

THE LOGIT EQUILIBRIUM

A PERSPECTIVE ON INTUITIVE BEHAVIORAL ANOMALIES

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

Department of Economics, Rouss Hall, University of Virginia, Charlottesville, VA 22903

revised, October 2000

Abstract. This paper considers a class of models in which rank-based payoffs are sensitive to "noise" in decision making. Examples include auctions, price-competition, coordination, and location games. Observed laboratory behavior in these games is often responsive to the asymmetric costs associated with deviations from the Nash equilibrium. These payoff-asymmetry effects are incorporated in an approach that introduces noisy behavior via a logit probabilistic choice function. In the resulting logit equilibrium, behavior is characterized by a probability distribution that satisfies a "rational expectations" consistency condition: the beliefs that determine players' expected payoffs match the decision distributions that arise from applying the logit rule to those expected payoffs. We prove existence of a unique, symmetric logit equilibrium and derive comparative statics results. The paper provides a unified perspective on many recent laboratory studies of games in which Nash equilibrium predictions are inconsistent with both intuition and experimental evidence.

JEL Classification: C72, C92.

Keywords: logit equilibrium, probabilistic choice, rank-based payoffs.

I. Introduction

With the possible exception of supply and demand, the Nash equilibrium is the most widely used theoretical construct in economics today. Indeed, almost all developments in some fields like industrial organization are based on a game-theoretic analysis. With Nash as its centerpiece, game theory is becoming more widely applied in other disciplines like law, biology, and political science and it is arguably the closest thing there is to being a general theory of social science, a role envisioned early on by von Neumann and Morgenstern (1944). However,

* This research was funded in part by the National Science Foundation (SBR-9617784) and (SBR-9818683).

many researchers are uneasy about using a strict game-theoretic approach, given the widespread documentation of anomalies observed in laboratory experiments (Kagel and Roth, 1995; Goeree and Holt, 2001). This skepticism is particularly strong in psychology, where experimental methods are central. In relatively non-experimental fields like political science, the opposition to the use of the "rational choice" approach is based in part on doubts about the extreme rationality assumptions that underlie virtually all "formal modeling" of political behavior.¹

This paper presents basic results for a relatively new approach to the analysis of strategic interactions, based on a relaxation of the assumption of perfect rationality used in standard game theory. By introducing some noise into behavior via a logit probabilistic choice function, the sharp-edged best reply functions that intersect at a Nash equilibrium become "smooth." The resulting *logit equilibrium* is essentially at the intersection of stochastic best-response functions (McKelvey and Palfrey, 1995). The comparative statics properties of such a model are analogous to those obtained by shifting smooth demand and supply curves, as opposed to the constant price that results from a supply shift with perfectly elastic demand (or vice versa). Analogously, parameter changes that have no effect on a Nash equilibrium may move logit equilibrium predictions in a smooth and often intuitive manner. Many general properties and specific applications of this approach have been developed for models with a discrete set of choices (e.g. McKelvey and Palfrey, 1995, 1998; McKelvey, Palfrey and Weber, 2000; Goeree and Holt, 2000a,b,c; Goeree, Holt, and Palfrey, 1999, 2000). In contrast, this paper considers models with a continuum of decisions, of the type that naturally arise in standard economic models where prices, claims, efforts, locations, etc. are usually assumed to be continuous. With noisy behavior and an continuum of decisions, the equilibria are characterized by densities of decisions, and the properties of these models are studied by using techniques developed for the analysis of functions, i.e. differential equations, stochastic dominance, fixed points in function space, etc. The contribution of this paper is to provide an easy-to-use tool kit of theoretical results (existence, uniqueness, symmetry, and comparative statics), which can serve as a foundation for subsequent applications. In addition, this paper gives a unified perspective on behavioral anomalies that have been reported in a series of laboratory experiments.

¹ See Green and Shapiro (1994) for a critical view and Ostrom (1998) for a more favorable view.

II. Stochastic Elements and Probabilistic Choice

Regardless of whether randomness, or noise, is due to preference shocks, experimentation, or actual mistakes in judgement, the effect can be particularly important when players' payoffs are sensitive to others' decisions, e.g. when payoffs are discontinuous as in auctions, or highly interrelated as in coordination games. Nor does noise cancel out when the Nash equilibrium is near a boundary of the set of feasible actions and noise pushes actions towards the interior, as in a Bertrand game in which the Nash equilibrium price equals marginal cost. In such games, small amounts of noise in decision making may have a large "snowball" effect when endogenous interactions are considered.² In particular, we will show that an amount of noise that only causes minor deviations from the optimal decision at the *individual* level, may cause a dramatic shift in behavior in a game where one player's choice affects others' expected payoffs.

The Nash equilibrium in the above mentioned games is often insensitive to parameter changes that most observers would expect to have a large impact on actual behavior. In a minimum-effort coordination game, for example, a player's payoff is the minimum of all player's efforts minus the cost of the player's own effort. With simultaneous choices, both intuition and experimental evidence suggest that coordination on desirable, high-effort outcomes will be harder with more players and higher effort costs, despite the fact that any common effort level is a Nash equilibrium (Goeree and Holt, 1999b). Another well known example is the "Bertrand paradox" that the Nash equilibrium price is equal to marginal cost, regardless of the number of competitors, even though intuition and experimental evidence suggest otherwise (Dufwenberg and Gneezy, 2000).

The rationality assumption implicit in the Nash approach is that decisions are determined by the *signs* of the payoff differences, not by the magnitudes of the payoff gains or losses. But the losses for unilateral deviations from a Nash equilibrium are often highly asymmetric. In the minimum-effort coordination game, for example, a unilateral increase in effort above a common (Nash) effort level is not very risky if the marginal cost of effort is small, while a unilateral decrease would reduce the minimum and not save very much in terms of effort cost. Similarly,

² Similarly, Akerlof and Yellen (1985) show that small deviations from rationality can have first-order consequences for equilibrium behavior. Alternatively, exogenous noise in the communication process may have a large impact on equilibrium outcomes, see, for instance, Rubinstein's (1989) electronic-mail game.

an effort increase is relatively more risky when effort costs are high. In each case, deviations in the less risky direction are more likely, and this is why effort levels observed in laboratory experiments are inversely related to effort cost.

Many of the counter-intuitive predictions of a Nash equilibrium disappear when some noise is introduced into the decision-making process, which is the approach taken in this paper. This randomness is modeled using a probabilistic choice function, i.e. the probability of making a particular decision is a smoothly increasing function of the payoff associated with that decision. One attractive interpretation of probabilistic choice models is that the apparent noisiness is due to unobserved shocks in preferences, which cause behavior to appear more random when the observed payoffs become approximately equal. Of course, mistakes and trembles are also possible, and these presumably would also be more likely to have an effect when payoff differences are small, i.e. when the cost of a mistake is small. In either case, probabilistic choice rules have the property that the probability of choosing the "best" decision is not one, and choice probabilities will be close to uniform when the other decisions are only slightly less attractive.

When a probabilistic choice function is used to analyze the interaction of strategic players, one has to model beliefs about others' decisions, since these beliefs determine expected payoffs. When prior experience with the game is available, beliefs will evolve as people learn. Learning slows down as observed decisions look more and more like prior beliefs, i.e. as surprises are reduced. In a steady state, systematic learning ceases when beliefs are consistent with observed decisions. Following McKelvey and Palfrey (1995), the equilibrium condition used here has the consistency property that belief probabilities which determine expected payoffs match the choice probabilities that result from applying a probabilistic choice rule to those expected payoffs. In other words, players take into account the errors in others' decisions.

Perhaps the most commonly used probabilistic choice function in empirical work is the logit model, in which the probability of choosing a decision is proportional to an exponential function of its expected payoff. This logit rule exhibits nice theoretical properties, such as having choice probabilities be unaffected by adding a constant to all payoffs. We have used the logit equilibrium extensively in a series of applications that include rent-seeking contests, price competition, bargaining, public goods games, coordination games, first-price auctions, and social

dilemmas with continuous choices.³ In the process, we noticed that many of the models share a common auction-like structure with payoff functions that depend on rank, i.e. whether a player's decision is higher or lower than another's. In this paper, we offer general proofs of theoretical properties based on characteristics of the expected payoff functions. Section III summarizes a logit equilibrium model of noisy behavior for games with rank-based outcomes. Proofs of existence, uniqueness, and comparative statics follow in section IV. In section V, we apply these results to a variety of models that represent many of the standard applications of game theory to economics and social science. Comparisons with learning theories and other ways of explaining behavioral anomalies are discussed in section VI, and the final section concludes.

III. An Equilibrium Model of Noisy Behavior in Auction-Like Games

The standard way to motivate a probabilistic choice rule is to specify a utility function with a stochastic component. If decision i has expected payoff π_i^e , then the person is assumed to choose the decision with the highest value of $U(i) = \pi_i^e + \mu \varepsilon_i$, where μ is a positive "error" parameter and ε_i is the realization of a random variable. When $\mu = 0$, the decision with the highest expected payoff is selected, but high values of μ imply more noise relative to payoff maximization. This noise can be due to either 1) errors, e.g. distractions, perception biases, or miscalculations that lead to non-optimal decisions, or 2) unobserved utility shocks that make rational behavior look noisy to an outside observer. Regardless of the source, the result is that choice is stochastic, and the distribution of the random variable determines the form of the choice probabilities. A normal distribution yields the probit model, while a double exponential distribution gives rise to the logit model, in which case the choice probabilities are proportional to exponential functions of expected payoffs. In particular, the logit probability of choosing

³ In a couple of these applications, the derivation of some theoretical properties are provided, but they rely on the special structure of the model being studied (Anderson, Goeree, and Holt, 1998a,b, 1999; Capra, *et al.*, 1999b, and Goeree, Anderson, and Holt, 1998). In other cases, theoretical results are absent, and the focus is on estimations that are based on a numerical analysis for the specific parameters of the experiment (Capra *et al.*, 1999a,b; Goeree and Holt, 1999a,b; and Goeree, Holt, and Laury, 1999). In contrast, this paper provides an extensive treatment of the theoretical properties of logit equilibria in a broad class of games that includes many of the applications discussed above as special cases. In addition, the existence proof in Proposition 1 applies to a general class of probabilistic choice functions that includes the logit model as a special case.

alternative i is proportional to $\exp(\pi^e(i)/\mu)$, where higher values of the error parameter μ make choice probabilities less sensitive to expected payoffs.⁴

With a continuum of decisions on $[\underline{x}, \bar{x}]$, the logit model specifies a choice density that is proportional to an exponential function of expected payoffs:

$$(1) \quad f(x) = \frac{\exp(\pi^e(x)/\mu)}{\int_{\underline{x}}^{\bar{x}} \exp(\pi^e(y)/\mu) dy},$$

where the integral in the denominator is a constant that makes the density integrate to one.⁵ Note that payoff differences do not matter as μ goes to infinity, since the argument of the exponential function in (1) goes to zero and the density becomes flat (uniform), irrespective of the payoffs. Conversely, payoff differences are "blown up" as μ goes to zero, and the density piles up at the optimal decision.^{6,7} Limiting cases are useful for providing intuition, but we will argue below

⁴ An alternative justification for use of the logit formula follows from work in mathematical psychology. Luce (1959) provides an axiomatic derivation of this type of decision rule; he showed that if the ratio of choice probabilities for any pair of decisions is independent of the payoffs of all other decisions, then the choice probability for decision i can be expressed as a ratio: $u_i/\sum_j u_j$, where u_i is a "scale value" number associated with decision i . If one adds an assumption that choice probabilities are unaffected by adding a constant to all payoffs, then it can be shown that the scale values are exponential functions of expected payoffs. Besides having these theoretical properties, the logit rule is convenient for estimation by providing a parsimonious one-parameter model of noisy behavior that includes perfect rationality (Nash) as a limiting case.

