Evolutionary Stability in the Traveler's Dilemma

Andrew T. Barker

Andrew T. Barker (andrew.barker@colorado.edu) is a Ph.D. student in Applied Mathematics at the University of Colorado, Boulder. His main research interests are in numerical analysis and parallel computation, with a current focus on the simulation of blood flow in human arteries. He can often be found drinking coffee, reading, or hiking the Colorado mountains, and sometimes doing all three at once.

The Traveler's Dilemma is something of a paradox in game theory. All of the theory points to a solution that is intuitively wrong, and intuition leads to strategies that analysis can definitively reject. In fact, the paper in which Kaushik Basu introduces the Traveler's Dilemma is subtitled "Paradoxes of Rationality in Game Theory" [**2**].

In experiments, human subjects faced with the Traveler's Dilemma do not act in the theoretically rational way, even if they are professional game theorists [**4, 7**]. It is easy to claim simply that humans are not rational until you realize that their irrational play leads to much higher cumulative payoffs than the alleged rational solution. In [**3**], Basu writes "There is something rational about choosing not to be rational when playing Traveler's Dilemma." In this note, we explore what that rational something is.

The Traveler's Dilemma and rational strategies

The best description of the game is in [**2**]; we give a brief overview here. The setup for the Traveler's Dilemma (TD) is two players returning from some tourist destination with identical souvenirs, both of which have been destroyed by inept baggage handlers. An airline representative, with no idea what the odd souvenir is worth, has each traveler write down its value (from 2 to 100 dollars) on a piece of paper. The representative assumes the lower value is more accurate, giving the player who writes it down that amount plus two dollars (as a reward for being honest) and the other player the lower number minus two dollars (as punishment for trying to scam the airline). If the two players write down the same number, they both receive that amount.

The natural strategy in TD is to pick a high number—after all, you can't realize a large payoff unless you write down a large number. The theoretical tools of game theory, however, reach a strikingly different answer: the only rational move is to write down 2, the lowest number that is allowed by the rules of the game. This conclusion is strange enough that we illustrate it in three ways, by discussing dominated and undominated strategies, by bringing in the concept of a Nash equilibrium, and by relating TD to the better known and closely related Prisoner's Dilemma.

A central assumption of game theory is that both players are rational, and that they both know that their opponent is rational. This means that if you're considering writing down 100, you will realize that playing 99 gives an equal or better payoff regardless of your opponent's strategy. So a rational player rejects the 100 strategy, and has to realize that their (rational) opponent will also reject it. Then the same logic applies to 99—it gets beat by 98. And so on. The race to the bottom ends only when the rules of the game force it to end, at the strategy 2. In the language of game theory, strategy j is weakly dominated by strategy $j-1$ once the higher strategies have been eliminated, which is to say that strategy $j-1$ never gets a smaller payoff than strategy j. The only undominated strategy in TD is 2.

Approaching the problem from the standpoint of the Nash equilibrium is very much the same. A Nash equilibrium is a strategy such that if both players choose it, neither can improve their outcome by changing strategy unilaterally. If player one chooses strategy j, player two wants to unilaterally change to strategy $j-1$. Then player one wants to switch to $j-2$. Again, this race ends at the bottom with the strategy 2. If both players write down 2, neither has incentive to switch unilaterally, so 2 is the only Nash equilibrium.

As a final way of illustrating the counterintuitive conclusion that 2 is the best rational strategy, we connect TD to the Prisoner's Dilemma. If the maximum damage you are allowed to claim is changed from 100 dollars to 3 dollars, then TD becomes exactly the Prisoner's Dilemma: claiming 2 dollars of damage is choosing "defect" and claiming 3 is choosing "cooperate". Regardless of your opponent's move, your payoff is higher if you defect than if you cooperate, so defection is the best strategy, both in theory and intuitively. The theory doesn't change when the maximum claim is increased from 3 to 100, but intuitively the game changes a great deal. Somehow the existing definitions in game theory can't "see" the difference between 3 and 100. Here we develop tools that do see this difference.

In what follows we will need to discuss more precisely the interaction between strategies and payoffs. To do so, we represent strategies as vectors and the game itself as a matrix, called the payoff matrix. Much of the notation below is borrowed from [**6**].

