

LOGIT EQUILIBRIUM MODELS OF ANOMALOUS BEHAVIOR:
WHAT TO DO WHEN THE NASH EQUILIBRIUM SAYS ONE THING
AND THE DATA SAY SOMETHING ELSE

Simon P. Anderson, Jacob K. Goeree, and Charles A. Holt*
May 2002

Every experimentalist will sooner or later come across a situation in which results from initial "baseline" treatments conform nicely to the Nash equilibrium, but subsequent changes in parameters push the data in ways not predicted by Nash. This may happen when one begins by giving theory its "best shot," reserving stress tests for later. Such tests often involve changing a parameter that, on the basis of intuition, is likely to alter behavior, but which has no effect on the Nash equilibrium. For example, behavior in a symmetric matching-pennies game conforms to the Nash prediction of mixing with equal probabilities. However, changing a player's own payoff parameters will typically change that player's choice probabilities (Ochs, 1995; Goeree, Holt, and Palfrey, 2002b), in spite of the fact that in a Nash equilibrium a player's mixed strategy only depends on the *other* players' payoffs. In "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," Goeree and Holt (2001) report a variety of games in which behavior conforms nicely to Nash predictions in one treatment (the "treasures") but deviates sharply from Nash predictions in the other treatment (the "contradictions").

A particularly striking example of the contrast between economic intuition and the cold logic of game theory is the "traveler's dilemma" described by Basu (1994). The dilemma is based on a situation in which two vacationers have purchased identical objects, which are then lost on the flight home. The airline tells them to fill out claim forms independently, with the promise that both claims will be paid if they match. Otherwise, both travelers are only reimbursed at the lower of the claims, with a small penalty for the high claimant and an equally small reward for the low claimant. Even with a very low penalty and reward, each person has an incentive to "undercut" any anticipated common claim amount, and so the only Nash equilibrium (in pure or mixed strategies) is at the lowest possible claim. The implausibility of this prediction becomes apparent when one considers very low values of the penalty and reward parameter.

Capra, Goeree, Gomez, and Holt (1999) report a traveler's dilemma experiment in which claims are required to be between 80 and 200 cents, so the Nash equilibrium involves claims of 80 cents.¹ With the penalty for the high claimant and the reward for the low claimant set at 50 cents, claims converged to near-Nash levels, as indicated by the blue bars in Figure 1. But with a relatively low penalty-reward level of 10 cents, the frequency distribution of claims (represented by the green bars) is much higher. Note that this distribution is concentrated at the *opposite* side of the feasible set from the Nash

* Department of Economics, 114 Rouss Hall, University of Virginia, Charlottesville, VA 22903-3328. This research was supported by a grant from the National Science Foundation (SES-0094800).

¹ See also Capra, Goeree, Gomez, and Holt (2002) for experiments that involve a similar game based on a model of imperfect price competition.

equilibrium (indicated by the red bar at 80 cents). What is needed is a theory that explains Nash-like behavior in some contexts and deviations in others, i.e. a generalization of the Nash solution concept.

This paper describes such a generalization and how it can be used to generate testable predictions for a variety of economic games. This approach involves introducing random elements, interpreted as either bounded rationality or unobserved preference shocks, into an equilibrium analysis. Individuals' choices are assumed to be positively, but not perfectly, related to expected payoffs, in that decisions with higher expected payoffs are more likely to be selected. With repeated (random) matchings, the choice probabilities of one player will affect the beliefs, and hence the expected payoffs, of others. The *equilibrium* is a fixed point: the choice probabilities that determine expected payoffs correspond to the probabilities determined by expected payoffs via a probabilistic choice rule (Rosenthal, 1989; McKelvey and Palfrey, 1995). The degree of bounded rationality is described by an error parameter, and the equilibrium probabilities converge to a Nash equilibrium as this parameter goes to zero.

We use this approach to explain stylized facts emerging from laboratory tests of game theory. The focus is on games with continuous strategy spaces: e.g. voluntary contributions games, all-pay auctions, continuous coordination games, and the traveler's dilemma game. The goal is to establish comparative statics properties that are consistent with patterns in laboratory data, especially when these patterns are intuitive but not explained by Nash predictions. See Anderson, Goeree, and Holt, 2002, for a formal derivation of existence, uniqueness, and comparative statics results for the logit quantal response equilibrium in games with continuous strategies.