⁵ An independent motivation for the equilibrium condition in (1) is provided by Anderson, Goeree, and Holt (1999), who postulate a directional-adjustment evolutionary model that yields (1) as a stationary state. The model is formulated in continuous time with a population of players. The primitive assumption is that each player adjusts the decision in the direction increasing expected payoff, at a rate that is proportional to the slope of the payoff function, plus some Brownian motion. Thus if the payoff function is flat, decisions change randomly, but if the payoff function is steep, then adjustments in an improving direction dominate the noise effect. We show that the stationary states for this process are logit equilibria. The advantage of a dynamic analysis is that it can be used to consider stability and elimination of unstable equilibria. Anderson, Goeree, and Holt (1999) show that the gradient-based directional adjustment process is globally stable for all potential games, with a Liauponov function that can be interpreted as a weighted combination of expected potential and entropy.

⁶ These effects of μ can be evaluated by taking ratios of densities in (1) for two decisions, x_1 and x_2 : $f(x_1)/f(x_2) = \exp((\pi^e(x_1) - \pi^e(x_2))/\mu)$.

⁷ Notice from (1) that a doubling of payoffs is equivalent to cutting the error rate in half. This property captures the intuitive idea that an increase in incentives will reduce noise in experimental data (see Smith and Walker, 1993, for supporting evidence). In fact, Goeree, Holt and Palfrey (2000) show that a logit equilibrium explains the behavioral response to a quadrupling of one player's payoffs in a matrix game. The predictions were obtained by estimating risk aversion and error parameters for a data set that included this and six other matrix games.

that it is the intermediate values of μ that are most relevant for explaining data of human subjects, who are neither perfectly rational nor perfectly noisy.⁸ In this case, choice probabilities are smoothly increasing functions of expected payoffs, so these probabilities will be affected by asymmetries in the costs of deviating from the payoff-maximizing decision.⁹

In order to apply this model to games, one must deal with the fact that distributions of others' decisions enter the expected payoff function on the right side of equation (1). A Nash-like consistency condition is that the belief distributions that determine expected payoffs on the right side of (1) match the decision distributions on the left that result from applying the logit rule to those expected payoffs. Thus the logit choice rule in (1) determines players' equilibrium distributions as a fixed point. This is known as a logit equilibrium, which is a special case of the "quantal response equilibrium" introduced by McKelvey and Palfrey (1995, 1998).

Differentiating both sides of (1) with respect to x (and rearranging) yields:

$$(2) \quad \pi^{e^i}(x) f(x) - \mu f'(x) = 0,$$

which provides the "logit differential equation" in the equilibrium choice density, first introduced in Anderson, Goeree, and Holt (1998a). This density has the same slope as the expected payoff function in equilibrium, so their relative maxima coincide, although the spread in the density around the payoff-maximizing choice is determined by μ . The use of (2) to calculate the equilibrium distribution is illustrated next in the context of an example that highlights the dramatic effects of adding noise to a standard Nash equilibrium analysis.

Example 1. Traveler's Dilemma

The game that has the widest range of applications in the social science literature is the

⁸ In contrast, most previous theoretical work on models with noise is primarily concerned with the limit as noise is removed to yield a selection among the Nash equilibria, e.g. "perfection" (Selten, 1975), "evolutionary drift" (Binmore and Samuelson, 1999), and "risk dominance" (Carlsson and van Damme, 1993).

⁹ Radner's (1980) " ϵ -equilibrium" allows strategy combinations with the property that unilateral deviations cannot yield payoff increases that exceed (some small amount) ϵ . Behavior in an ϵ -equilibrium is "discontinuous" in the sense that deviations do not occur unless the gain is greater than ϵ , in which case they occur with probability one. In contrast, the probabilistic choice approach in (1) is based on the idea that choice probabilities are smooth, increasing functions of expected payoffs.

social dilemma in which the unique Nash equilibrium yields an outcome that is worse for all players than a non-equilibrium cooperative outcome. Unlike the familiar prisoner's dilemma game, the traveler's dilemma is a social dilemma in which the Nash strategy is not a dominant strategy. This game describes a situation in which two people lose identical objects and must make simultaneous loss claims in a pre-specified interval (Basu, 1994). Each player is reimbursed at a rate that equals the minimum of the two claims, with a fixed penalty amount $\$R$ being transferred from the high claimant to the low claimant if the claims are unequal. This penalty gives each an incentive to "undercut" the other, and the unique Nash equilibrium is for both to claim the lowest possible amount, despite the fact that there is little risk of making a high claim when R is small. The traveler's dilemma game is important precisely because of this sharp difference between economic intuition and the unique Nash prediction.

The expected payoff function for the traveler's dilemma game is:

$$(3) \quad \pi_i^e(x) = \int_{\underline{x}}^x (y - R) f_j(y) dy + \int_x^{\bar{x}} (x + R) f_j(y) dy, \quad i, j = 1, 2, \quad j \neq i,$$

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 can be expressed:

$$(4) \quad \pi_i^{e'}(x) = 1 - F_j(x) - 2Rf_j(x), \quad i, j = 1, 2, \quad j \neq i.$$

The $1 - F_j(x)$ term picks up the probability that the other's claim is higher, i.e. that a unilateral increase will raise the minimum. The final term in (4) is due to the payoff discontinuity at equal claims: $-2R$ is the payoff reduction involved in "crossing over" the other's claim, i.e. losing the R reward and paying the R penalty. This crossover occurs with a probability that is determined by the density $f_j(x)$. In most of the applications considered in section IV below, the marginal expected payoff function will have terms with distribution functions, reflecting the probabilities of being higher or lower than the others, and terms involving the densities, reflecting "cross-over" probabilities when there are payoff discontinuities.

In order to solve for the equilibrium distribution, substitute the expected payoff derivative

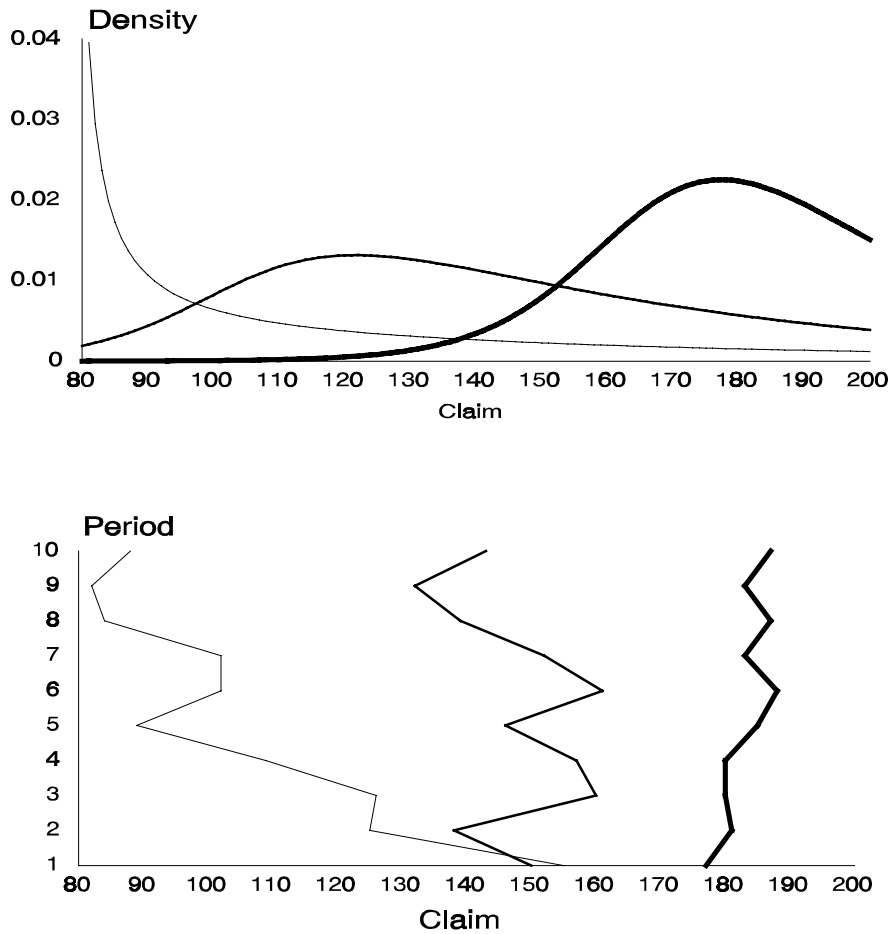


Figure 1. The Traveler's Dilemma.
 Key: Predicted Claim Densities (Top) and Actual Claim Averages (Bottom)
 for $R = 50$ (Left), $R = 20$ (Center), and $R = 10$ (Right)

(4) into the logit differential equation, which yields a second-order differential equation in the equilibrium $F(x)$. Although no analytical solutions exists, this equation can easily be solved numerically for a given value of μ . The top part of Figure 1 shows the equilibrium densities for $\mu = 8.5$ (Capra *et al.*, 1999) and penalty/reward parameters of 10, 25, and 50. Notice that the predictions of this model are very sensitive to changes in R . With $R = 50$, the density piles up near the unique Nash prediction of 80 on the left side of the graph, but the density is concentrated at the opposite side of the set of feasible claims when $R = 10$. The general pattern

of deviations from the Nash prediction shows up in the bottom part of Figure 1, which shows the data averages for each treatment, as a function of the period number on the vertical axis.

