete binary preference relation defined on the set of goods they could consume. The relation is assumed to be continuous and transitive (and convex) and with those assumptions it can be proven that such preferences can be represented by a utility function. Demand functions can be derived assuming that people maximize under constraints and can do the calculations needed by the theory. So it is actually hard work to prove that our original innocuous statement was wrong. Without the deductive process embedded in rational choice theory we have no way of knowing if what comes out of our mouths is logically consistent or the circumstances in which it is. So to my mind one cannot object to the process of creating rational choice theory or modelling if that is the objective. In fact, I am agnostic with respect to the assumptions that one makes when building models. While economic theory has been wedded to certain types of assumptions I would be happy to allow a wide variety of assumptions to be used as long as the process generating results was kept intact.

What I think is probably where people get confused by the rational choice methodology is when they believe that the results of the theorems proven in this fashion are correct predictors of human behavior. Mature thinkers understand that this is not true. Models are created simply to lay out the conditions upon which statements (theorems) are logically valid. From that point on it is an empirical question as to whether we should use these theorems as descriptive of human behavior or not and later in this article I will discuss how experimental economics tries to sort these things out. What I will claim is that using rational choice theory to structure this empirical investigation is an efficient and fruitful way to proceed. Rational choice theory does not provide us with the truth about human behavior but rather with a very compelling platform with which we can seek the truth.

In summary, what I have said has been said in various ways by Reinhard Selten: game theory or rational choice theory is a tool for proving theorems mostly about what rational individuals will do in hypothetical situations. Whether that theory applies to the behavior of an ordinary person off the street is an empirical (experimental) question. Still, in our empirical search for the truth I will argue that starting from a rational choice model is a good way to proceed and that the combination of rational choice models and empirical support is the only way to gather knowledge in the social sciences.

3. Niels Bohr and Being Logical

There is a story that Niels Bohr, the great physicist, was once in a fierce argument about a point in physics. At a moment of great frustration he turned to his colleagues and said, ‘You are not making sense you’re just being logical’. To my mind Bohr’s statement gives us an insight into where rational choice theorist and economists get into trouble – we are many times more interested in being logical than in making sense.

There are many examples that can be used to illustrate this point. For example, take the finitely repeated prisoners’ dilemma game. Because there is a dominant strategy equilibrium in each sub-game of this repeated engagement, backward induction along with common knowledge of rationality predicts that people will defect for all periods. Of course we know from experimental evidence and casual empiricism that such long-winded backward induction arguments make no sense to perfectly sensible people. Rather, people tend to cooperate in the beginning of such interactions and toward the end defect but certainly don’t defect all the time.

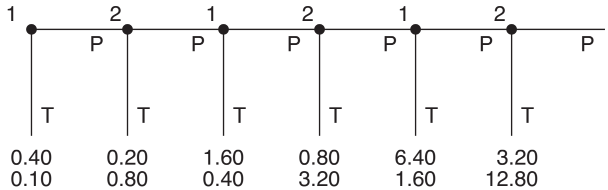

Figure 1. The Centipede Game

Another example comes from the centipede game as studied by McKelvey and Palfrey (1992) shown in Figure 1.

Here again, backward induction leads to predictions which are logical but not what sensible people would do. More precisely, backward induction prescribes a choice of down on the first move of the game. Yet Player 1 must realize that if he moves right and Player 2 passes at least once, he will never be worse off than if he moved down on the first move. So all he needs to do is trust the other player for a short period of time in order to have no regrets about passing on the first move. In fact, in the centipede game the backward induction solution is invariant to the length of the game yet clearly as the game gets longer, greater and greater incentives are created for players not to move down on the first move. When this game was played by McKelvey and Palfrey (1992) only 6 per cent of the subjects moved down on the first move. The backward induction predictions were logical but made no sense.