1. BACKGROUND: THE LOGIT APPROACH

The idea of bounded rationality - that agents are limited in their ability to evaluate their environments - has been around for decades (the term was coined by Nobel-prize-winner Herb Simon). Little attempt, however, has been made to use the idea to generate formal equilibrium predictions until recently, when Rosenthal (1989) and McKelvey and Palfrey (1995) incorporated decision error into an analysis of non-cooperative games. To see how this is done, suppose that there are two decisions, D_1 and D_2 , with associated expected payoffs B^e_1 and B^e_2 . Under perfect rationality, the decision with the higher expected payoff is always chosen. Bounded rationality can be modeled by adding a random element in the expected-payoff comparison. In particular, suppose that the probability of choosing D_1 is:

$$(1) \quad \Pr(D_1) = \Pr(\mathbf{p}^e_1 - \mathbf{p}^e_2 > \boldsymbol{\mu}\boldsymbol{\epsilon}),$$

where $\boldsymbol{\epsilon}$ is a random variable with mean zero, and $\boldsymbol{\mu}$ is an error parameter that determines the importance of the error term. If $\boldsymbol{\mu}$ is small enough, the effects of errors vanish and the decision with the highest payoff is always chosen. If $\boldsymbol{\mu}$ is large enough, decisions are equally likely to be selected, irrespective of their expected payoffs. The particular characteristics of the model depend on the distribution of the random component. If $\boldsymbol{\epsilon}$ is uniform on $[-.5, .5]$, for example, $\Pr(D_1) = .5 + (B^e_1 - B^e_2)/2\boldsymbol{\mu}$, unless the expected payoff difference is so great as to push the probability to 0 or 1. This is known as a linear

probability model because choice probabilities are linear in payoff differences, which is the assumption made in Rosenthal (1989).

McKelvey and Palfrey (1995) propose a more general *quantal-response* equilibrium model, in which choice probabilities are increasing, but possibly non-linear functions of expected-payoff differences. Using an exponential function, which corresponds to a logistic distribution of μ , in (1), yields the familiar logit form that is widely used in

empirical work. In this case, the choice probabilities are:

$$(2) \quad \Pr(D_i) = \frac{\exp(\mathbf{p}_i^e / \mathbf{m})}{\exp(\mathbf{p}_1^e / \mathbf{m}) + \exp(\mathbf{p}_2^e / \mathbf{m})}, \quad i = 1, 2,$$

where the denominator ensures that the probabilities sum to one. The continuous analogue to (2) applies when there is a continuum of alternatives, with the probabilities in (2) replaced by a probability density for choosing an action x from a feasible set $[a, b]$:

$$(3) \quad f_i(x) = \frac{\exp(\mathbf{p}_i^e(x) / \mathbf{m})}{\int_a^b \exp(\mathbf{p}_i^e(s) / \mathbf{m}) ds},$$

where the denominator is a constant of integration that is independent of x . This continuous formulation was first used by Lopez (1995). Equation (3) does not provide explicit solutions since the expected payoff on the right side of (3) depends on the choice densities of the other players. Rather, it gives *equilibrium consistency* conditions: the choice densities that determine expected payoffs should match the choice densities determined by expected payoffs via the logit choice rule.

Differentiating both sides of (3) with respect to x establishes that the equilibrium densities satisfy the *logit differential equation*:

$$(4) \quad \mathbf{m} f_i'(x) = f_i(x) \mathbf{p}_i^e'(x),$$

where the prime denotes the derivative with respect to x . When μ goes to zero, equation (4) dictates that $\mathbf{B}_i^e(x) f_i(x) = 0$, which is the condition needed for an interior Nash equilibrium: either the necessary condition for payoff maximization is satisfied at x , or else the density of decisions is zero. As μ goes to infinity in (4), the noise effect dominates and the equilibrium density is uniform with $f_i' = 0$. In the next section we will use (4) to derive the logit equilibrium densities for specific games.

The focus above has been on the equilibrium, and not on some learning process that presumably leads to it. This analysis can be supported by dynamic adjustment models that explain convergence to a quantal response equilibrium, e.g. the fictitious play learning model in Chen, Friedman, and Thisse (1996) and the noisy evolution model in

Anderson, Goeree, and Holt (1998). The equilibrium analysis presented here, however, applies to the long-run situation in which players are familiar with the distributions of others' decisions, e.g. when behavior has stabilized in the later rounds of an experiment.

2. HOW TO FIND A LOGIT EQUILIBRIUM

In this section we will sketch how the differential equation (4) can be used to derive the logit equilibrium densities for specific games. One starts by writing down the expression for a player's expected payoff as a function of that player's own decision and of the other players' densities. Substituting this expected payoff into (4) yields a differential equation for the equilibrium density. For a number of games, this equation admits a closed-form solution.