These large deviations from the unique Nash prediction are relatively insensitive to the error parameter and would occur even if this parameter were halved or doubled from the level used in the figure ($\mu = 8.5$). To get a feel for the effects of an error parameter of this magnitude, suppose there are two decisions, 1 and 2, that yield an expected payoff difference of about 25 cents, i.e. $\pi^e(2) - \pi^e(1) = 25$ cents (which is sometimes thought to be about as low as you can go in designing experiments with salient payoffs). In this case, the logit probability of choosing the incorrect decision 1 is: $1/(1 + \exp((\pi^e(2) - \pi^e(1))/\mu)) = 1/(1 + \exp(25/8.5))$, which is about $1/(1 + e^3)$, or about 0.05. Similarly, it can be shown that the probability of making an error when the expected payoff difference is 50 cents is 0.002.

The numerical calculations used to construct the upper part of Figure 1 only pertain to the particular parameters used in the experiment, which raises some interesting theoretical issues: will a logit equilibrium generally exist for this game and others like it, will the equilibrium be unique, symmetric, and single-peaked, and will increases in the incentive parameter R always reduce claim distributions? These theoretical issues were not addressed in the original paper (Capra *et al.*, 1999), but are resolved by the propositions that follow.

Rank-Based Payoffs and the Local Payoff Property

In the traveler's dilemma example, the payoff function consists of two parts, where each part is the integral of a payoff function that is relevant for the case of whether the player's decision is the higher one or not. This rank-based payoff also arises naturally in other contexts: in price competition games where the low-priced firm gains more sales, in minimum-effort coordination games where the common part of the production depends on another's effort only when it is lower than one's own effort, and in location games on a line where the market divides with the left part going to the firm with the left-most location. These applications can be handled with a rank-based expected payoff function that has two parts. First consider two-person games and let $\alpha_H(x)$ and $\alpha_L(x)$ be payoff components associated with one's own decision when it is higher or lower than the other's decision. Similarly, let $\beta_H(y)$ and $\beta_L(y)$ be payoff components associated with the other player's decision when one's own rank is high or low. Then the

traveler's dilemma payoff function in (4) is a special case of:

$$(5) \quad \pi_i^e(x) = \int_{\underline{x}}^x [\alpha_H(x) + \beta_H(y)] f_j(y) dy + \int_x^{\bar{x}} [\alpha_L(x) + \beta_L(y)] f_j(y) dy,$$

with $\alpha_H(x) = -R$, $\beta_H(y) = y$, $\alpha_L(x) = x$, and $\beta_L(y) = R$. The formulation in (5) also includes cases where the payoffs are not dependent on the relative rank, as in the public goods game discussed below. As long as these two component payoff functions are additively separable and continuous in own and other's decision, it is straightforward to verify that the expected payoff derivative will have the "local" property that it depends on the player's own decision x and on the other's distribution and density functions evaluated at x . In this case, we can express the expected payoff derivative as: $\pi_i^{e*}(F_j(x), f_j(x), x, \alpha)$, as is the case in equation (4), where the α notation represents an exogenous shift parameter that corresponds to the penalty parameter in the traveler's dilemma.

Equation (5) is easily adapted to the N -player case in which one's payoff depends on whether one has the highest (or lowest) decision. If having the highest decision is critical, as in an auction, then the H and L subscripts represent the case where one's decision is the highest or not, respectively, and the density $f(y)$ used in the integrals is replaced by the density of the maximum of the $N-1$ other decisions. In a second-price auction for a prize with value V , for example, $\alpha_H(x) = V$, $\beta_H(y) = -y$, and $\alpha_L(x) = \beta_L(y) = 0$. Given the assumed additive separability of the α and β functions, it is straightforward to verify that (5) (with the density of the maximum (or minimum) of the others' decisions substituted for $f(y)$) yields an analogous "local" property for N -player games. In other words, the expected payoff derivative, $\pi_i^{e*}(F_{-i}(x), f_{-i}(x), x, \alpha)$, depends on the distribution and density functions of all $N - 1$ other players, $j = 1, \dots, N, j \neq i$. We will use the term *local payoff property* for games in which the expected payoff derivative is only a function of probability-like terms, i.e. of the other players' distribution and density functions evaluated at x , and possibly of x and exogenous parameters.

IV. Properties of Equilibrium: Existence, Uniqueness, and Comparative Statics

The expected payoff derivatives for particular games, e.g. (4), can be used together with the logit differential equation (2) to calculate equilibrium choice distributions for given values of the exogenous payoff and error parameters. These calculations are vastly simplified if we know that there exists a solution that is symmetric across players.

Proposition 1. (Existence) There exists a logit equilibrium for all N -player games with a continuum of feasible decisions when players' expected payoffs are bounded and continuous in others' distribution functions. Moreover, the equilibrium distribution is differentiable.

The proof in Appendix A is obtained by applying Schauder's fixed point theorem to the mapping in (1). In fact, the proof applies to the more general case where the exponential functions in (1) are replaced by strictly positive and strictly increasing functions, which allows other probabilistic choice rules besides the logit/exponential form.

Uniqueness

When the expected payoff derivative satisfies the local-payoff property, the logit differential equation in (1) is a second-order differential equation with boundary conditions $F(x) = 0$ and $F(\bar{x}) = 1$.¹⁰ We will show that for many games with rank-based payoffs the symmetric logit equilibrium is unique. The method of proof is by contradiction: we start by assuming that there exists a second symmetric logit equilibrium, and then show that this is impossible under the assumed conditions. There are several "directions" in which one can obtain a contradiction, which explains why there are alternative sets of assumptions for each proposition. These alternative assumptions will enable us to evaluate uniqueness for an array of diverse examples in the next section. Parts of the uniqueness proof are included in the text here because they are representative of the symmetry and comparative statics proofs that are found in the appendices. In particular, all of these proofs have graphical "lens" structures, as indicated below.

¹⁰ Incidentally, it is a property of the logit choice function that all feasible decisions in the interval $[x, \bar{x}]$ have strictly positive chance of being selected.

Proposition 2 (Uniqueness). Any symmetric logit equilibrium for a game satisfying the local payoff property is unique if the expected payoff derivative, $\pi^e(F, f, x, \alpha)$ is either

- a) strictly decreasing in x , or
- b) strictly increasing in the common distribution function F , or
- c) independent of x and strictly decreasing in f , or
- d) a polynomial expression in F , with no terms involving f or x .

Proof for Parts (a) and (b). Suppose, in contradiction to the statement of the proposition, that there exist (at least) two symmetric logit equilibrium distributions, denoted by F_1 and F_2 . Without loss of generality, assume $F_1(x)$ is lower on some interval, as shown in Figure 2.




Figure 2. Horizontal Lens Proof: A Configuration that Yields a Contradiction when $F_2 > F_1$

Case (a) is based on a *horizontal lens proof*. Any region of divergence between the distribution functions will have a maximum horizontal difference, as indicated by the horizontal line in Figure 2 at height $F^* = F_1(x_1) = F_2(x_2)$. The first- and second-order necessary conditions for the distance to be maximized at F^* are that the slopes of the distribution functions be identical at F^* , i.e. $f_1(x_1) = f_2(x_2)$, and that $f_1'(x_1) \geq f_2'(x_2)$. In Case (a), $\pi^e(F, f, x, \alpha)$ is decreasing in x ,




Figure 3. Vertical Lens Proof: A Configuration that Yields a Contradiction when $F_2 > F_1$.

and since the values of the density and distribution functions are equal, it follows that

$$(6) \quad \pi_1^{e'}(F_1(x_1), f_1(x_1), x_1, \alpha) < \pi_2^{e'}(F_2(x_2), f_2(x_2), x_2, \alpha).$$

Then the logit differential equation in (2) implies that $f_1'(x_1) < f_2'(x_2)$, which yields the desired contradiction of the necessary conditions for the distance between F_1 and F_2 to be maximized.

Case (b) is proved with a *vertical lens proof*. If there are two symmetric distribution functions, then they must have a maximum vertical distance at x^* as shown in Figure 3. The first-order condition is that the slopes are equal, so the densities are the same at x^* . Under assumption (b), $\pi^{e'}(F, f, x, \alpha)$ is strictly increasing in F , and it follows from (1) that $f_1'(x_1) < f_2'(x_2)$, which yields the desired contradiction. Q. E. D.

The proof of Proposition 2(c) in Appendix B can be skipped on a first reading since it involves a transformation-of-variables technique that is not used in any of the other proofs that follow. Note, however, that Proposition 2(c) implies uniqueness for the traveler's dilemma example, since the expected payoff derivative in (4) is independent of x and decreasing in f . Proposition 2(d), also proved in Appendix B, is based on observation that the logit differential equation (1) can be integrated directly when the expected payoff derivative is a polynomial in

F , and the resulting expression for the density produces the desired contradiction.

Even when the symmetric equilibrium is unique, there may exist asymmetric equilibria for some games, e.g. those with asymmetric Nash equilibria. In experiments we often restrict attention to symmetric equilibria when subjects are matched from single-population protocols and have no way of coordinating on asymmetric equilibria (Harrison and Hirshleifer, 1989). In other games it is possible to use properties of the expected payoff function and its slope, $\pi^{e'}(F_j, f_j, x, \alpha)$, to prove that an equilibrium is necessarily symmetric. The symmetry result in Proposition 3, which is stated and proved in Appendix B, is based on the assumption that $\pi^{e'}(F_j, f_j, x, \alpha)$ is strictly decreasing in F_j , as is the case in the traveler's dilemma game.

Comparative Statics

It is apparent from (1) that the logit equilibrium density is sensitive to all aspects of the expected payoff function, i.e. choice propensities are affected by magnitudes of expected payoff differences, not just by the signs of the differences as in a Nash equilibrium. In particular, the logit predictions can differ sharply from Nash predictions when the costs of deviations from a Nash equilibrium are highly asymmetric, and when deviations in the less costly direction make further deviations in that direction even less risky, creating a feedback effect. These asymmetric payoff effects can be accentuated by shifts in parameters that do not alter the Nash predictions. Since the logit equilibrium is a probability distribution, the comparative statics will be in terms of shifts in distribution functions. Our results pertain to shifts in the sense of first-degree stochastic dominance, i.e. the distribution of decisions increases in this sense when the distribution function shifts down for all interior values of x . We assume that the expected payoff derivative, $\pi^{e'}(F, f, x, \alpha)$, is increasing in an exogenous parameter α , *ceteris paribus*. The next proposition shows that an increase in α raises the logit equilibrium distribution in the sense of first-degree stochastic dominance. Only monotonicity in α is required, since any parameter that decreases marginal profit can be rewritten so that marginal expected payoff is strictly increasing in the redefined parameter.

