Comment on Tullock, Hechter, and Wildavsky

Eric Rasmusen

Published: *Rationality and Society* (January 1992) 4: 83-94.

Abstract

Game theory has been criticized as neglecting key aspects of individual behavior and as relying too heavily on special assumptions. It can, in fact, handle individual heterogeneity if the modeler is willing to carefully specify how people are different, but to the extent that such things as heterogeneity and culture are important, the desire for a single unified model is impossible to satisfy. At the same time, game theory's approach is very useful for building specialized models.

Draft: 2.1 (Draft 1.1, June 1990)

2000: Eric Rasmusen, Professor of Business Economics and Public Policy and Sanjay Subhedar Faculty Fellow, Indiana University, Kelley School of Business, BU 456, 1309 E 10th Street, Bloomington, Indiana, 47405-1701. Office: (812) 855-9219. Fax: 812-855-3354. Erasmuse@indiana.edu. Php.indiana.edu/∼erasmuse.

A theme common to the three criticisms of game theory in this issue of *Rationality and Society* is that it is a constricted theory, relying too heavily on narrow assumptions to be useful in explaining how the world works. Small changes in assumptions produce big changes in conclusions, and, in particular, game theory has trouble with the heterogeneity of human beings—their differences in tastes, information, and culture.

The game theorist's reply must be, I think, that game theory's building blocks do apply to a wide variety of situations, and are, in fact, particularly well-suited to heterogeneity, but that the situations being modelled do not lend themselves to general theories. To the extent that heterogeneity is important, models ought to be narrow, for a general model is effectively saying that situations are homogeneous. If situations apparently similar are actually different, then tailoring the game to the facts is better than relying on a single game for all facts. In a good theory, at least some seemingly minor details make a difference, for how else are we to discover that anything but the obvious is important?

To make this concrete, let us follow Gordon Tullock's good example and analyze a particular game. Tullock uses the following matrix:

In its stark form, this matrix conjures up players in a laboratory, whose behavior, while analyzable, has no intrinsic interest for us. But to call the properties of a bare mathematical matrix unrealistic is not a well-posed statement; to build a model, something must exist to be modelled, something whose essence is to be captured. Starting with just the matrix, we are in the position of the RAND game theorists in the early 1950's who were perplexed by a certain two-by-two matrix that generated perverse results. Albert Tucker, on being asked to give a talk on game theory to the Stanford psychology department, decided to attach a story to the numbers. The result was the Prisoner's Dilemma, and a deeper understanding than the mathematics alone could give (Straffin, 1980).

The Tullock matrix is not a prisoner's dilemma, but a story can nonetheless be attached to it. Let us try discussing it as a model of the conflict between offense and defense in a football game.[1] The offensive team is trying

[1] If your response is: "He's using a game to model a game! ", mentally substitute a

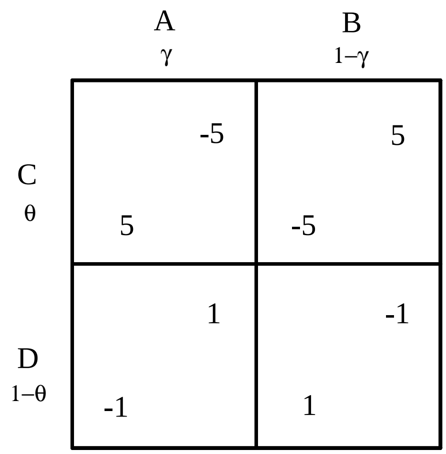

Figure 1: The Tullock Game

Figure 1:

to decide between passing the ball and running with it, and the opposing team must decide whether to set up a defense against passing or running. The offensive team would like to do the unexpected, since that is the way to advance the ball towards the goal. But more is at stake with passing than with running—5 yards instead of 1, if the numbers above are retained.[2] Let us call this "The Football Game." Its Nash equilibrium is in mixed (random) strategies. The Offense passes with probability 1/6, and the Defense uses a running defense with probability 1/2.[3]

---

military conflict between offense and defense where the attack can come on the right or left flank.