Example 1 (linear payoffs). Consider a linear public goods game in which individuals are given an endowment, T . When player i contributes an amount x_i to the public good, the player earns $T - x_i$ for the part of the endowment that is kept. In addition, every player receives a constant (positive) fraction m of the total amount contributed to the public good. So the payoff to player i is: $B_i = T - x_i + mX$, where X is the sum of all contributions including those of player i . Since the derivative of expected payoff with respect to x_i is $(m-1)$, equation (4) reduces to $\mu f_i'(x) = (m-1) f_i(x)$ for $x \in [0, T]$. Therefore, f_i'/f_i equals a constant, $(m-1)/\mu$, and the equilibrium density is simply a (truncated) exponential: $f_i(x) = K \exp(x/\mu)$, where K is a constant of integration and $\mu = (m-1)$. Indeed, the logit equilibrium density is exponential with parameter μ for any game in which players' payoffs are linear in own decisions with constant slope μ .

Example 2 (quadratic payoffs). When the expected payoff function is quadratic in one's own decision, the solution to (4) is proportional to an exponential of a quadratic function, i.e. a (truncated) normal density. We use this observation to characterize the equilibrium for public goods games with quadratic payoffs (Anderson, Goeree, and Holt, 1998a). The logit equilibrium density for other quadratic models, such as Cournot competition with linear demand, is also given by a truncated normal. This normality result is important, given the pervasiveness of quadratic payoffs both in theoretical models and experimental work.

Example 3 (all-pay auction). In this game, a prize is awarded to the person making the highest effort or "bid," but all competitors incur the costs of their own bids, whether or not they are successful. Consider a symmetric two-person auction in which the prize is worth V dollars. If $F(x)$ denotes the equilibrium distribution of bids, then the expected payoff for a bid of x is given by $VF(x) - cx$, where c is an effort cost parameter. Hence, $B^e(x) = VF(x) - cx$, and it follows from (4) that the logit equilibrium density satisfies $\mu f'(x) = Vf(x)^2 - cf(x)$. It is straightforward to verify that the solution to this differential equation is: $f(x) = 1/[1 - K \exp(x/\mu)]$, where K is a constant of integration. In Anderson, Goeree, and Holt (1998b) we extend this analysis to cover asymmetric, n -player all-pay auctions.

Example 4 (minimum-effort coordination game). A player's payoff in this game is the minimum of all efforts minus the cost of that player's effort: $B_i = \min_{j=1, \dots, n} \{x_j\} - cx_i$, where the cost parameter c is less than one and $x_i \in [0, \infty)$. Given the linear payoff structure, the Pareto-dominant equilibrium is for each player to provide the maximum effort, ∞ , but *any* other common effort level is also a Nash equilibrium.

Consider the symmetric two-player case, in which each player is characterized by a distribution of efforts, $F(x)$. The probability that the other's effort is below x is then given by $F(x)$ with density $f(x)$. Recall that a player's payoff is the minimum effort minus the cost of the player's own effort:

(5)

$$p^e(x) = \int_0^x y f(y) dy + x(1-F(x)) - cx,$$

where the first term is the benefit when another player's effort is below the player's own effort, x , and the second term is the benefit when the player's own decision determines the minimum effort. Differentiation of (5) establishes that $B^e(x) = (1 - F(x)) - c$, which together with (4) results in the logit differential equation: $\mu f'(x) = f(x) (1 - F(x) - c)$. In Anderson, Goeree, and Holt (2001) we show that the solution is a truncated logistic distribution of the type used in epidemiological models of the transmission of disease over time.

Example 5 (traveler's dilemma). The expected payoff for this game consists of two terms, depending on whether or not the person's claim is lower:

(6)

$$p^e(x) = \int_a^x (y - R)f(y) dy + (x + R)(1 - F(x)),$$

where the first term on the right corresponds to the case where the penalty R is paid, and the second term corresponds to the case where the reward R is obtained. The derivative of expected payoff is thus: $B^e(x) = 1 - F(x) - 2Rf(x)$, so it follows from (4) that $\mu f'(x) = f(x) (1 - F(x) - 2Rf(x))$. In this case there is no closed-form solution, although numerical methods can be used to plot the equilibrium density. For instance, the curved blue and green lines in Figure 1 are plots of the logit density for an error parameter $\mu = 6$ (see Capra, Goeree, Gomez, and Holt, 1999).

These examples demonstrate that the logit equilibrium approach can be applied to a wide variety of interesting economic contexts. Closed-form solutions are available in some cases, but even when they are not, it is often possible to derive comparative static properties of the logit equilibrium. This is the topic of the next section.