Proposition 4 (Comparative Statics for a Symmetric Equilibrium). Suppose that the shift parameter increases marginal expected payoffs, i.e. $\partial\pi^e(F, f, x, \alpha)/\partial\alpha > 0$, for a symmetric game satisfying the local payoff property. Then an increase in α yields stochastically higher logit equilibrium decisions (in the sense of first-degree stochastic dominance) if either

a) $\partial\pi^e/\partial x \leq 0$, or

b) $\partial\pi^e/\partial F \geq 0$.

The proof is provided in Appendix C. Case (a), which is proved with a horizontal lens argument, is based on a weak concavity property that will be satisfied by all of the games considered in this paper. In the traveler's dilemma game, for example, $\partial\pi^e/\partial x$ is exactly 0, so case (a) applies. Let $\alpha = -R$. Since the expected payoff derivative in (4) is decreasing in R , it follows that a decrease in R will raise α and hence will raise claims in the sense of first-degree stochastic dominance, which is consistent with the data in Figure 1. This increase in claims, however intuitive, is not predicted by standard game theory, since the Nash equilibrium is the minimum feasible claim as long as R is strictly positive. The logit result is intuitive given that a reduction in the penalty parameter raises the slope of the expected payoff function and makes it less risky and less costly to raise one's claim unilaterally.

Finally, consider the effects of changes in the error parameter μ . Although one would not normally think of the error parameter as being under the control of the experimenter, it is apparent from (1) that multiplicative scaling up of all payoffs corresponds to a reduction in the error parameter, i.e. multiplying expected payoffs by γ is equivalent to multiplying μ by $1/\gamma$. Error parameter effects may also be of interest if one believes that noise will decline as subjects become experienced, and the purification of noise might provide a selection criterion (McKelvey and Palfrey, 1998). The effects of changes in μ are generally not monotonic, since the whole π^e function in (2) is divided by μ , but the case when marginal payoffs are everywhere positive (negative) can be handled (the proof is essentially the same as for Proposition 4).

Proposition 5 (Effects of a Decrease in the Error Parameter). Suppose that marginal expected payoffs, $\pi^e(F, f, x)$, are everywhere positive (negative) for a symmetric game satisfying the local payoff property. Then a decrease in μ yields stochastically higher (lower) logit equilibrium

decisions (in the sense of first-degree stochastic dominance) if either

a) $\partial\pi^e/\partial x \leq 0$, or

b) $\partial\pi^e/\partial F \geq 0$.

This result is intuitive: when expected payoffs are increasing in x , so is the density determined by (1) is increasing, and an increase in noise "flattens" the density "pushing" mass to the left. Conversely, if the expected payoff derivative is negative, the density is decreasing and an increase in noise pushes mass to the right and causes a stochastic increase in decisions.

So far we have confined attention to games in which the payoff functions are symmetric across the two firms. However, specific asymmetries are readily introduced. In particular, suppose the functional forms of $\pi_1^e(F_2, f_2, x, \alpha_1)$ and $\pi_2^e(F_1, f_1, x, \alpha_2)$ are the same but $\alpha_1 > \alpha_2$.

Proposition 6 (Comparative Statics for Asymmetric Payoffs). Suppose that the shift parameter increases marginal expected payoffs, i.e. $\partial\pi^e(F, f, x, \alpha)/\partial\alpha_i > 0$, and let $\alpha_2 > \alpha_1$ in a game satisfying the local payoff property. Then player 2's logit equilibrium distribution of decisions is stochastically higher than that of player 1, i.e. the distribution function for player 2 is lower at each interior value of x , if either

a) $\partial\pi^e/\partial x \leq 0$, or

b) $\partial\pi^e/\partial F \geq 0$.

The proofs in Appendix C are again lens proofs, horizontal for case (a) and vertical for case (b). In a traveler's dilemma game with individual-specific R_i parameters, this proposition would imply that the person with higher penalty-reward parameter would have stochastically lower claims.

Other Properties

For many applications, it is possible to show that the symmetric logit equilibrium density function that solves (1) is single peaked. Since this proposition pertains to symmetric equilibria, the player subscripts are dropped.

Proposition 7. (Single Peakedness) If the logit equilibrium for a game satisfying the local payoff property is symmetric and the expected payoff derivative, $\pi^e'(F, f, x, \alpha)$ is non-increasing in x and strictly decreasing in the common F function, then the equilibrium density that solves (2) will be single peaked.

The proof in Appendix D is based on assumed concavity-like properties of expected payoff function, which ensure that expected payoffs are single peaked, and hence that the exponential (or any other continuously increasing) functions of those expected payoffs in (1) are single peaked. Of course, the "single peak" maximum may be at a boundary point if the density is monotonic, as with the traveler's dilemma for high R values in Figure 1.

V. APPLICATIONS

The applications in this section include many types of games that are commonly used in economics and other some other social sciences: coordination, public goods, bargaining, auctions, and spatial location. These applications illustrate the usefulness of the theoretical propositions and the contrasts between logit equilibrium analysis and the special case of a Nash equilibrium.

Example 2: Minimum-effort Coordination Game

Coordination games, which date back to Rousseau's stag hunt problem, are perhaps second only to social dilemma games in terms of interest to economists and social scientists. Coordination games possess multiple Nash equilibrium, some of which are worse than others for all players, which raises the issue of how a group of people (or even a whole economy) can become mired in an inefficient equilibrium. First consider the minimum effort game described above, with a payoff equal to the lowest effort minus the cost of a player's own effort. Letting $f_j(x)$ and $F_j(x)$ denote the density and distribution functions associated with the *other* player's decision, it is straightforward to write player i 's expected payoff from choosing an effort level, x :

$$(7) \quad \pi_i^e(x) = \int_0^x y f_j(y) dy + \int_x^{\bar{x}} x f_j(y) dy - cx, \quad i, j = 1, 2, \quad j \neq i,$$

where the first term on the right side pertains to the case where the other's effort is below the player's own effort, x , and the second term pertains to the case where the player's own effort is the minimum. In order to work with the logit differential equation (2), consider the derivative of this expected payoff with respect to x :

$$(8) \quad \pi_i^{e'}(x) = 1 - F_j(x) - c, \quad i, j = 1, 2, \quad j \neq i.$$

The intuition behind (8) is clear, since $1 - F_j(x)$ is the probability that the other's effort is higher, this is also the probability that an increase in effort will raise the minimum, but such an increase will incur a cost of c . The expected payoff derivative in (8) is positive if $F_j(x) = 0$, and it is negative if $F_j(x) = 1$, so *any* common effort is a pure-strategy Nash equilibrium, even though all players prefer higher common efforts. Also, notice that the effort cost c determines the extent of the asymmetry in loss incurred by deviating from any common effort.

Proposition 2(d) implies uniqueness, and the conditions of the comparative statics Proposition 4 are also satisfied. Since π^{e^*} is strictly decreasing in effort cost c , efforts are stochastically lower in a minimum effort coordination game if the effort cost is increased, despite the fact that changes in c do not alter the set of Nash equilibria (as long as $0 < c < 1$). Goeree and Holt (1999b) report results for a two-person minimum effort experiment in which an increase in effort cost from .25 to .75 lowered average efforts from 159 to 126. The logit predictions, based on an estimated $\mu = 7.4$, were 154 and 126 respectively. The estimated μ had a standard error of .3, so the null hypothesis of $\mu = 0$ (Nash) can be rejected at any conventional level of significance.

Coordinating on high-effort outcomes is far more difficult in experiments with larger numbers of players, so consider the effect of having more than two players. With $N-1$ other players, the increase in effort will only raise the minimum when all $N-1$ others are higher, so the right side of (8) would become the product of all $1 - F_j(x)$ terms for the others, with the addition of a term, $-c$, reflecting the cost effect as before. In a symmetric equilibrium, $\pi^{e^*}(x) = (1 -$

$F(x)^{N-1} - c$, which is decreasing in N , so an increase in the number of players will result in a stochastic reduction in effort. Again, this intuitive result is notable since the set of Nash equilibria is independent of N .

Example 3: The Median Effort and Other Order-Statistic Coordination Games

The minimum-effort game is only one of many types of coordination games. Consider a three-person, *median*-effort coordination game in which each player's payoff is the median effort minus the cost of their own effort. Instead of writing out the expected payoff function and differentiating, the marginal expected payoff can be obtained directly since the marginal effect of an effort increase is the probability that one's effort is the median effort minus the effort cost:

$$(9) \quad \pi_i^e / (x) = 2F(x)(1 - F(x)) - c.$$

The number 2 on the right side of (9) reflects the fact that there are two ways in which one player can be below x and one can be above x , and each of these cases occurs with probability $F(x)(1-F(x))$. A similar expression is obtained for an N -player game in which the payoff is the k th order statistic minus the own effort cost. The marginal value of raising one's effort is the probability that an effort increase is relevant, which is the probability that $k-1$ others are above x and $N - k$ others below x . This probability again yields a formula for the marginal expected payoff that is an N th order polynomial in F , with a cost term, $-c$, attached. These intuitive derivations of expected payoff derivatives are useful because they serve as a check on the straightforward but tedious derivations based on differentiation.

These coordination games have the local payoff property, since the expected payoff derivative depends only on powers of the cumulative distribution function. This ensures existence of a symmetric equilibrium, and by Proposition 2(d), uniqueness. The expected payoff derivative is non-increasing in x (holding F constant), so Proposition 4(a) implies that the common effort distribution is stochastically increasing in $-c$, or decreasing in c . This intuitive effort-cost effect is supported by the data for 3-person median effort experiments in Goeree and Holt (1999b), where average efforts in the final three periods decreased from 157 to 137 and again to 113 as effort cost was raised from .1 to .4 and then to .6. There is a continuum of

asymmetric equilibria in the median effort game (with the top two efforts being equal and the lowest one at the lower bound), so the intuitive effort cost effects cannot be explained by a Nash analysis.