[2]These numbers are not quite realistic, though their ordinal rank fits football. Partly this is a matter of the units of measurement. If the numbers in the matrix are doubled and one yard is added to each entry, they become more realistic. Such a change will not affect the optimal strategies in the slightest— it is just a change of measurement units.

[3]To calculate these, use the fact that in a mixed strategy equilibrium a player is indifferent between his pure strategies. The expected payoffs for the Offense from his two pure strategies are $\pi_o(pass) = 5\gamma - 5(1-\gamma)$ and $\pi_o(run) = -1\gamma + (1)(1-\gamma)$. Equating these yields $\gamma = .5$. Similarly, equating $\pi_d(running) = -5\theta + (1)(1-\theta)$ and $\pi_d(passing) = 5\theta + (-1)(1-\theta)$ yields $\theta = 1/6$.

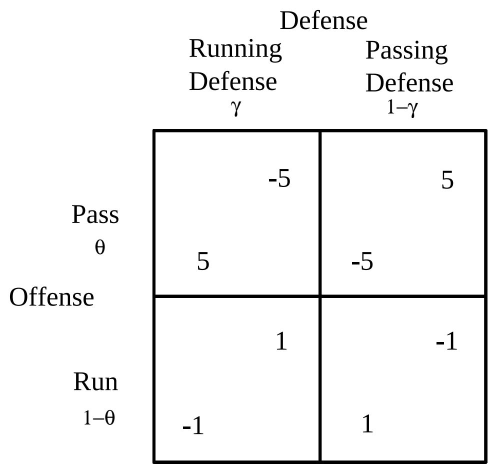

Figure 2: The Football Game

Figure 2:

After discussing the characteristics of the original matrix, Tullock objects that the players should be, like most people, risk-averse. Gaining five yards may not please as much as losing five yards displeases— at least if one's team is currently ahead in the game. This means that the numbers in Fiure 2 are no longer valid, for they no longer represent the payoffs—utility— but an instrumental means to those payoffs—yards. The obvious solution is to change the numbers to fit risk-averse payoffs. But how can this be done? It requires knowledge of how much each player values winning the overall game, and how much the gain and losses of yardage affect winning.

The difficulty of measuring possible payoffs is a common criticism of game theory, but empirical work in any subject runs into measurement problems, and game theory presents no special difficulties—only the standard, hard, ones. The same problem faces the economist who is asked how many more oranges John will buy if the price of apples rises. About both fruit and football, he can make an informed guess using general knowledge and ten minute's thought, or he can use a government grant and three years' work to come up with a somewhat better answer. In the absence of a grant (and

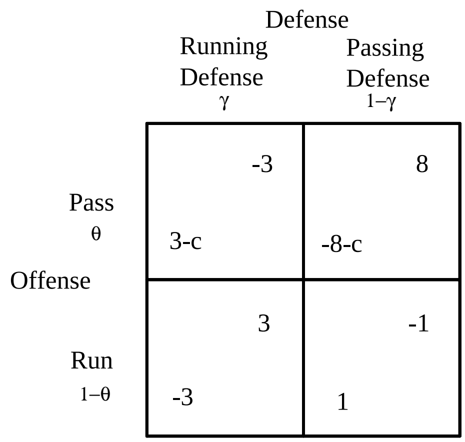

Figure 3: The Football Game with Risk Aversion and Jealousy

Figure 3:

three years), let us make a guess, based on the assumption that the Offensive player is already ahead on points, and replace our earlier matrix with Figure 3 (I have also added variable c, which will be explained shortly):

The game is no longer zero-sum, but the Nash equilibrium can still be calculated. It will depend on the term c, which is added to answer another of Tullock's questions: what if the row player does not like playing his upper action? For the Football Game, the question can be made more concrete: what if the Offensive team captain becomes jealous when his pass receivers get attention from the crowd, and suffers disutility pangs of c when he uses the passing strategy? Jealousy is easy to incorporate; that is what c depicts. The mixed-strategy equilibrium is for the Offence to pass with probability $\theta = 4/15$ and the Defence to use a running defence with probability $\gamma = c/15+3/5$.[4] This yields the interesting prediction that if the Offensive captain