I obviously could go on and on with examples like this but the point should be clear: rational choice theory is a tool for clarifying statements not a tool for predicting behavior. Now comes a question whose answer is actually the point of my article. If rational choice theory is likely to be wrong, why would we ever want to use it as the starting point for empirical or experimental work? How can it serve as a platform for seeking the truth? Because it is a strong theory, that is, one capable of making strong point predictions, I will argue that almost by definition it must be wrong. For instance, the theory could predict well if the world we test it on fulfilled all of the assumptions underlying the theory. However, the theory was not constructed to prove a behavioral regularity in the real world but rather to outline the type of world that would need to exist in order for the theory to be true. Those two worlds, the real world and the theoretically necessary world, need not coincide at all so there is no reason to think the theory will predict well. Constructing the theory offers no insight into the likelihood of the theory being correct. Note also that I am giving up the typical positivist justification for rational choice theory which is that while the assumptions are wrong the predictions it makes are, on average, correct. I am saying almost the opposite. So what's so good about being strong and wrong?

4. Strong and Wrong in Experimental Work

Consider the Ultimatum game. In this game one person is given $100 and asked to split it with another anonymous person. He, the Proposer, makes a proposal x indicating what he wants to keep and offers the remainder, $1 - x$, to the other person, the Responder. The Responder either accepts or rejects. If she rejects they both get 0 and if she accepts, the Proposer gets x and the Responder $1 - x$.

This is an extremely simple game and one that touches aspects of all social science. What I find amazing is that the only social science that gives a point prediction as to what is likely to happen here is economics. Despite the years of studying human behavior neither psychology, sociology or anthropology make a point prediction even in this simple situation. If that is the case, it is hard to believe that they could generate the types of testable hypotheses that one needs to seek knowledge.

According to economic theory, a rational Proposer will offer an assumed rational Responder ε and keep $1 - \varepsilon$ for himself. This is true because once the offer is made the Responder has a choice between ε and nothing and since it is assumed that all agents' utility functions are monotonically increasing in money and all players are selfish and care only about their own payoffs and not their relative payoffs, ε will be accepted. So if this experiment were done in the lab a true believer would predict that all offers would be of ε. Of course this is not what we find. What is found in the lab is that the modal offer is between 40 per cent and 50 per cent of whatever pie there is (see Camerer, 2003) and that offers below 20 per cent are rare. So the theory is not supported by data yet is logically consistent using the assumptions made.

The fact that the theory predicts incorrectly does not mean that it is not valuable. The theory must only be considered a good starting point for predictions about behavior. In some circumstances it organizes the data well, while in others it does not. But it is essential in allowing us to organize our thinking about what our expectations are and formulating new theories to explain the data when things go wrong.

The Ultimatum game is actually a perfect place to help explain what I am talking about. Since the theory is predicated on the assumption that people are selfish and care only about their own payoffs and not their relative payoff as compared to others, they are unaffected by feelings of fairness, inequality, and so on. Therefore, we could have predicted that such a theory would fail since we know that generally people do care about others and their comparisons to them. What is interesting is how badly the theory fails. In other words, how badly are the assumptions about human nature. In this case they are rather bad but that only means we need to revise the theory on the basis of the data, not abandon the approach. To do this a number of scholars have proposed different types of utility functions that may be appropriate for economic agents.

For example, the common rational choice approach assumes that people have utility functions of the form $U_i = U_i(y_i)$, with $U_i'(y_i) > 0$ and typically $U_i''(y_i) < 0$. Here note that the utility that an economic agent receives from any social engagement is a function of their own payoff and is independent of how much others receive. This is the selfishness assumption that has dominated economic theory for 200 years and comes from Adam Smith's assertion that it is not from the benevolence of the butcher or baker that we get our meat or bread but rather from their selfish desire to improve their own situation. Using this utility function leads us to the poor predictions discussed earlier. Recently Ernst Fehr and Klaus Schmidt (1999), and Gary Bolton and Axel Ockenfels (2000) have proposed other utility functions which if used could explain the data observed in these experiments and many more.