Example 4. Spatial Competition

The Hotelling model of spatial competition on a line has had wide applications in industrial organization, and generalizations of this model constitute the most common application of game theory in political science. Suppose that voters are located uniformly on a line of unit length in a single dimension (e.g. representing preferences on government spending). Two candidates choose locations on the line, and voters vote for the candidate who is closest to their preferred point on the line. If the two locations are x_1 and x_2 , then the division point that determines vote shares is the midpoint: $(x_1+x_2)/2$. The unique Nash equilibrium is for each to locate at the midpoint of the line, which is an example of the "median voter theorem." To make this model more interesting, let's assume that this is a primary, and that candidates incur a cost in the general election when they move away from the extreme left point (0), since the extreme left for this party is the center for the general electorate. Let this cost be denoted by cx , where x is the distance from the left side of the line. We chose this example because the unique Nash equilibrium is independent of c and remains at the midpoint as long as $c < 1/2$.¹¹

The logit equilibrium will be sensitive to the payoff asymmetries associated with the location costs. To see this, let $f_j(x)$ denote the choice density for the other candidate, then the expected vote share in the primary for location x is:

$$(10) \quad \pi_i^e(x) = \int_{\underline{x}}^x [1 - (x+y)/2] f_j(y) dy + \int_x^{\bar{x}} [(x+y)/2] f_j(y) dy - cx,$$

where the left term represents the case where the other candidate is to the right, the middle integral represents the case where the other candidate is to the left, and the final term is the

¹¹ To see this, note that the two locations should be adjacent in any Nash equilibrium, any adjacent locations away from the midpoint would give the person with the smaller share an incentive to move a small distance to capture the larger share. When $c > 1/2$, the unique Nash equilibrium is for both candidates to locate at the left boundary and share the vote.

location cost. In a symmetric equilibrium, the expected payoff derivative can be expressed:

$$(11) \quad \pi_i^{e'}(x) = -F(x)/2 + (1 - F(x))/2 + f(x)(1 - 2x) - c.$$

The first term on the right side is the probability of having the "higher" x , times the $-1/2$ that is the marginal *loss* from moving to the right, i.e. away from the other candidate's location. The second term is the analogous share *gain* from moving to the right when this is in the direction of the other candidate's position. The third term represents the probability of a crossover, measured by the density $f(x)$, times the effect of crossing over at x , i.e. of losing the vote share x to the left and gaining the vote share $1 - x$ to the right, for a net effect of $1 - 2x$.

Since $f(x)$ determined by logit probabilistic choice rule in (1) will always be strictly positive, it follows that the expected payoff derivative in (11) is strictly decreasing in x , holding the other (F, f) arguments constant, so the uniqueness and comparative statics theorems apply. It can be shown that the equilibrium density is symmetric around $1/2$ if $c = 0$, and the implication of Proposition 4 is that increases in c shift the densities the left.¹²

Example 5. Bertrand Competition in a Procurement Auction

Consider a model in which N sellers choose bid prices simultaneously, and the contract is awarded to the low-priced seller (ties occur with probability zero in a logit equilibrium with a continuum of price choices). With zero costs, it is straightforward to express the expected payoff for a bid of x in a symmetric equilibrium as: $x(1 - F(x))^{N-1}$, i.e. the price times the probability that all others are higher. Differentiation yields:

$$(12) \quad \pi_i^{e'}(x) = (1 - F(x))^{N-1} - x(N - 1)(1 - F(x))^{N-2}f(x),$$

where the first term on the right represents the probability that a price increase will be relevant (the others are higher), and the second term is the "cross-over" loss at x associated with the chance of overbidding in a symmetric equilibrium. Since the expected payoff derivative is decreasing in x , the symmetric equilibrium will be unique. The formula in (12) however is not

¹² Goeree and Holt have also applied these techniques to the analysis of three-person location problems (work in progress) to explain laboratory results that do not conform to Nash predictions.

Table I. Predicted Low Bids in Bertrand Game

	$N = 2$	$N = 3$	$N = 4$
$\mu = 1$	9.6	7.4	6.5
$\mu = 5$	23.7	16.9	13.9
$\mu = 8$	28.1	19.8	16.0
Laboratory Data ^a	26.4	19.0	15.2

^a Duwfenberg and Gneezy (1999)

decreasing in N , and in fact, an increase in the number of bidders does not result in a stochastic decrease in prices for any value of μ .¹³ However, we have calculated the expected value of the winning (low) bid for various values of μ that are in the range of μ values estimated from other experiments. An increase in the number of bidders from 2 to 3 to 4 lowers the procurement cost in this range, see Table I. With an error parameter of about 8, the logit predicted minimum bids are close to those reported by Dufwenberg and Gneezy (1998), and are inconsistent with the "Bertrand paradox" prediction that price will be driven to marginal cost (zero in this case) even for the case of two sellers.¹⁴ Baye and Morgan (1999) have also pointed out that prices above the Bertrand prediction can be explained by a (power function) quantal response equilibrium.

Example 6. Imperfect Price Competition with Meet-or-Release Clauses

The Bertrand paradox has inspired a number of models that relax the assumption that the firm with the low price makes all sales. Suppose that there is a number α_i of loyal buyers who purchase one unit from firm i . The remaining consumers, numbering β , purchase from the firm with the lowest price. For simplicity, assume that $\alpha_1 = \alpha_2 = \alpha$. Thus α represents the expected sales of the firm with the high price, and $\alpha + \beta$ represents the low-price firm's sales, which we will normalize to 1. Loyalty has its limits, and therefore buyers have implicit or explicit "meet-

¹³ This is because, for given N , the slope of the equilibrium density at the highest allowed bid must equal the slope for the lowest allowed bid, which ensures that the distribution functions will cross.

¹⁴ A small discrepancy is that the average bids predicted by the logit equilibrium are slightly higher than those reported by Duwfenberg and Gneezy (1999).

or-release" assurances that the high price firm must meet the lower price or allow the buyers to switch. Since the market share is higher for the low-price firm, and since their final sales prices are identical, the unique Bertrand/Nash equilibrium for a one-shot price competition game involves lowering price to marginal cost, regardless of the size of α . Intuition and laboratory evidence, however, suggests that price competition would be stiff for low values of α and that prices would be much higher as the market share of the high-price firm approaches 1/2. This intuition is again counter to the predictions of the unique Nash equilibrium. The expected payoff consists of two terms, depending on whether or not the firm has the higher price and sells α , or has the lower price and sells $\alpha + \beta = 1$:

$$(13) \quad \pi_i^e(x) = \alpha \int_0^x y f_j(y) dy + x(1 - F_j(x)),$$

which can be differentiated to obtain:

$$(14) \quad \pi_i^{e'}(x) = -(1 - \alpha) x f_j(x) + [1 - F_j(x)].$$

This is non-increasing in x and increasing in α , so Proposition 4 ensures that prices will be stochastically increasing in α , which measures the sales of the high-price firm. In the Capra *et al.* (2000) experiments, prices were restricted to the interval [60, 160], and an increase in α from .2 to .8 raised average prices from 69 to 129 in the final five periods. The unique Nash prediction is 60 for both treatments, which contrasts with the logit predictions of 78 (± 7) and 128 (± 6) respectively, based on an error parameter estimated from a previous traveler's dilemma paper (Capra *et al.*, 1999).¹⁵

Example 7. Capacity-Constrained Price Competition

Market power can arise when capacity constraints are introduced into the standard

¹⁵ A new estimate of the error parameter for this imperfect price competition experiment yields $\mu = 6.7$ with a standard error of 0.5, which again allows rejection of the null hypothesis associated with the Nash equilibrium (no errors). This estimated error parameter is quite close to the estimates of 7.4 for the minimum effort coordination game data (Goeree and Holt, 1999b) and 8.5 for the traveler's dilemma data (Capra *et al.*, 1999). These were repeated game experiments with random matching; we have obtained higher error parameter estimates for games only played once.

Bertrand duopoly model of price competition. Suppose that demand is inelastic at $K + D_r$ units at any price below \bar{x} , where K is the capacity of each firm and D_r is the residual demand obtained by the high-price firm. In a symmetric equilibrium, the expected payoff for a price of x is $[1 - F(x)]Kx + F(x)D_r x$, so $\pi^e = K - F(x)(K - D_r) - f(x)(K - D_r)x$, which satisfies the assumptions of Propositions 1 and 2, so the symmetric logit equilibrium exists and is unique. The implication of Proposition 4 is that an increase in firms' common capacity, K , will result in a stochastic reduction in prices. This intuitive prediction is also a property of the mixed-strategy Nash equilibrium obtained by equating expected profit to the safe earnings obtained by selling the residual demand at the highest price: $D_r \bar{x}$.¹⁶

Example 8. Public Goods

In a linear public goods game, each person makes a voluntary contribution, x_i , and the payoff depends on this contribution and on the sum of the others' contributions:

$$(15) \quad \pi_i(x_i) = E - x_i + R_I x_i + R_E \sum_{j \neq i} x_j,$$

where E is the endowment, R_I is the "internal return" received from one's own contribution, and R_E is the "external return" received from the sum of others' contributions. It is typically assumed that $R_I < 1$, so it is a dominant strategy not to contribute. The internal return may be greater than the external return if one's contribution is somehow located nearby, e.g. a flower garden will be seen more by the owner than by those passing on the street. Notice that this is a trivial special case of the rank-based payoffs in (5), since the payoffs do not depend on whether or not one's contribution is higher or lower than the others. In any case, the marginal expected return is a constant, $R_I - 1$, so uniqueness follows from Proposition 2(d). The constant marginal expected payoff is non-increasing in x , so the comparative statics implications of Proposition 4 are that an increase in the internal return will result in a stochastic increase in contributions, even though full free riding is a dominant-strategy Nash equilibrium. Dozens of linear public goods experiments have been conducted for the special case of (15) in which $R_I = R_E$, which is then

¹⁶ It can be shown that the logit and Nash models have different qualitative predictions in an asymmetric capacity model, since a firm's logit price distribution will be sensitive to changes in its own capacity. In contrast, a change in one firm's capacity will only affect the *other* firm's price distribution in a mixed equilibrium.

called the marginal per capita return (MPCR). The most salient result from this literature is the positive MPCR effect (Ledyard, 1995), which is predicted by the logit equilibrium and not by a Nash analysis.

Goeree, Holt, and Laury (1999) report experiments in which the internal and external returns are varied independently, since only the internal return affects the cost of contributing, whereas the external return may be relevant if one cares about other's earnings. The strongest treatment effect in the data was associated with the internal return, although contributions did increase with increases in the external return as well. Econometric analysis of the data suggests that the addition of an altruism factor to the basic preference structure explains the data well, and the estimated error parameter is highly significant, allowing rejection of the null hypothesis that the error rate is zero.