We index matrices and vectors starting with 2, which we hope will relieve more confusion than it causes. The entry u_{ij} of the payoff matrix U is given by

$$u_{ij} = \begin{cases} i+2 & i < j, \\ j-2 & j < i, \\ i & i = j. \end{cases}$$

In this matrix, the entry u_{ij} corresponds to the payoff a player with strategy i receives when faced with an opponent choosing strategy j. Let N be the maximum strategy allowed in the TD; in the canonical TD, $N = 100$, and in Prisoner's Dilemma, $N = 3$. Then the payoff matrix for TD is $U \in \mathbb{R}^{n\times n}$ where $n = N - 1$ (strategies from 2 to N are allowed, inclusive). For $N = 6$, for example, the payoff matrix looks like

$$U = \begin{pmatrix} 2 & 4 & 4 & 4 & 4 \\ 0 & 3 & 5 & 5 & 5 \\ 0 & 1 & 4 & 6 & 6 \\ 0 & 1 & 2 & 5 & 7 \\ 0 & 1 & 2 & 3 & 6 \end{pmatrix}.$$

So far we have only considered pure strategies, but the game can be played with mixed strategies, where, for example, you play 2 with probability $1/2$ and 4 with probability $1/2$. We can represent strategies of this kind as vectors in $\mathbb{R}^n$ whose entries add up to 1. Then the strategy just considered can be represented as $(\frac{1}{2}\ 0\ \frac{1}{2}\ 0\ 0 \cdots)^T$. In this framework a pure strategy is a unit basis vector. Then we can write the expected payoff for playing strategy p against strategy q as

$$W(p,q) = p^T U q. \tag{1}$$

Using this notation, a strategy p is a Nash equilibrium if

$$W(q, p) \le W(p, p)$$

for all $q \neq p$. If the inequality is strict, we call p a strict Nash equilibrium.

The standard definitions in game theory are binary, which is to say that a strategy in a game is either a Nash equilibrium or it is not. In order to understand TD, we need a quantitative measure of how stable an equilibrium is, a measure with a range of possible values. We can develop such a measure using evolutionary game theory and the notion of an Evolutionarily Stable Strategy (ESS), a concept developed by John Maynard Smith and George Price [**8, 9**].

Enter evolutionary game theory

Evolutionary game theory is an extension of game theory where we interpret payoff as evolutionary fitness and a strategy vector as a population. Now the strategy vector $(\frac{1}{2}\ 0\ \frac{1}{2}\ 0\ 0 \cdots)^T$ represents a population where half the members use strategy 2 and half use strategy 4; this interpretation is equivalent to the probability interpretation earlier. What is different is that the strategy changes in response to evolutionary pressure; the strategies with higher payoff produce more offspring, and so the population changes over time.

If almost all $(1 - \epsilon)$ of a population uses strategy p, and it cannot be successfully invaded by a small number (ϵ) of competitors using any possible alternative strategy, then we call p an ESS. More formally, a strategy p is said to be an ESS if there exists $0 < \epsilon < 1$ such that

$$W\big(q, \epsilon q + (1 - \epsilon)p\big) < W\big(p, \epsilon q + (1 - \epsilon)p\big) \tag{2}$$

for all $q \neq p$. If (2) holds for some ϵ_0, then it also holds for all $\epsilon \le \epsilon_0$, which is to say that p is robust to invasion by small groups of competitors. This definition says nothing, however, about invasion by a large number.

Let $q_2 = (1, 0, 0, \cdots)^T$. This represents the pure strategy of choosing 2 all the time, which we have seen is the only Nash equilibrium in TD. This is true even if mixed strategies are allowed, and it is easy to see that this strategy is also the only ESS. But our concern is with *how* stable the equilibrium is, not a binary question of whether it is stable or not. Let $\mu(p, q)$ be the largest ϵ for which (2) holds, and define $\mu(p, q) = 0$ if there is no such $\epsilon > 0$, that is, if p is not an ESS. Then $\mu(p, q)$ is a measure of how stable p is with respect to invasion by q. In general, $\mu(p, q)$ is what proportion of the population must be invaders for the invaders to displace the original population. If $\mu(p, q) = 1$, then p cannot be invaded by even a large force employing strategy q; if $\mu(p, q) = 0$, then even an infinitesimally small force can overcome the population. Of course, if $\mu(p, q) > 1/2$ then the division into "invaders" and "orginal population" makes little sense, and in this case we can consider p to be very stable. But if $\mu(p, q)$ is a small number greater than zero, p can be invaded fairly easily, even if it is an ESS.