In this utility function people care about their own material payoff as well as inequality of income both to themselves and others. For example, take the two-person version of the problem. Here you either receive more or less income than your partner. If you receive less, given your income, you suffer by α_i times the difference and if you receive more you suffer by β_i times the difference. Assuming $\alpha_i \geq \beta_i$ means that a player suffers more from inequality if it is to his disadvantage, $\beta_i \geq 0$ means that we rule our people who like to be better off than others, while $\beta_i = 0.5$ means that the person is indifferent between keeping an extra dollar for himself or giving it up to the other person. Finally, an assumption that $\beta_i > 1$ would mean person is willing to throw away a dollar to decrease his relative advantage over the other person. They rule this out.

With this utility function, and maintaining the maximization hypothesis of rational choice theory, Fehr and Schmidt explained considerably more of the data than do models based on the selfishness axiom.

Equally interesting is the fact that this same theory can be used to explain what appears to be selfish, or at least competitive, behavior in markets. For example, in a paper by Roth et al. (1991), it is observed that when you play a multi-person version of the Ultimatum game where $n > 2$ and Proposers simultaneously make offers to split a pie, say $100, with the highest offer being accepted or rejected by the Responder, the offers made tend to reach the competitive level where the Responder receives all of the surplus. In other words, competition leads to a very unfair outcome. For such a game it is interesting that both the Fehr–Schmidt and Bolton–Ockenfels model predict such behavior. In other words, people who care about inequality are very able to behave as if they were selfish.

It is my claim that the benefit of rational choice theory is that it helps structure our search for models of human behavior that organize both field and experimental data. The benefit of the approach is not so much that it is strong but rather that the manner with which it is wrong may be very instructive in allowing us to structure our search for the truth by providing us with a very well defined model whose elements can be stripped away one by one until we find that part which leads to prediction failures.

For example, as stated earlier, in the Fehr–Schmidt case, what failed was the selfishness axiom, not the maximization principle. People don't fail to offer ε because they can't do the backward induction, they fail to do it because the backward induction leads to unfair results or results where it is very cheap to reject unfair offers. Having a well defined but wrong model helps us come to this conclusion. Put differently, rational choice theory must be considered a platform from which we can seek the truth, not a repository of logically valid yet untested predictions.

This is the role that rational choice theory has played in experimental economics and, I think, the role that any theory should play in the process of discovering the truth. There are only two other ways to seek knowledge. One is a purely empirical way which tends to be the path chosen by psychology while the other is a method that shuns both theory and empirical analysis. In the former approach, intuition leads to experimental examination which leads to results which modifies intuition etc. After a plethora of small experiments are done, one begins to see the forest fore the trees and pieces knowledge together in this completely empirical manner. While I believe that theory-guided empiricism is more efficient, I must say that over the years I have learned that the purely empirical approach can provide another way to seek the truth. The second approach, while possibly a source of inspiration to many, cannot offer us any guidance as to what statements are true or false.

Let me offer another example for how strong and wrong a benefit may be. Game theory is destined to fail empirically because in many situations, like the Ultimatum game, it makes point predictions while we know that observed behavior, even observed highly rational behavior, is many times noisy. So the theory is not always well calibrated but rather highly resolved. It gives sharp or strong predictions in many, if not all, situations. People make mistakes, fail to be rational to some extent, but not fully, and so on. So a theory that allows for no variance in behavior is destined to fail. Again, it would be wrong to take this failure as a weakness. On the contrary it must be considered a strength if the process of discovering its weakness leads us to construct a better theory.

One example of where this may be true is the theory of Quantal Response Equilibrium. To motivate our discussion here, consider the Traveler's Dilemma game first discussed by Basu (1994) and investigated experimentally by Capra et al. (1999). In this game two travelers travelling with identical luggage lose their bags at the same time. The bags are identical as are the contents. The airline knows from bitter past experience that when bags are lost people tend to claim the maximum value for them no matter what their true value is. To eliminate this moral hazard, the airlines institute the following reimbursement scheme. Both travelers must state an amount that their bag is worth. The airlines takes each stated amount and compares them. It pays each of the travelers the minimum of the amount stated but penalizes the person stating the bigger amount by R and transfers that amount to the other. Hence, if the bags are worth something in the interval [180, 300] and person A states a value of \$190 and the other, person B, states a value of \$240 and $R = \$20$, then person A will be paid \$190 + \$20 while person B will be paid \$190–\$20.