Example 9: The Best-Shot Game

In the minimum-effort game, the common payoff factor is determined by the lowest effort, or the "weakest link." The converse situation, which has been applied in some public goods problems, is known as the "best-shot" game, where the maximum decision determines the common payoff element (Harrison and Hirshleifer, 1989). Consider an asymmetric version where the contribution costs may differ. The marginal expected payoff is the probability that one's decision is the relevant best shot, minus the cost of the payoff increase:

$$(16) \quad \pi_i^e / (x) = F_j(x) - c_i.$$

These derivatives are non-increasing in x , so the player with the higher cost will have a stochastically lower distribution of contributions, by Proposition 6. This intuitive "own-cost" effect is not predicted in a Nash equilibrium. There are pure-strategy equilibria in which one player contributes nothing and the other makes a full contribution. It can also be shown that there is a mixed-strategy equilibrium, but it is a necessary property of a mixed equilibrium that changes in one's own payoff parameter do not affect one's own mixed distribution, which must remain fixed to keep the other player indifferent over the range of randomization.

The logit equilibrium for this game has a particularly interesting structure when $c_1 = c_2 = 1/2$, so that the logit differential equation, $f' = (Ff - f/2)/\mu$, can be integrated to obtain:

$$(17) \quad f(x) = f(0) - [F(x) - F(x)^2]/2\mu = f(0) - F(x)[1 - F(x)]/2\mu.$$

This is the formula describing the progression of an epidemic, where $F(x)$ represents the fraction of uninfected people at time x , so the rate of new infections, $-f(x)$, is a linear function of $F(x)[1 - F(x)]$, which is the probability that an uninfected person meets an infected person. It is well known that this dynamic system traces out a logistic curve, and therefore, the logit equilibrium will be a (truncated) logistic distribution.

Summary

Propositions 1 and 2 guarantee the existence of a unique, symmetric equilibrium for all examples considered (including the symmetric version of the best-shot game). Moreover, all examples satisfy the conditions of Proposition 4, so theoretical comparative statics results can be determined, except the numbers effect in the Bertrand game, which we analyzed numerically. There are laboratory experiments to evaluate the qualitative comparative statics predictions for six of these games, as summarized in Table 2. The left column shows the expected payoff derivative, and the second column indicates the sign of the comparative statics effect associated with each variable, where the + (or -) sign indicates that an increase in the exogenous variable results in an increase (or decrease) in decisions in the sense of first-degree stochastic dominance. The third column summarizes the directions of comparative statics effects reported in the experiments cited in the footnotes. For comparison, the comparative statics properties of the symmetric Nash equilibrium are listed in the right-hand column. In all cases, the reported effects for laboratory experiments correspond to the logit equilibrium predictions. Most important, none of the comparative statics effects listed are explained by the Nash equilibrium for that game. This contrast is due to the fact that the shift variables listed in the table change the magnitudes of payoff differences but not the signs, so the Nash equilibria are invariant to changes in these variables.

VI. Relationship with Other Approaches to Explaining Anomalies in Game Experiments

The noisy equilibrium models developed in this paper are complemented by noisy models of learning, evolution, and adjustment. Learning models with probabilistic choices will be

Table II. Summary of Comparative Statics Results with Supporting Laboratory Evidence

Game: expected payoff derivative	Logit Comparative Statics	Laboratory Treatment Effects	Nash Comparative Statics
Traveler's Dilemma $1 - F_j - 2Rf_j$	$R (-)$	$R (-)^a$	R (no effect)
Coordination Game $\prod_{j \neq i} (1 - F_j) - c$	$c (-)$ $N (-)$	$c (-)^b$ $N (-)^b$	c (no effect) N (no effect)
Median Effort CG $2F_j(1 - F_k) - c$	$c (-)$ $N (-)$	$c (-)^a$	c (no effect)
Bertrand Game $1 - F_j - xf_j$	$N (-)$	$N (-)^c$	N (no effect)
Imperfect Price Competition $-(1-\alpha)xf_j + [1 - F_j]$	$\alpha (+)$	$\alpha (+)^d$	α (no effect)
Public Goods Game $1 - R_1$	$R_1 (+)$	$R_1 (+)^e$	R_1 (no effect)

^a Capra, *et al.* (1999).

^b Goeree and Holt (1999b).

^c Comparative statics based on numerical calculations; laboratory data from Dufwenberg and Gneezy (1998).

^d Capra, *et al.* (2000).

^e Goeree, Holt, and Laury (1999).

responsive to asymmetries in the costs of directional adjustments, just as the logit equilibrium will be sensitive to expected payoff asymmetries. These learning models include reinforcement learning (Erev and Roth, 1995), where ratios of choice probabilities for two decisions depend on ratios of the cumulated payoffs for those decisions. Even closer to the logit approach is the use of fictitious play or other weighted frequencies of past observed decisions to construct "naive" beliefs, and thereby obtain expected payoffs that are filtered through a logit choice function (e.g. Mookherjee and Sopher, 1997; Fudenberg and Levine, 1998).¹⁷ Indeed, we have used these methods to predict and explain the directional patterns of adjustment in the traveler's dilemma, imperfect price competition, and coordination games (Capra *et al.*, 1999, 2000; Goeree and Holt, 1999a). For example, a version of fictitious play with a single learning (forgetting) parameter,

¹⁷ See Camerer and Ho (1999) for a hybrid model that combines elements of reinforcement and belief learning models.

together with a logit choice function, explains why average claims in Figure 1 fall over time in the $R = 50$ treatment, stay the same in the $R = 20$ treatment, and rise in the $R = 10$ treatment. Simulations using estimated learning and error parameters both track these patterns in the traveler's dilemma (Goeree and Holt, 1999a) and were used to predict the directions of adjustment in the subsequent coordination and imperfect price competition experiments.

On the other hand, learning models that only specify partial or directional adjustments to best responses to previously observed decisions need to be augmented with probabilistic choice, since otherwise they are not sensitive to payoff asymmetries. For example, the best response to previous decisions in the traveler's dilemma game is the other's claim, independent of R , and the best response in the minimum effort coordination game is the minimum of other's efforts, independent of effort cost, so directional best-response and partial adjustment models cannot explain the strong treatment effect in these games unless payoff-based (e.g. logit) errors are included.

Of course, learning models provide lower prediction errors since they use data up to round t to predict behavior in round $t + 1$. Simulations of learning models are quite powerful prediction tools, and we sometimes use them to predict dynamic data patterns for possible treatments before we run them with human subjects (e.g. Capra *et al.*, 2000). These learning and simulation models are *complementary* with equilibrium models, which predict steady state distributions when learning slows down or stops, as in the last five periods in Figure 1. To summarize, learning models are used to predict adjustment patterns and selection in the case of multiple equilibria; whereas equilibrium models are used to predict the steady state distributions and how they shift in response to changes in exogenous parameters.

A second approach to the analysis of behavioral anomalies involves relaxing the standard preference assumptions. Positive contributions in public goods games, for example, are often attributed in part to concerns about others' payoffs. Lottery-choice anomalies have been attributed to non-linear probability weighting. Overbidding relative to Nash predictions has been attributed to risk aversion. In bargaining experiments, the tendency for inequitable offers to be rejected has been attributed to inequity aversion. These generalized preference models will be more convincing if the estimated parameters turn out to be somewhat stable across different

experiments, e.g. a risk aversion explanation of overbidding in private value auctions will be more appealing if similar degrees of risk aversion are estimated from experiments with similar payoff levels. Indeed, Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) have developed models with inequity aversion that are intended to explain behavior in a wide class of games and markets.

Without any added noise, these preference-based theories will suffer from the same problem that plagues the Nash equilibrium with perfect rationality, i.e. that choice tendencies depend on the signs, not on the magnitudes, of payoff differences. For two players, for example, the Fehr and Schmidt model replaces own payoffs, π_i , with a function that depends on whether the other person has a higher or lower payoff, i.e. with $\pi_i - \alpha(\pi_j - \pi_i)$ if $\pi_j - \pi_i > 0$, and with $\pi_i - \beta(\pi_i - \pi_j)$ if $\pi_j - \pi_i < 0$. Here, the "envy" parameter, α , is greater than or equal to the "guilt" parameter, β , which is assumed to be non-negative. Consider the application of this model to the minimum effort game. A unilateral increase from any common effort will lower own payoff due to the increased effort cost, and since the other's payoff is not changed, this will create an envy cost. Conversely, a unilateral decrease will decrease both players' earnings, but the decrease will save on own effort cost, which creates an additional loss due to the guilt effect. Thus the effect of the envy and guilt parameters is to increase deviation losses in both directions, so the set of equilibria is unchanged. As before, any common effort level is an equilibrium with these generalized preferences, irrespective of the effort cost, so this model of inequity aversion would not explain the strong (effort-cost) treatment effects observed in this game.

Fortunately, generalized preference models can be combined with logit and other probabilistic choice models. Fairness and relative earnings considerations are salient in bargaining. In our own work, inequity aversion explains the strong effects of asymmetric money endowments on behavior in alternating offer games, where both the inequity and error parameters estimated from laboratory data are highly significant (Goeree and Holt, 2000a). Similarly, we have found that noise alone does not explain why bidders bid above the Nash equilibrium in private value auctions, but a hybrid model yields highly significant error and risk aversion estimates (Goeree, Holt, and Palfrey, 1999, 2000).

Finally, it is well known that subjects in experiments are sometimes subject to systematic biases, and that complex problems may be dealt with by applying rules of thumb or heuristics.

In a common-value auction, for example, bidders fail to realize that having the high bid contains unfavorable information about the unknown prize value, and overbidding with losses can occur. When there is a single identifiable bias, it should be modeled, perhaps with probabilistic choice appended. When there is not single source of error that can be feasibly modeled, the standard practice is to put the un-modeled effects into the error term.

VII. Conclusion

The standard techniques for characterizing a Nash equilibrium are well developed and understood, but the Nash concept fails to explain the most salient aspects of data from a wide array of laboratory experiments. For example, a large reduction in the penalty rate in a traveler's dilemma does not alter the unique Nash prediction at the lowest claim, but moves the distribution of observed claims toward the opposite end of the set of feasible decisions. Similarly, increases in effort cost sharply reduce distributions of observed efforts in experiments, despite the fact that these cost reductions do not alter the set of Nash equilibria. In both cases, the most salient feature of the data is not being explained by a Nash analysis.