[4] To calculate this, use the fact that in a mixed strategy equilibrium a player is indifferent between his pure strategies. The expected payoffs for the Offensive player from his two pure strategies are $\pi_o(pass) = (3-c)\gamma + (-8-c)(1-\gamma)$ and $\pi_o(run) = -3\gamma + (1)(1-\gamma)$. Equating these yields $\gamma = c/15 + 3/5$. Similarly, equating $\pi_d(running\ defence) = -3\theta + (3)(1-\theta)$ and $\pi_d(passing\ defence) = 8\theta + (-1)(1-\theta)$ yields $\theta = 4/15$.

comes to hate his receivers more (c increases), he will not pass any the less, but the Defence will use the passing defence less often, knowing that the Offense is reluctant to take advantage of the opportunity. Thus, not only can risk aversion and preferences for certain actions be incorporated into the game, but the model can make predictions about how behavior varies depending on those parameters.

I have already mentioned that measurement is difficult for the modeller, but it is is also difficult for the players, a distinctly different objection. So far the Football Game has assumed that they know each others' utility functions. Is this justified? On a general level, it seems that individuals do act as if they know something about each others' utility functions. When a businessman opens up a shoe factory, it is under the belief that he knows pretty well how much other people value shoes— enough so that in tight competition, he can still make a normal profit. He may be mistaken, but he is willing to bet on his knowledge, and since the information is very important to him, he has invested some effort in acquiring information (even more, perhaps, than a scholar writing an article on the subject). So the question is really what happens if the players know each others' payoff functions, but imperfectly.

A complete model requires specification of what it is that the players know and do not know. In the Football Game, let us see what happens if the Offensive captain knows how jealous he is, but the Defensive captain does not. Assume that (a) the Defensive captain does not know the exact value of c, and attaches equal probabilities to $c = 0$ and $c = 1$; (b) the Offensive captain knows that the Defensive captain has those beliefs; and (c) in actual fact, $c = 1$.[5] In equilibrium, the Offensive player will pass with probability $\theta_{c=1} = 0$ and the Defensive player uses a running defence with probability $\gamma = 3/5$. The equilibrium also specifies the behavior that the Offensive player would adopt if $c = 0$; it is to pass with probability $\theta_{c=0} = 8/15$.[6] From the point of view of the Defensive player, who does

[5]If the Offensive captain is uncertain about the Defensive captain's beliefs, that too can be incorporated into the model, at the cost of extra complexity. If he is certain but wrong, it is even simpler to incorporate.

[6]The equilibrium is calculated as follows. Whatever value of γ is picked, it cannot be the case that the Offensive player would mix for both $c = 0$ and $c = 1$; in one case or the other he will use a pure strategy. In computing the equilibrium with known c, it was established that if the value of θ was greater than $4/15$, then $\gamma = 1$. Hence, it cannot be

not know the value of c, the probability of the Offensive player passing is $\overline{\theta} = .5(0) + .5(8/15) = 4/15$. The change in assumptions has changed the Offensive player's behavior, from $\theta = 4/15$ to $\theta_{c=1} = 0$. But this what we should expect: the Offensive player should pass less often when facing an opponent who overestimates his incentive to pass. Thus, game theory easily adapts to differences in tastes and knowledge, but by the very fact that it does so, predicting different outcomes under different circumstances, it tells us that a perfectly general theory is impossible.

Even if one accepts that game theory can accommodate different sorts of utility functions and beliefs, however, one might still quarrel with the very idea of the mixed-strategy equilibrium— as, indeed, Tullock does. Not only do mixed strategies involve "carefully random" behavior, but the equilibrium is weak in the sense that each player is indifferent between at least two of his pure strategies or he would not be willing to mix between them. Yet if the equilibrium is to exist, it seems this indifferent player must carefully pick just the right probabilities for each action.

A mixed-strategy equilibrium does not, however, actually require any randomization. What it requires is that the mixing player's actions *seem* random to the other players, whether this results from literal randomization or not. A football captain might not throw dice in the huddle , but he surely wishes to take actions unpredictable to the other team. If the captain chooses his plays based on a deterministic device such as whether the time remaining in the game is an odd or even number, the game theorist is justified in modelling the choice as random if it seems random to the Defense. The captain may even decide in advance to pass on the first play of the game with probability one, but if the other team believes that in general only proportion θ of captains pass on the first play, the result is the mixed-strategy equilibrium.