More generally, consider the fitness of q_2 with respect to a potential invader with pure strategy q_r, where the rth component of q_r is 1 and all other components are 0. (Allowing q_r to be a more general strategy does not change the results, but makes the argument slightly more complex.) We measure the fitness as $\mu(q_2, q_r)$, which is the largest ϵ so that

$$W(q_r, \epsilon q_r + (1 - \epsilon)q_2) < W(q_2, \epsilon q_r + (1 - \epsilon)q_2).$$

We have

$$W(q_r, \epsilon q_r + (1-\epsilon)q_2) = \epsilon W(q_r, q_r) + (1-\epsilon)W(q_r, q_2) = \epsilon W(q_r, q_r) = \epsilon r$$

and

$$W(q_2, \epsilon q_r + (1-\epsilon)q_2) = \epsilon W(q_2, q_r) + (1-\epsilon)W(q_2, q_2) = 4\epsilon + (1-\epsilon)2.$$

This means that

$$\mu(q_2, q_r) = \min\left\{1, \frac{2}{r-2}\right\}.$$

This function is plotted in Figure 1. Note how rapidly the fitness decays, suggesting that q_2 is not very stable with respect to large values of r.

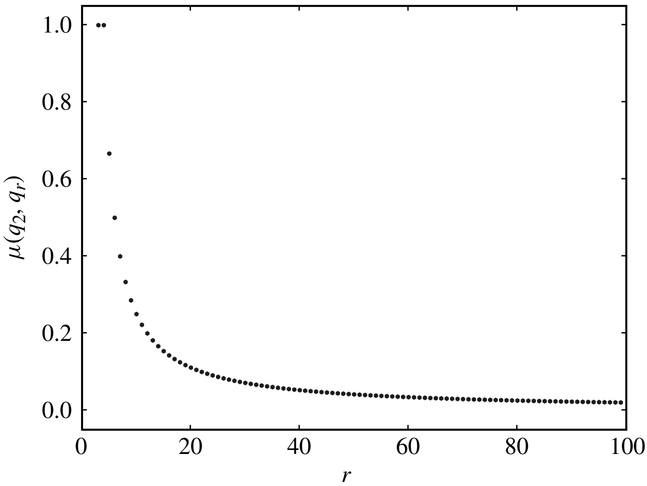

Figure 1. Relative fitness $\mu(q_2, q_r)$ of the ESS compared to other pure strategies. The strategy q_2 is perfectly stable with respect to strategies q_3 and q_4, but rapidly becomes less stable against the higher-valued strategies, which could invade the ESS if introduced in sufficient numbers.

How stable is a strategy?

Let

$$\mu_0(p) = \inf_{q \neq p} \mu(p, q).$$

Then $\mu_0(p)$ is a quantitative measure of how stable a strategy is, not only against a single opponent q but with respect to all possible competing strategies. Note that $\mu_0(p) > 0$ if and only if p is an ESS. Since $\mu(q_2, q_r)$ is a monotone decreasing function of r and contains no explicit dependence on N, we can interpret Figure 1 as a plot of $\mu_0(q_2)$ versus N; that is

$$\mu_0(q_2) = \min\left\{1, \frac{2}{N-2}\right\}.$$

For the canonical TD, with $N = 100$, $\mu_0(q_2) = 1/49$. In contrast, for the Prisoner's Dilemma with $N = 3$, $\mu_0(q_2) = 1$.

As we saw, the two games are nearly equivalent in traditional game theory, but our stability measure shows a dramatic difference, which begins to explain the paradox of TD. Part of what makes the TD counterintuitive is the high value of N. It is fairly reasonable to claim q_2 as the optimal strategy in Prisoner's Dilemma, but saying the same thing in TD when there are potential payoffs of 100 seems unreasonable. In traditional game theory it doesn't matter if $N = 3$ or 100 or 10^{10}; but now we see that as N goes to infinity, the Nash equilibrium becomes evolutionarily unstable. To complete the picture, we need to consider why higher strategies, strategies that the theory definitively rejects, might be attractive.

Let q_r be the pure strategy which chooses r every time in the TD, and compare its fitness to q_s, which chooses s. We have

$$W\big(q_s, \epsilon q_s + (1-\epsilon)q_r\big) = \epsilon W(q_s, q_s) + (1-\epsilon) W(q_s, q_r)$$

and

$$W\big(q_r, \epsilon q_s + (1-\epsilon)q_r\big) = \epsilon W(q_r, q_s) + (1-\epsilon) W(q_r, q_r).$$

Then we have (details omitted)

$$\mu(q_r, q_s) = \begin{cases} \dfrac{s-(r-2)}{s-r} & s < r-2, \\ 0 & r-2 \leq s \leq r-1, \\ 1 & r+1 \leq s \leq r+2, \\ \dfrac{2}{s-r} & s > r+2. \end{cases}$$

This function is plotted in Figure 2 for $r = 50$. Note the relatively larger region of high fitness for q_{50} as compared to q_2; in fact, the only two population mutants that can invade at infinitessimally small invasion sizes are strategies 48 and 49.