Game theory makes two predictions for this situation. First it predicts that each person will claim the least amount of value for the bag. Second, it predicts that the first result is invariant to the amount R that is used as a bonus or punishment. So there are two predictions. One is a point prediction that all offers will be identical and minimal and the second is a comparative static prediction that as R changes the players do not respond.

This theory fails when confronted with the data. As we see in Figure 2, while the theory predicts that all of the messages sent would be concentrated

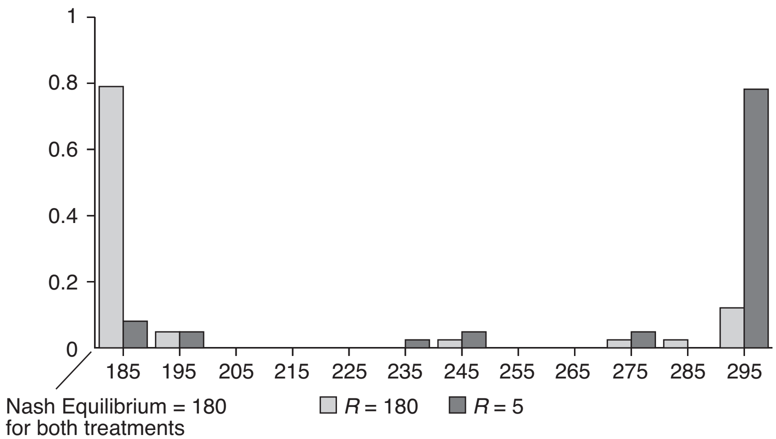

Figure 2. Claim Frequencies in a Traveler's Dilemma for $R = 180$ (light bars) and $R = 5$ (dark bars)

around the minimum value, so that all of the frequency should be in the interval [180, 185], in fact the values are distributed over the entire interval. Further, note that unlike the theory, the data do respond to R. When R is small most of the submitted values are high while when it is high most of them are small. What this calls for then is a theory that can simultaneously explain the variance of behavior as well as its sensitivity to the underlying parameters of the problem, that is, R.

Such a theory was proposed by Richard McKelvey and Tom Palfrey (1995) and labeled by them Quantal Response Theory. What this theory basically says is that instead of making a sharp best response to one's beliefs about one's opponent, as conventional game theory suggests, people may choose a stochastic best response by choosing strategies that have higher expected utilities with higher likelihood yet still leaving open the possibility that they will choose worse alternatives as well. For example, say that in a game where each player has two strategies each player chooses his strategies using the following logit choice function,

$$p(D_1) = \frac{e^{(\pi_1^e)/\lambda}}{e^{(\pi_1^e)/\lambda} + e^{(\pi_2^e)/\lambda}}.$$

Here π_j^e, $j = 1, 2$, is the expected profit that player i expects to receive if he or she uses strategy j. Now this expectation is taken over this player's beliefs about the other player and if we find a pair of beliefs for which the distributions of choices match the distribution of beliefs about choices (i.e., a fixed point), we have reached a Quantal Response equilibrium. The parameter μ describes the precision with which subjects best response to their beliefs about the other players. If the strategy set was continuous, as is true in the Traveler's Dilemma Game, we would have

$$f_i(x) = \frac{e^{(\pi_1^e(x))/\lambda}}{\int_{\underline{x}}^{\bar{x}} e^{(\pi_1^e(y))/\lambda} dy}.$$

What does this theory offer us? Basically it has the two features we desired before. First, rather than making point predictions this theory offers us equilibrium choice distributions as an equilibrium concept. Second, as Figure 3 indicates, it allows us to track the movement of the data as R changes across treatments. For example, look at Figure 2.