The Harsanyi (1973) explanation for mixed strategies cited by Tullock is

that either $\theta_{c=0} = 1$ or $\theta_{c=1} = 1$, or the average would exceed 1/2 and we would have $\gamma = 1$. Since it is the player with $c = 1$ who is more reluctant to pass, it must be that $\theta_{c=1} = 0$, and $\theta_{c=1} \in (0, 1)$. If γ is to be between zero and one, it must be that the average value of θ equals 4/15, so $4/15 = .5(\theta_{c=1}) + .5(0)$, which yields $\theta_{c=1} = 8/15$. Since γ must be chosen to make the player with $c = 0$ willing to mix, it is chosen using the formula $\gamma = c/15 + 3/5$ found in the earlier equilibrium, which since $c = 0$ gives $\gamma = 3/5$.

similar in flavor, but based on player heterogeneity. Harsanyi suggests that there are always small features of the situation which would push the deciding player to one or the other pure strategy, but which cannot be observed by the other players. The Football Game with unknown c hints at this, because what to the Defensive player seems randomization with a 4/15 probability is actually probability 8/15 for one type of Offensive player ($c = 0$) and probability 1 for the other type ($c = 1$). If there were a continuum of types from $c = 0$ to $c = 2$, practically all types would be using pure strategies; but to the Defensive player, who cannot observe c, it would seem that the Offensive player was randomizing. The differences in types need not even be large; there could be a continuum of types from $c = .98$ to $c = 1.02$, and while they would all be choosing one pure strategy or the other,they would appear to the Defensive player (and the modeller) be randomizing.

It should be kept in mind, too, that a mixed-strategy equilibrium is still a Nash equilibrium: no player can profitably deviate from his assigned strategy, even though he may be indifferent about it. This is illustrated by Tullock's other example, the Hillman-Samet rent-seeking game. In that game, two players simultaneously offer bribes to an official who will grant 100 dollars to the highest bidder, but keep both bribes. In the mixed-strategy equilibrium, the two players randomize their choices of bribes between 0 and 100. Tullock asks what happens if the game is repeated, and a third player enters and bids 90 each time. This new strategy will do very badly. If the original two players fail to react, they too will do badly, but that does not make the "Bid 90" strategy rational unless the entrant is malicious. If the original players do react, they can achieve higher payoffs than the entrant by mixing between bids of 0 and bids on the interval between 90 and 100.[7] They will not be driven from the game; often each will bid 0, sometimes one will bid high and win, and sometimes both will bid high and one will lose his bid without reward. The example only shows that rational players should adapt

[7] Let $F(x)$ be the cumulative probability that a player bids up to x. In equilibrium, the payoffs from the pure strategies are equal, and the pure strategies are postulated to be 0 and the $[90, 100]$ interval.. The payoff from bidding 0 is 0, and the probability of winning with a bid of $x > 90$ is the probability $F(x)$ that the other rational player has bid less than x, so $\pi(0) = 0 = \pi(x) = -x + 100F(x)$, and $F(x) = x/100$. This implies that a player bids 0 with probability .9, and spreads the rest of his probability over the interval $[90, 100]$. The "Bid 90" player will win with probability $.9^2$, so his expected payoff is $.81(100) - 90 = -8.9$.

their behavior to the behavior of irrational ones, not that irrational ones do better.[8]

I have spent so much time on the Football Game because it illustrates the sort of sensitivity analysis that is useful for honing intuition. Where changing the assumptions makes a difference, it should make a difference, and shows that apparently minor assumptions are not so minor. This approach places the burden on the modeller to describe the game carefully, but that burden is inescapable under any method of analyis; a car that moves along fine on a solid highway will spin its wheels uselessly when you try to drive it on sand.

Let me turn now to the other two papers. Wildavsky says that game theory ignores culture, and that culture is important, because "without a supportive cultural context, no strategy makes sense." The rational man of game theory is not the reasonable man of law; he is a sociopath, without preexisting values or relations to defend. In particular, the prisoner's dilemma is not nearly as useful as has been claimed, since most people are actually not sociopaths. People in different cultures play the prisoner's dilemma differently, and a given situation will be a prisoner's dilemma in one culture but not in another.