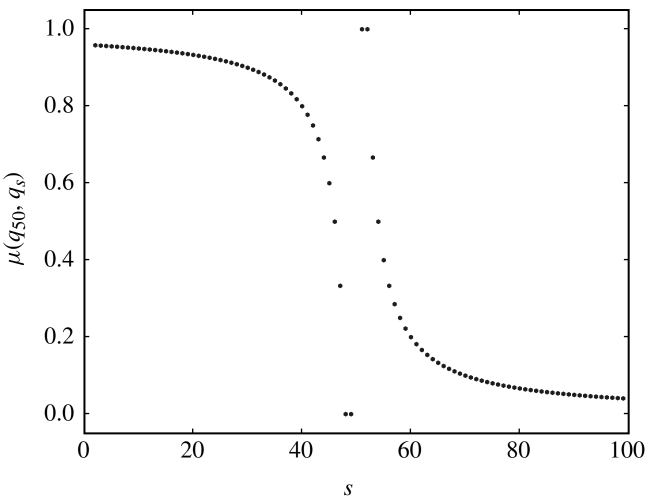

Figure 2. Relative fitness $\mu(q_{50}, q_s)$ with s varying. μ is zero for $s = 48, 49$, so q_{50} is not an ESS, but it is fairly stable with respect to values of $s < 45$ anyway.

The preceding analysis suggests a solution to the paradox of TD. Though q_2 is the only evolutionarily stable strategy, it is in some sense not very stable. An analogy can be made to the concept of a singular matrix. For a certain kind of pure mathematician, a matrix is either singular or it is not; but for those who actually do matrix computations, there is a measure of how singular a matrix is, its condition number. Similarly, a strategy is either evolutionarily stable or not; but it is possible to define a meaningful measure of its stability.

The condition number of a matrix A can be defined as the ratio of its largest eigenvalue to its smallest, and in some sense measures the amplification of error in the solution of linear systems involving A. We refer to a matrix as numerically singular if its condition number is so high that rounding error, or some other kind of noise, makes it effectively impossible to invert. In [**2**], Basu proposes a possible solution to the TD is to use ill-defined categories rather than exact integers, in which case the strategy pair (large, large) is a kind of ill-defined Nash equilibrium. This suggests a rounding error in the categories, where q_{98} is seen as essentially the same as q_{94}. If this rounding error is large enough, the ESS is effectively unstable. Similarly, in [**4**], Becker et. al. show that if even a few (3 percent, or $\mu_0(q_2)$ in our notation) players are known to play "irrationally" high numbers, it becomse rational for all players to do likewise—this is another kind of noise, in this case affecting the mutual rationality assumption of game theory, and this kind of noise also affects the stability of the Nash equilibrium.

A final word

For those interested in evolutionary game theory, the classic texts are [**8**] and [**12**]. Evolutionary ideas have been used very often in the related Prisoner's Dilemma (see [**1**], [**10**], and the helpful review article [**11**]), but rarely for TD. The idea of some kind of noise in evolutionary game theory has been addressed in [**5**], but as a way of choosing among multiple equilibria, not as a way of measuring the stability of an equilibrium. Here we have only considered TD, but it would be interesting to see if measures of evolutionary stability are useful in the analysis of other games.

References

1. R. Axelrod, *The Evolution of Cooperation*, Basic Books, New York, 1984.
2. K. Basu, The traveler's dilemma: paradoxes of rationality in game theory, *Am. Econ. Rev.* **84** (1994) 391–395.
3. ———, The traveler's dilemma, *Sci. Am.*, June 2007.
4. T. Becker, M. Carter, and J. Naeve, Experts playing the traveler's dilemma, Diskussionspapiere aus dem Institut für Volkswirtschaftslehre der Universität Hohenheim 252/2005, Department of Economics, University of Hohenheim, Germany (2005), available at `http://ideas.repec.org/p/hoh/hohdip/252.html`.
5. K. Binmore and L. Samuelson, Muddling through: Noisy equilibrium selection, *J. Econom. Theory* **74** (1997) 235–265.
6. N. F. Britton, *Essential Mathematical Biology*, Springer, 2003.
7. C. M. Capra, J. K. Goeree, R. Gomez, and C. A. Holt, Anomalous behavior in a traveler's dilemma, *Am. Econ. Rev.* **89** (1999) 679–690.
8. J. Maynard Smith, *Evolution and the Theory of Games*, Cambridge, 1982.
9. J. Maynard Smith and G. Price, The logic of animal conflict, *Nature* **246** (1973) 15–18.
10. M. Nowak and K. Sigmund, A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game, *Nature* **364** (1993) 56–58.
11. ———, Evolutionary dynamics of biological games, *Science* **303** (2004) 793–799.
12. J. W. Weibull, *Evolutionary Game Theory*, MIT Press, Cambridge, 1995.