Anomalous experimental results would be less damaging to the Nash paradigm if there were no obvious alternative, but here we argue for an approach is based on probabilistic choice functions that introduce some "noise" that can represent either error and bounded rationality (Rosenthal, 1989) or unobserved preference shocks (McKelvey and Palfrey, 1995). In games, relatively small amounts of noise can have a snowball effect if deviations in the "less risky" direction make further deviations in that direction more attractive. The logit probabilistic choice function allows decision probabilities to be positively but not perfectly related to expected payoffs, and the logit equilibrium incorporates the feedback effects of noisy behavior by requiring belief distributions that determine expected payoffs to match logit choice distributions for those expected payoffs.

The logit equilibrium is essentially a one-parameter generalization of Nash, obtained by not requiring the error parameter to be exactly zero. Since the logit model nests the Nash model, it is straightforward to evaluate them with maximum likelihood estimation based on laboratory data. In fact, any econometric estimation requires some incorporation of random noise, and the quantal response approach provides a structural framework that is natural for games, since it

allows choice probabilities to be affected by the interaction of others' errors and own payoff effects.

The particular logit specification can be generalized or parameterized differently (a power function specification is used in Goeree, Holt, and Palfrey, 1999), but it is difficult to think an alternative error specification that makes more sense. Simply assuming that players make noisy responses to beliefs that others will use their Nash equilibrium strategies ("noisy Nash") is clearly inadequate in games like the Traveler's Dilemma where behavior can deviate so sharply from Nash prediction. One issue is the stability of estimated error rates; we have estimated μ values of 8.5, 7.4, and 6.7 for three of the games discussed above (traveler's dilemma, coordination, and imperfect price competition), but these were games of similar complexity, with the same random matching protocol. The predicted patterns of behavior in these games is somewhat insensitive to error rate changes in this range, and the qualitative comparative statics properties hold for all error rates. Nevertheless, one would expect error rates to be lower for simple individual choice tasks, and higher for complex experiments with asymmetric information and high payoff variability across decisions. One important task for the future will be to develop models of the decision process that allow us predict error rates, i.e. a model of endogenous error rates.

The experience with generalized expected utility theory in the last fifteen years, however, indicates that a generalized approach simply will not be used if it is too messy. The logit analysis, at first glance, is messy; the equilibria are always probability distributions, which complicates analysis of existence and uniqueness. Similarly, comparative statics results pertain to relationships among distributions. In this paper, we provide a general existence result for games with a continuum of decisions, and for auction-like games we show how symmetry, uniqueness, and comparative static results can be obtained from a series of related proofs by contradiction, based on "lens" graphs. The theoretical propositions are then used to characterize the comparative statics properties of the logit equilibria for a series of games. All of the logit comparative statics results in Table 2 are as predicted, and *none* are explained by the relevant Nash equilibrium. Although anomalous from a Nash perspective, these theoretical and experimental results are particularly important because they are consistent with simple economic intuition that deviations from best responses are more likely in the less risky direction.

Finally, the complexity of the theoretical calculations will naturally rise the issue of how

boundedly rational players will learn to conform to these predictions, even as a first approximation. Remember that individuals do not solve the equilibrium differential equations any more than traders in a competitive economy solve the general equilibrium system. Just as traders respond to price signals in a multi-market economy, players in a game may adjust behavior via relatively myopic evolutionary or learning rules that reinforce profitable behavior. The evolutionary model in Anderson, Goeree, and Holt (1999), for example, postulates a population of agents that adjust decisions in the direction of payoff increases, subject to noise (Brownian motion), and the steady state is shown to be a logit equilibrium. Similarly, Goeree and Holt (2000d) discuss conditions under which a naive model of generalized fictitious play will have a steady state that is well approximated by a logit equilibrium; the approximation is better as beliefs become less "overresponsive" to recent experience. In fact, learning models enjoy considerable predictive success (e.g. Camerer and Ho, 1999), especially in terms of explaining the direction of adjustment toward equilibrium (Capra *et al.*, 1999, 2000). This observation may lead some to wonder about the usefulness of equilibrium models like the ones developed in this paper. When you look through the literature, the vast majority of useful predictions in applied work are based on equilibrium models, since (unlike psychologists) economists are primarily interested in behavior in markets and strategic interactions that tend to be repeated. To rely exclusively on learning and evolutionary models is like using generalized cobweb models of market behavior without ever intersecting supply and demand. The theoretical results of this paper are intended to facilitate the application of equilibrium models of bounded rationality with continuous decisions of the type that are commonly used in standard economic models.

Appendix A: Proof of Proposition 1 (Existence of Equilibrium)

Unlike most of the other propositions in the paper, the existence result only requires that expected payoffs are bounded and continuous in others' distribution functions. The latter condition certainly holds for the "local" payoff functions considered in this paper, but is true more generally. We also generalize the logit rule by writing the choice density function as:

$$(A1) \quad f(x) = \frac{g(\pi^e(x)/\mu)}{\int_{\underline{x}}^{\bar{x}} g(\pi^e(y)/\mu) dy},$$

with $g(x)$ a continuous function that is strictly positive everywhere and strictly increasing in x . Note that boundedness of the expected payoff implies that (for $\mu > 0$) there will be no mass points, i.e. the equilibrium distribution functions will be continuous.

Proposition 1. (Existence) For all games in which players' expected payoffs are bounded and continuous in other's distribution functions, there exists a symmetric logit equilibrium for N -player games with a continuum of feasible decisions, and the equilibrium distribution is differentiable.

Proof. Let $\underline{F}(x)$ denote the vector of choice distributions, whose i -th entry, $F_i(x)$, is the distribution of player i , for $i = 1, \dots, n$. Integrating the left and right-hand sides of (A1), yields an operator T that maps a vector $\underline{F}(b)$ into a vector $\underline{TF}(b)$, with components:

$$(A2) \quad TF_i(x) \equiv \frac{\int_{\underline{x}}^x g(\pi_i^e(y)/\mu) dy}{\int_{\underline{x}}^{\bar{x}} g(\pi_i^e(y)/\mu) dy}.$$

The vector of equilibrium distributions is a fixed point of this operator, i.e., $TF_i(b) = F_i(b)$ for all $x \in [\underline{x}, \bar{x}]$, and $i = 1, \dots, n$. As noted above the equilibrium distributions are continuous, so there is no loss of generality in restricting attention to $C[\underline{x}, \bar{x}]$, the set of continuous functions on $[\underline{x}, \bar{x}]$. In particular, consider the set: $S \equiv \{F \in C[\underline{x}, \bar{x}] \mid \|F\| \leq 1\}$, where $\|\cdot\|$ denotes the sup norm. The set S , which includes all continuous cumulative distributions, is an infinite-

dimensional unit ball, and is thus closed and convex. Hence, the n -fold (Cartesian) product $S^n = S \times \dots \times S$, is a closed and convex subset of $C[\underline{x}, \bar{x}] \times \dots \times C[\underline{x}, \bar{x}]$, the set of all continuous n -vector valued functions on $[\underline{x}, \bar{x}]$. This latter space is endowed with the norm $\|\underline{F}\|_n = \max_{i=1..n} \|F_i\|$. The operator T maps elements from S^n to itself, but since S^n is not compact, we cannot rely on Brouwer's fixed point theorem. Instead, we use the following fixed point theorem due to Schauder (see for instance Griffel, 1985):

Schauder's Second Theorem: If S^n is a closed convex subset of a normed space and H^n is relatively compact subset of S^n , then every continuous mapping of S^n to H^n has a fixed point.

To apply the theorem, we need to prove: (i) that $H^n \equiv \{\underline{TF} \in S^n\}$ is relatively compact, and (ii) that T is a continuous mapping from S^n to H^n . The proof of (i) requires showing that elements of H^n are uniformly bounded and equicontinuous on $[\underline{x}, \bar{x}]$. From (A2) it is clear that the mapping $TF_i(x)$ is non-decreasing. So $|TF_i(x)| \leq TF_i(\bar{x}) = 1$ for all $x \in [\underline{x}, \bar{x}]$, $F_i \in S$, and $i = 1, \dots, n$, and elements of H^n are uniformly bounded. To prove equicontinuity of H^n , we must show that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that $|TF_i(x_1) - TF_i(x_2)| < \varepsilon$ whenever $|x_1 - x_2| < \delta$, for all $F_i \in S$, $i = 1, \dots, n$. Consider the difference:

$$(A3) \quad |TF_i(x_1) - TF_i(x_2)| = \frac{\left| \int_{x_1}^{x_2} g(\pi_i^e(y)/\mu) dy \right|}{\int_{\underline{x}}^{\bar{x}} g(\pi_i^e(x)/\mu) dy}.$$

Let π_{\min} and π_{\max} denote the lowest and highest possible payoffs for the game at hand. We can bound the right side of (A3) by:

$$|TF_i(x_1) - TF_i(x_2)| \leq \frac{g(\pi_{\max}) |x_1 - x_2|}{g(\pi_{\min}) (\bar{x} - \underline{x})}.$$

Thus the difference in the values of TF_i is ensured to be less than ε for all $F_i \in S$, $i = 1, \dots, n$, by setting $|x_1 - x_2| < \delta$, where $\delta = \varepsilon (\bar{x} - \underline{x}) g(\pi_{\min}) / g(\pi_{\max})$. Hence, \underline{TF} is equicontinuous for all $\underline{F} \in S^n$.

Finally, consider continuity of T . By assumption the expected payoffs are continuous in others' distributions and g is continuous, so $g(\pi_i^e(x)/\mu)$ is continuous in the others' distributions and so are integrals of $g(\pi_i^e(x)/\mu)$. And since $g(\pi_{\min}/\mu)$ is bounded away from zero, so is the ratio of integrals in (1). Hence T is a continuous mapping from S^n to H^n .

Finally, consider differentiability. Each player's expected payoff function in (8) is a continuous function of x for any vector of distributions of the others' efforts. A player's effort density is a continuous function of expected payoff, and hence each density is a continuous function of x as well. Therefore the distribution functions are continuous, and the expected payoffs are differentiable. The effort densities in are differentiable transformations of expected payoffs, and so these densities are also differentiable. Thus all vectors of densities get mapped into vectors of differentiable densities, and any fixed point must be a vector of differentiable density functions. Q.E.D.