In this figure we see the equilibrium distributions of the Quantal Response Equilibrium theory as parametrized by the experiment as well as the time path of mean choices of subjects as the game is iterated 10 times. What is clear from this is that the change in the observed data obviously moves in the direction of the equilibrium predictions as we change R from 50 to 20 to 10.

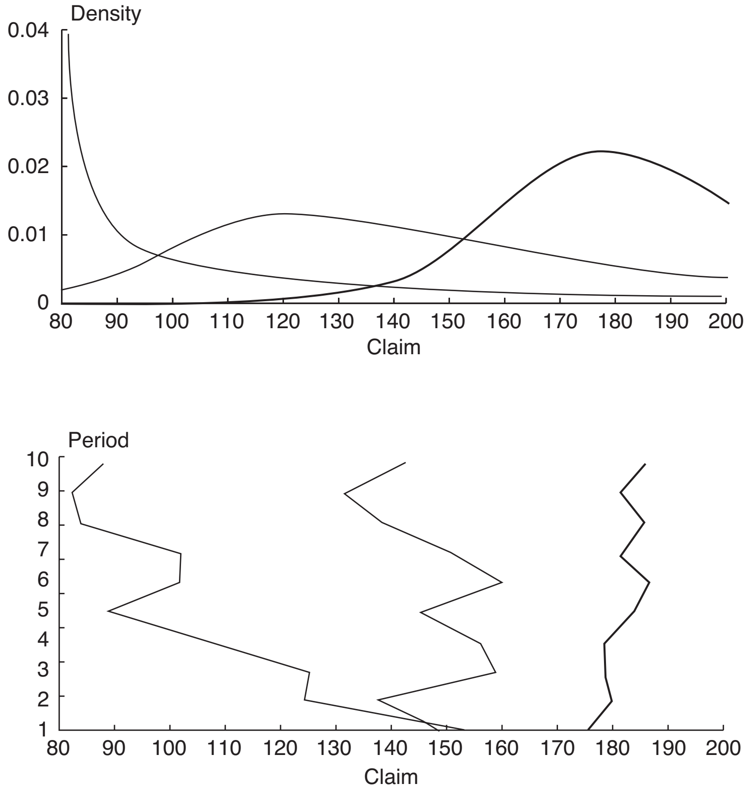

Figure 3. The Traveler's Dilemma: Predicted Claim Densities (Top) and Actual Claim Averages (Bottom) for $R = 50$ (Left), $R = 20$ (Center), and $R = 10$ (right)

This example again illustrates the value of rational choice theory. It is the theory that offers us a platform to understand what is failing in our assumptions and to change those elements without throwing the rest of the theory away. While in the Ultimatum game case we had to alter the utility function of the players while leaving maximization intact, here we leave the utility function intact but generalize what maximization means by expanding it to include stochastic best responses rather than deterministic best responses. It would be hard for me to envision how we would go about explaining this behavior without a starting point that offers what we might expect in the hypothetical world of rational choice theory. The theory allows to classify data as deviations from a baseline and hence tells us what needs to be explained. With no underlying theory it would be hard to separate the observations that are to be expected from those that need explanation.

5. Unpacking What Goes Wrong

From our discussion so far it should be clear that rational choice theory can fail for many reasons. For example, it requires that normal people be able to perform a variety of mental tasks which they may not be able to do. For instance, as discussed earlier, we know that some rational choice arguments require that people be able to string together long backward induction arguments of any finite or infinite duration. Laboratory evidence tends to suggest that three stages may be the most that people are capable of doing (see Camerer, 2003). People are also supposed to be able to iteratively eliminate dominated actions yet we know from a variety of laboratory experiments that people consistently fail at such tasks (most notably see the literature on the 2/3rd's guessing game of Nagel (1995)). People fail to process probabilities correctly (Kahneman and Tversky, 1972), treat losses differently than gains (Kahneman and Tversky, 1979), and on and on. Still all of these 'anomalies' are interesting only because we have an underlying theory from which deviations can occur – there can be no interesting anomalies without an interesting (or uninteresting) theory to generate them. We have made so much progress precisely because we have a strong and wrong theory to function as a baseline from which we measure deviations.