The issue comes down to what is similar about humans and what is different. The position of game theory is that everyone, whatever their culture, is best analyzed as a rational maximizer, but what is rational depends on the particular preferences and constraints available to the individual. To quote O'Rourke (1988, p. 4), "A Japanese raised in Riyadh would be an Arab. A Zulu raised in New Rochelle would be a dentist." Japanese, Arabs, Zulus, and dentists are all rational actors, even if we who are outside their particular cultures do not share their tastes and beliefs.

As Wildavsky says, one physical situation might call for different models for different cultures, or even within one culture, if preferences differ. The Football Game's variants— the basic game and the variants with risk aversion, jealousy, and asymmetric information— all were based on the same physical situation. Game theory does treat people as if they had preexisting

[8]In other contexts, however, such as bargaining, irrationality can be a positive advantage. See Rasmusen, 1989, Chapter 10.)

values; the payoffs are literally values, and the model takes them as preexisting. Wildavsky hits the point precisely when he says the modeler's position should be "Tell me how individuals understand their situation in terms of their preferences, and I will then model the strategic aspects of this situation..." Once provided with the rules of the game—players, payoffs, actions, and initial beliefs— game theory predicts the outcome. If the rules of the game are those of the prisoner's dilemma, the outcome is that both players will confess, whatever their culture may be. But if people in the culture enjoy prison, or prefer confession to lying, then the payoffs are different and the game is no longer a prisoner's dilemma. This is not a matter of how the game is analyzed, but of how its rules are specified.

Culture theorists may object to the entire idea that people respond to incentives, believing instead that their responses are preconditioned. This certainly makes for generality, since it leads to predictions independent of the parameters of the particular situation. But it seems much like an extreme form of the Tullock objection that a player may not enjoy playing one of his actions, which, as seen above, is easily handled by game theory. The objection must then become that the theory is unimportant, because the answer is determined by the empirical measurement of the disutility of the action. That is unobjectionable; theory and empirical investigation each have their place.

In some contexts, however, one must wonder whether "culture" really just refers to situations with different payoff matrices. Wildavsky's cultural groups—hierarchists who trust authorities, individualists who trust market exchange, egalitarians who trust voluntary groups, and fatalists who trust no one— might be identical human beings facing different incentives. Consider, for example, a game in which Smith must decide whether to buy a computer from Jones, and Jones must decide whether to sell him a working computer or a broken one. Smith will be a fatalist and not buy the computer if there is no enforcement mechanism penalizing Jones for selling a broken one. Smith will be a hierarchist, and buy, if Jones has a boss who will punish him for selling broken equipment. Smith will be an individualist, and buy, if warranty law would force Jones to replace a defective computer. Smith will be an egalitarian, and buy, if he and Jones are members of the same university department and Jones's reputation will be ruined if he defrauds Smith.

Although these are not all prisoner's dilemmas, and each story changes the game, the prisoner's dilemma is the basic building block for them all. One must only remember that a house requires more than just the basic concrete blocks.

Finally, I turn to Michael Hechter's discussion of the repeated prisoner's dilemma. Essentially, he has two objections: that the game has multiple equilibria and that cooperation is difficult to achieve under incomplete information. The first objection is quite valid, and often forgotten. The Folk Theorem tells us that a huge variety of outcomes can occur in equilibria of the infinitely repeated prisoner's dilemma, including *Cooperate* each period and *Fink* each period. Arguments have been put forward for why cooperation might be more likely, but the debate is still very much alive.[9]

That cooperation is more difficult when information is incomplete, however, is dubious. Incomplete information, in fact, is the most widely accepted explanation for why cooperation might ensue. Kreps, Milgrom, Roberts, & Wilson (1982) show that even if the prisoner's dilemma is only finitely repeated (and *a fortiori* if it is infinitely repeated), the unique equilibrium outcome can be cooperation until close to the end of the game. The extra assumption they add is a small probability that one player (let us call him Smith) is not rational— Smith plays the tit-for-tat strategy whether it helps him or not. The rational player knows of this possibility, and so he cooperates till near the end of the game. If Smith is indeed irrational, the other player need not fear an unprovoked confession; but even if Smith is rational, it is to his advantage to pretend to be irrational till near the end of the game. The argument is subtle, but the conclusion is simple: uncertainty over the player's payoffs can make cooperation more likely, not less.[10] Under