3. COMPARATIVE STATIC PROPERTIES

Our research is motivated by experimental evidence that cannot be explained by a standard Nash equilibrium analysis. We are particularly interested in deviations from Nash predictions that are consistent with economic intuition. This section summarizes intuitive comparative statics results from a number of specific applications of the logit equilibrium to games with a continuum of feasible decisions.

Example 1 (linear payoffs). Suppose the expected payoff is linear in one's own contribution. In the standard public goods game, the slope equals $m-1$. This slope is negative when the marginal value of the public good, m , is less than the cost of contribution, 1, and the Nash equilibrium involves zero contributions (free riding). Some contributions are observed in laboratory experiments, however, even after repetition.

These deviations from the Nash prediction are often consistent with loose economic intuition. For instance, average contribution levels in public goods experiments increase with the marginal value of the public good, even though the Nash prediction is unaffected by (non-critical) changes in this parameter. The logit equilibrium provides one explanation of this effect: since $B^e(x) = m-1$, an increase in m will increase the slope of the density in (4), which in equilibrium, raises the average level of contributions. This result is derived formally in Anderson, Goeree, and Holt (1998a).

Of course, there are many other possible explanations of the positive relationship between m and average contributions, e.g. altruism. These explanations need not be mutually exclusive: altruistic concerns for others' earnings can be added to the payoff functions, in which case the total benefit to others from one's own contribution increases with the number of people who enjoy the public good. An increase in the number of participants is then predicted to increase average contributions, a prediction that is also roughly consistent with laboratory evidence, at least for low values of m .

Example 2 (quadratic payoffs). Quadratic payoffs have been introduced in public goods games by using either a declining marginal value of the public good or an increasing marginal cost of contribution. The purpose of introducing non-linear payoffs was to move the Nash equilibrium away from the boundary of the feasible set. Experiments with interior Nash equilibria have yielded average contribution levels that are between the Nash prediction and the midpoint of the feasible set. Recall that the logit equilibrium density for the quadratic case is a truncated normal on $[a,b]$. When the mode of this density is to left or to the right of the midpoint of the feasible set, the symmetry of the normal implies that the mean is between the mode and the midpoint. This is the intuition behind the logit equilibrium prediction that the average contribution is "sandwiched" between Nash and the midpoint of the feasible set (Anderson, Goeree, and Holt, 1998a).

Example 3 (all-pay auction). The winner-take-all nature of these contests makes the outcome sensitive to decision errors. When the bids represent costly efforts (e.g. lobbying) to obtain the prize, rent dissipation is determined by comparing the value of the prize with the expected effort costs. The logit equilibrium model predicts that the extent of rent dissipation increases with the number of players and the cost of bidding, which is not the case for the Nash equilibrium (Anderson, Goeree, and Holt, 1998b). To evaluate comparative statics, consider the form of the expected-payoff function for the case of two players: $VF(x) - cx$. Thus an increase in c lowers the slope of expected payoff and thus the slope of the equilibrium density by (4). This is why an increase in the bid cost causes stochastically lower bids.

Over-dissipation of rents is impossible with fully rational players, since it implies negative expected payoffs, which can be avoided by non-participation. However, over-dissipation is observed in some laboratory experiments (e.g., Davis and Reilly, 1998). Such over-dissipation is predicted for high enough values of n and c in our model with endogenous decision errors.

Example 4 (minimum-effort coordination game). The classic dilemma in coordination games is that better outcomes require higher effort and more risk of coordination failure. Uncertainty about others' actions is a central element of such situations. The coordination game has a continuum of pure-strategy Nash equilibria, which means that the Nash equilibrium has no predictive power. In contrast, the logit

model can be used to determine a unique probability distribution of effort decisions, and so generates testable hypotheses (Anderson, Goeree, and Holt, 1997b). For example, increases in c and n result in lower average efforts. The proof is by contradiction, but the underlying intuition is again that increases in these parameters will lower the slope of the expected payoff, and therefore the slope of the equilibrium density determined by (4). These comparative static results are consistent with some experimental findings of Van Huyck, Battalio, and Beil (1990). These authors conducted laboratory experiments with a minimum-effort structure, with seven effort levels and seven corresponding Pareto-ranked pure-strategy Nash equilibria. The intuition that coordination is more difficult with more players is apparent in the data: behavior approaches the "best" equilibrium with two players, but the "worst" Nash outcome has more drawing power with large numbers of players. Similarly, the experiments confirm that higher effort costs reduce the level of effort.

Example 5 (traveler's dilemma). Recall that the slope of the expected payoff in the symmetric logit equilibrium is: $B^e'(x) = 1 - F(x) - 2Rf(x)$. Hence, an increase in the penalty and reward parameter, R , decreases this slope and therefore the slope of the equilibrium density by (4). This provides the intuition behind the comparative-static result that increases in R result in a stochastic decrease in claims (see Capra, Goeree, Gomez, and Holt, 1999).