6. What Happens When the Theory Works: The Positivist Argument Revisited

Up until now I have only discussed situations where rational choice theory has failed to predict behavior. What happens when, God forbid, it predicts well? Well, one might think that at that point we might break out the champagne and smugly celebrate. This of course would be the wrong response. In order for rational choice theory or any behavioral theory to be valid, it must not only predict well but do so for the right reason, that is, the reasons laid out by the theory. This is important for many reasons. First, of course, rational choice theory and game theory lay out a process of thinking that is supposed to lead a player to his or her decision. If the player makes the right choice but follows another heuristic, he or she is bound to go astray from the theory when the circumstances he or she faces changes (i.e., we change a parameter defining the game) and the two heuristics, the theoretical one and the one being used, make different predictions. For example, consider the following game:

	Column Player 1	Column Player 2
Row Player 1	80, 40	40, 80
2	40, 80	80, 40

This is a constant sum game with a unique mixed strategy equilibrium. Here the Nash equilibrium is for Row to play strategies 1 and 2 with probabilities (.5, .5) and Column to play strategies 1 and 2 with probabilities (.5, .5).

When this game is offered to subjects in a lab, they tend to play 1 and 2 about 50 per cent of the time (actually 48% for row and 52% for column) as they are supposed to (see Holt and Goeree, 2001). So if this were the only experiment you ran you might conclude that the theory does well. But there are many reasons why this theory could be predicting well. One is that the players are using the game theoretic logic and mixing their strategies so as to make their opponent indifferent between using either of his or her two strategies. Another is that since the game is symmetric, a symmetric strategy seems focal or natural and it is chosen for that reason. If this is the case, the conformity of the data with the theory is a Pyrrhic victory. To separate these two explanations, Holt and Goeree (2001) ran the following game:

	Column Player 1	Column Player 2
Row Player 1	320, 40	40, 80
2	40, 80	80, 40

Notice that this game is identical to the original one but here the payoff to Row Player is increased dramatically in the upper left-hand corner cell. Game theory predicts that despite this change Row Player will not alter his play from (.5, .5) since in equilibrium each player mixes to make the other player indifferent and here the other player's payoffs have stayed the same and since that is the only ingredient into Row Player's mixing strategy, he should not change his behavior. The new equilibrium strategy after the change is for Row to choose strategies 1 and 2 with probabilities (.5, .5) and Column to play (.10, .90).

This is not what happens in the lab. The Row Player chooses Strategy 1 96 per cent of the time and Strategy 2 only 4 per cent while the Column Player chooses Strategies 1 and 2 16 per cent and 84 per cent of the time, respectively. This presents clear evidence that, while in the original game players were behaving according to theory, they were not doing so for the right reasons. Hence, the positivist claim that we do not need to look at the realism of assumptions or models but only at the accuracy of the theory's predictions needs to be challenged. Good predictions alone do not validate a theory. What is needed is a theory that predicts well for the right reasons. This is why laboratory experimentation is so important. The world rarely offers us a chance to test a theory under two different controlled circumstances. The laboratory does and hence it is an ideal place to test theory.

Another example of this phenomenon is given by Partow and Schotter (1993). They investigate a more complicated signalling game of incomplete information where the relevant concept is the intuitive criterion. This refinement of Nash equilibrium asked subjects to go through a fairly complicated process to eliminate certain equilibria as possibilities where multiple equilibria exist. What Partow and Schotter show is that if we play this game by giving players information only about their own payoffs, so that it would be impossible to follow the complex logic prescribed by the intuitive criterion, the play the game in a virtually identical manner. This indicates that there must be another heuristic that they are following to determine their behavior. Again, we get the right game theoretical or rational choice outcome for the wrong reasons. Still, none of these results are possible with the starting point of the theory.