[9] Technical notes, however: (1) The strategy of tit-for-tat is *not* a credible (perfect) equilibrium strategy, because there is insufficient incentive to punish a *confess* deviation with *Confess* the next period. See Kalai, Samet & Stanford (1988). (2) Taylor (1990) is cited as claiming that for large but not too large discount rate there is a unique efficient equilibrium. In fact, (Always *Confess*) remains an equilibrium outcome for any discount rate for which (Always *Cooperate*) is an equilibrium outcome. (3) The repeated prisoner's dilemma is not a game of perfect information. It includes simultaneous moves, so information sets are not singletons. See Rasmusen (1989), Chapter 2. (4) Discounting hurts for cooperation, rather than enhancing it.

[10] For elaboration, see Chapter 5 of Rasmusen (1989) .

complete information, cooperation is difficult, so muddying the waters can hardly hurt.

Hechter and Wildavsky would probably agree with Tullock when he says "My objection to game theory is as a formal body of mathematics allegedly applying to human action, not as a heuristic which makes it easier to think about certain problems." Game theory should indeed be a branch of storytelling, not of mathematics— storytelling with the i's dotted and the t's crossed. One can tell stories that are consistent, and stories that are not; and formalism helps spot the inconsistencies by forcing the storyteller to tell one story at a time. When Tullock says that "The problem is that parties with no knowledge of formal game theory are likely to go through the same process as the trained game theorist," he pays a complement to formal game theory, for economists have learned only slowly that participants in the marketplace are often wiser than scholars. Long experience, inherited tradition, or the careful deliberation that self-interest motivates often lead to behavior whose usefulness is not apparent to the outsider. But theory has an advantage similar to that of the factory worker over the craftsman: the theorist can make do without experience, tradition, and self-interest, and use superior capital to produce a product that is cheaper and more uniform, if perhaps not so reliable. Let us not be Luddites; without theoretical tools, only the born craftsman is able to produce decent scholarship.

A criticism common to all three critics is that game theory relies too heavily on situation-specific assumptions. We all would prefer generalizing theory to exemplifying theory, a model of what *must* happen instead of what *can* happen, as Franklin Fisher (1989) puts it in his own critique of game theory. The danger in this is that the modeler may try to force-fit a model to situations for which it is unsuited, as Wildavsky says is done with the prisoner's dilemma. But sociology, political science, and anthropology are in a wonderful position to escape this danger. These disciplines have been heavily data-driven, with many descriptive studies of particular situations, which is the empirical analog of the exemplifying style of game theory. Game theory can explain not only what happened in a case study, but what might have happened had conditions been different and what will happen if the parameters change. If game theory is used in this way, I think that the critics will be happier.

REFERENCES

Fisher, Franklin. 1989. Games economists play: a noncooperative view. *Rand Journal of Economics* 20: 113-124 .

Harsanyi, John. 1973. Games with randomly disturbed payoffs: a new rationale for mixed strategy equilibrium points. *International Journal of Game Theory.* 2: 1-23.

Hillman, A. and D. Samet. 1987. Dissipation of contestable rents by small numbers of contenders. *Public Choice* 54: 63-82.

Kalai, Ehud, Dov Samet and William Stanford. 1988. Note on reactive equilibria in the discounted prisoner's dilemma and associated games. *International Journal of Game Theory* 17: 177-186.

Kreps, David, Paul Milgrom, John Roberts, & Robert Wilson. 1982. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory* 27: 245-52.

O'Rourke, P.J. 1988. *Holidays in hell.* New York: Atlantic Monthly Press.

Rasmusen, Eric. 1989. *Games and information.* Oxford: Basil Blackwell.

Rasmusen, Eric. Forthcoming. Folk theorems for the observable implications of repeated games. *Theory and Decision.*

Straffin, Philip. 1980. The prisoner's dilemma. *UMAP Journal* 1: 101-3.