The main comparative results discussed in this paper are summarized in Table 1. In each of the models listed on the left, the Nash equilibrium predicts no effect of the exogenous treatment variables, as shown by the zeros in the Nash column. In all cases, the logit equilibrium makes a prediction, positive or negative, that is qualitatively consistent with the laboratory data. (In the case of the all-pay auction, the logit equilibrium can explain observed over-dissipation of rents, but there is no direct evidence on the relationship between the amount of dissipation and effort costs or the number of rent seekers.) The logit equilibrium also explains some data patterns that are not summarized in the table, e.g. the "sandwich" result for quadratic games.

Other explanations have been offered for some, but by no means all, of these empirical anomalies, but the ability of the logit equilibrium to track these patterns is striking. Many aspects of the data do not correspond to the predictions of any one-parameter family of models, but the logit equilibrium predictions are superior to those of the Nash equilibrium that is the workhorse of standard game theory, and indeed, of much of economic theory.

Game	Treatment Variable	Nash Prediction	Logit Prediction	Laboratory Evidence
Linear Public Goods (effect on contributions)	MPCR number of players	0 0	+ +	+ +
Minimum Effort Coordination (effect on effort levels)	effort costs number of players	0 0	- -	- -
All-Pay Auction (effect on rent dissipation)	effort cost number of players	0 0	- -	na na
Traveler's Dilemma (effect on claims)	penalty/reward	0	-	-

REFERENCES

- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998a) "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games," *Journal of Public Economics*, 70(2), 297-323.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998b) "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction," *Journal of Political Economy*, 106(4), 828-853.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1999) "Stochastic Game Theory: Adjustment to Equilibrium Under Noisy Directional Learning," Working Paper, University of Virginia.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2001) "Minimum-Effort Coordination Games: Stochastic Potential and Logit Equilibrium," *Games and Economic Behavior*, 34(2), 177-199.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2002) "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," *Southern Economic Journal*, 69(1), 21-47.
- Basu, Kaushik (1994) "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory," *American Economic Review*, 84(2), 391-395.
- 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), 678-690.
- Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (2002) "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition," *International Economic Review*, 43(3), August 2002, 613-636.
- Chen, Hsiao-Chi, James W. Friedman, and Jacques-François Thisse (1997) "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, 18, 32-54.
- Davis, Douglas D. and Robert J. Reilly (1998) "Do Too Many Cooks Always Spoil The Stew? An Experimental Analysis of Rent-Seeking and the Role of a Strategic Buyer," *Public Choice*, 95, 89-115.
- Goeree, Jacob K. and Charles A. Holt (1999) "Stochastic Game Theory: For Playing Games, Not Just For Doing Theory," *Proceedings of the National Academy of Sciences*, 96, 10564-10567.
- Goeree, Jacob K. and Charles A. Holt (2000) "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," *European Economic Review*, 44, 1079-1089.
- Goeree, Jacob K. and Charles A. Holt (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402-1422.
- Goeree, Jacob K., Charles A. Holt and Susan K. Laury (2002) "Private Costs and Public Benefits: Unraveling the Effects of Altruism and Noisy Behavior," *Journal of Public Economics*, 83(2), 257-278.
- Goeree, Jacob K., Charles A. Holt and Thomas R. Palfrey (2002a) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), 247-272.

- Goeree, Jacob K., Charles A. Holt and Thomas R. Palfrey (2002b) "Risk Averse Behavior in Generalized Matching Pennies Games," Working Paper, University of Virginia.
- Isaac, Mark, and James M. Walker (1994) "Nash as an Organizing Principle in the Voluntary Provision of Public Goods: Experimental Evidence," Working Paper, University of Arizona.
- Lopez, Gladys (1995) *Quantal Response Equilibria for Models of Price Competition*, unpublished Ph.D. dissertation, University of Virginia.
- McKelvey, Richard D., and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- Ochs, Jack (1995) "Games with Unique Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.
- Rosenthal, Robert W. (1989) "A Bounded Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273-292.
- Smith, Vernon L. (1997) "Monetary Rewards and Decision Cost in Experimental Economics: An Extension," Working Paper, University of Arizona.
- Smith, Vernon L., and James M. Walker (1993) "Monetary Rewards and Decision Cost in Experimental Economics," *Economic Inquiry*, 31, 245-261.
- Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1990) "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80, 234-248.

Figure 1. Data and Theoretical Predictions for the Traveler's Dilemma Game