7. Conclusion

In this article I have argued that rational choice theory is an efficient platform from which to study human behavior. I have stressed that we must not judge it by its ability to predict human behavior but rather take it as a baseline that offers us a set of predictions which can be used to evaluate theories. Like it or not, rational choice theory is the only theory existing in the social sciences that provides such a platform. However, what is interesting about the theory is not its ability to be right but the interesting ways in which it is wrong.

This is obviously a chauvinistic viewpoint. First, it basically says let's see how much of the variance of human behavior can be explained by rational choice modelling before even looking for other variables that may help increase our explanatory power. This is not actually the case. We start with rational choice theory not because we expect it to be correct but because it has the precision to make point predictions that can be tested and is based on a set of assumptions that have been clearly laid out. Hence when its predictions fail to be confirmed we have a place to start looking for help to explain what went wrong and what can be added to make things right. The things that are added may involve information from a wide variety of fields but are not the basis upon which our experiments are designed.

Finally, I have stressed that while being wrong can be useful in discovering the limitations of one's theory, being right is not the end of the story. We must be right for the right reasons. Data consistent with theory must be so for the reasons stipulated in the theory or else the theory can not be expected to be accurate in predicting comparative static changes.

REFERENCES

Basu, Kaushik (1994) 'The Traveler's Dilemma: Paradoxes of Rationality in Game Theory', *American Economic Review* 84(2): 391–5.

Bolton, Gary and Axel Ockenfels (2000) 'ERC: A Theory of Equity, Reciprocity and Competition', *American Economic Review* 90: 166–93.

Camerer, Colin (2003) *Behavioral Game Theory*. Princeton, NJ: Princeton University Press.

Capra, C. Monica, Jacob K. Goeree, Rosario Gomez and Charles A. Holt (1999) 'Anomalous Behavior in a Traveler's Dilemma?', *American Economic Review* 89(3): 678–90.

Fehr, Ernst and K. Schmidt (1999) 'A Theory of Fairness, Competition and Cooperation', *Quarterly Journal of Economics* 114: 817–68.

Holt, Charles A. and Jacob Goeree (2001) 'Ten Little Treasures of Game Theory, and Ten Intuitive Contradictions', *American Economic Review* 91: 1402–22.

Kahneman, Daniel and Amos Tversky (1972) 'Subjective Probability: A Judgement of Representativeness', *Cognitive Psychology* 3: 430–54.

Kahneman, Daniel and Amos Tversky (1979) 'Prospect Theory: An Analysis of Decisions Under Risk', *Econometrica* 47: 313–27.

McKelvey, Richard D. and Thomas R. Palfrey (1992) 'An Experimental Study of the Centipede Game', *Econometrica* 60: 803–36.

McKelvey, Richard D. and Thomas R. Palfrey (1995) 'Quantal Response Equilibria for Normal Form Games', *Games and Economic Behavior* 10: 6–38.

Nagel, Rosemarie (1995) 'Unraveling in Guessing Games: An Experimental Study', *American Economic Review* 85(5): 1313–26.

Partow, Z. and A. Schotter (1993) 'Does Game Theory Predict Well for the Wrong Reasons: An Experimental Investigation', C. V. Starr Research Report 93–46, NYU.

Roth, A., V. Prasnikar, M. Okuno-Fujiwara and S. Zamir (1991) 'Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study', *American Economic Review* 81: 1068–95.

ANDREW SCHOTTER is Professor of Economics at New York University and Director of the NYU Center for Experimental Social Science. His main areas of research are economic theory, game theory, and experimental economics. He is the author of *The Economic Theory of Social Institutions* and author or co-author of a number of articles in *American Economic Review*, *Journal of Political Economy*, *Quarterly Journal of Economics*, and other journals. ADDRESS: New York University, Department of Economics, 269 Mercer St., 7th Floor, New York, NY, 10003–6687, USA
[email: andrew.schotter@nyu.edu].