Appendix B: Proof of Proposition 2, parts (c) and (d) (Uniqueness) and Symmetry

Proposition 2. Any symmetric logit equilibrium for a game satisfying the local payoff property is unique if the expected payoff derivative, $\pi^e(F, f, x, \alpha)$ is either: a) strictly decreasing in x , or b) strictly increasing in the common distribution function F , or c) independent of x and strictly decreasing in f , or d) a polynomial expression in F , with no terms involving f or x .

Proof of Parts (c) and (d). Case (c) is based on a *transformation of variables* that allows a more direct application of the logit differential equation (2), since it will produce a graph in which the transformed densities have slopes that are exactly equal to the π^e/μ functions that are so central in these arguments. Notice that raising the height of the horizontal slice in Figure 2 will alter the slopes of the distribution functions at that height. Let y denote the height of the slice in Figure 2, and let $f^*(y)$ denote the density as a function of y . Thus we are considering the transformed density, $f^*(y)$ where $F(x) = y$, and therefore $dx/dy = 1/f(x)$. To derive the slope of the transformed density as a function of the height of the slice, note that $df^*(y)/dy = [df(x)/dx][dx/dy] = f'(x)/f(x) = \pi^e(x)/\mu$, where the final equality follows from the logit differential

equation (2). Thus when we graph the transformed density as a function of y , we get a function with a slope that equals the expected payoff derivative divided by μ . Suppose that there are two symmetric equilibrium distributions denoted by F_1 and F_2 , with the transformed density $f_1^*(y)$ being above $f_2^*(y)$ for low values of y , as shown on the left side of Figure 4. These densities must cross, or the distribution functions will never come together, as they must at \bar{x} , if not before. In any neighborhood to the right of the crossing, it must be the case that $f_1^*(y) < f_2^*(y)$. But since π^{e^*} is assumed to be independent of x and strictly decreasing in the density, it follows that $\pi^{e^*}(y, f_1(x), x_1, \alpha) > \pi^{e^*}(y, f_2(x), x_2, \alpha)$, and therefore, the slope of $f_1^*(y)$ is greater than the slope of $f_2^*(y)$ at all points where f_1 is lower, i.e. to the right of the crossing, which is a contradiction.




Figure 4. A Configuration of Transformed Densities that Yields A Contradiction.

Case (d) is based on a *cone proof*: If the π^{e^*} is a polynomial in F of the form: $A + BF + CF^2 + \dots$, then when it is multiplied by $f(x)$ in (2), we get an expression for $f'(x)$ that can be integrated directly to obtain:

$$(B1) \quad f(x) = f(0) + AF(x) + BF^2/2 + CF(x)^3/3 + \dots$$

Obviously, any solution to (10) is determined by the initial condition, $f(0)$. Suppose that there are two solutions, and without loss of generality, $f_1(0) > f_2(0)$. The two distribution functions must intersect at least once since they must intersect at the upper bound of the support, if not before. Let x^* be the lowest intersection point. Then at any point where the distribution

functions cross, i.e. where $F_1(x^*) = F_2(x^*)$, it follows from (B1) that $f_1(x^*) - f_2(x^*) = f_1(0) - f_2(0) > 0$. This contradicts the fact that $F_1(x)$ must cut $F_2(x)$ "from above" when they cross. Q.E.D.

Proposition 3. (Symmetry) Any logit equilibrium for a game satisfying the local payoff property is necessarily symmetric across players if the expected payoff derivative, $\pi^{e'}(F_j, f_j, x, \alpha)$, is either (a) strictly decreasing in the F_j functions for all other players, or (b) weakly decreasing in the F_j and f_j functions.

Proof. Case (a): First consider the case of two players and suppose, in contradiction, that their equilibrium distributions are not the same. Without loss of generality, assume $F_1(x)$ is lower on some interval, as shown in Figure 3. Any region of divergence between the distribution functions will have a maximum *vertical* difference, as indicated by the vertical line at x^* . The necessary first- and second-order conditions for the distance to be maximized at height x^* are that the slopes of the distribution functions be identical, i.e. $f_1(x^*) = f_2(x^*)$, and that $f_1'(x^*) \geq f_2'(x^*)$. However, since the densities are equal at x^* and $\pi_i^{e'}(F_j, f_j, x, \alpha)$ is decreasing in the *other* player's distribution, F_j , it follows that

$$(D2) \quad \pi_1^{e'}(F_2(x^*), f_2(x^*), x^*, \alpha) < \pi_2^{e'}(F_1(x^*), f_1(x^*), x^*, \alpha).$$

Then the logit differential equation in (2) implies that $f_1'(x^*) < f_2'(x^*)$, which yields the desired contradiction. This proof generalizes to the N player case, since the others' density and distribution functions, evaluated at x^* , would have the same effect on both distribution functions.

Case (b): Consider the asymmetric configuration in Figure 3 again. Just to the right of the left-side crossing of the distribution functions, there must be an interval where $F_2 > F_1$, and $f_2 > f_1$. For any x in this interval, it follows from assumption (b) that $\pi_2^{e'}(F_1(x), f_1(x), x, \alpha) \geq \pi_1^{e'}(F_2(x), f_2(x), x, \alpha)$, and hence that $f_2(x) \pi_2^{e'}(F_1(x), f_1(x), x, \alpha) \geq f_1(x) \pi_1^{e'}(F_2(x), f_2(x), x, \alpha)$. But this latter inequality is, by (2), a condition that $f_2'(x_2) \geq f_1'(x_1)$, so the horizontal distance between F_1 and F_2 will never decrease, in contradiction of the fact that these distributions must meet, at the uppermost value of x if not before. Q.E.D.

Appendix C: Comparative Statics Proofs

Proposition 4 (Comparative Statics for a Symmetric Equilibrium). Suppose that the shift parameter increases marginal expected payoffs, i.e. $\partial \pi^e(F, f, x, \alpha) / \partial \alpha > 0$, for a symmetric game satisfying the local payoff property. Then an increase in α yields stochastically higher logit equilibrium decisions (in the sense of first-degree stochastic dominance) if either a) $\partial \pi^e / \partial x \leq 0$, or b) $\partial \pi^e / \partial F \geq 0$.

Proof. Suppose that $\alpha_2 > \alpha_1$, and let the corresponding symmetric equilibrium distributions be denoted by $F_1(x)$ and $F_2(x)$. The proof requires showing that $F_2(x)$ dominates $F_1(x)$ in the sense of first-degree stochastic dominance, i.e. that $F_1(x) > F_2(x)$ for all interior x . Suppose, in contradiction, that $F_1(x)$ is lower on some interval, as shown in Figure 2. First consider case (a). Any region of divergence between the distribution functions will have a maximum horizontal difference, as indicated by the horizontal dashed line at the height of F^* . As in the proof of Proposition 4(a), the necessary first- and second-order conditions for the distance to be maximized at height $F^* = F_1(x_1) = F_2(x_2)$ are that the slopes of the distribution functions be identical at F^* , i.e. $f_1(x_1) = f_2(x_2)$, and that $f_1'(x_1) \geq f_2'(x_2)$. In order to obtain a contradiction, recall that the distribution functions satisfy the differential equation (2), evaluated at the appropriate level of α :

$$(C1) \quad \mu f_i'(x) = \pi^{e_i}(F_i, f_i, x, \alpha_i) f_i(p), \quad i = 1, 2.$$

Since $F_1(x_1) = F_2(x_2)$ and $f_1(x_1) = f_2(x_2)$, everything except for α_1 and α_2 and the arguments x_1 and x_2 on the right sides of the equations in (C1) are identical, when these equations are evaluated at x_1 and x_2 respectively. The assumption for Case (a), together with $\alpha_2 > \alpha_1$ and $x_2 < x_1$, implies that $f_1'(x_1) < f_2'(x_2)$, which contradicts the second-order condition for the maximum horizontal difference. Next consider Case (b), in which the payoff derivative is non-decreasing in the distribution function. Any region of divergence between the distribution functions will have a maximum *vertical* difference, as indicated by the vertical dashed line at x^* Figure 3, where the two distributions for $\alpha_2 > \alpha_1$ are now denoted by F_1 and F_2 . The necessary first- and second-order conditions for the distance to be maximized at height x^* are that the slopes of the

distribution functions be identical, i.e. $f_1(x^*) = f_2(x^*)$, and that $f_1'(x^*) \geq f_2'(x^*)$. However, since $\pi^{e'}(F_j, f_j, x, \alpha_j)$ is increasing in F_j and $F_1(x^*) < F_2(x^*)$ by assumption, it follows that

$$(C2) \quad \pi^{e'}(F_1(x^*), f_1(x^*), x^*, \alpha_1) < \pi^{e'}(F_2(x^*), f_2(x^*), x^*, \alpha_2).$$

Then the logit differential equation in (2) implies that $f_1'(x^*) < f_2'(x^*)$, which yields the desired contradiction. These arguments apply to the N player case, since by symmetry, the density and distribution functions of all players are identical and have the same value at x^* . Q.E.D.

Proposition 6 (Comparative Statics for Asymmetric Payoffs). Suppose that the shift parameter increases marginal expected payoffs, i.e. $\partial \pi^{e'}(F, f, x, \alpha) / \partial \alpha_i > 0$, and let $\alpha_2 > \alpha_1$ in a game satisfying the local payoff property. Then player 2's logit equilibrium distribution of decisions is stochastically higher than that of player 1, i.e. the distribution function for player 2 is lower at each interior value of x , if either: a) $\partial \pi^{e'} / \partial x \leq 0$, or b) $\partial \pi^{e'} / \partial F \leq 0$.

Proof. Suppose that $\alpha_2 > \alpha_1$, and let the corresponding equilibrium distributions be denoted by $F_1(x)$ and $F_2(x)$ for players 1 and 2 respectively. We wish to show that $F_1(x) > F_2(x)$ for all interior x (i.e., $F_2(x)$ dominates $F_1(x)$ in the sense of first-degree stochastic dominance). Suppose not, so that $F_1(x)$ is lower on some interval, as per Figure 3. As in the proofs of Proposition 1, the necessary first- and second-order conditions for the vertical distance $F_2(x) - F_1(x)$ to be maximized at action x^* imply that $f_1(x^*) = f_2(x^*)$, and that $f_1'(x^*) \geq f_2'(x^*)$. This in turn implies that we must have $\pi_1^{e'}(F_2, f_2, x^*, \alpha_1) \geq \pi_2^{e'}(F_1, f_1, x^*, \alpha_2)$; but, since $\alpha_2 > \alpha_1$ and $\pi_i^{e'}$ is increasing in α , it follows from the assumption in case (a) that this can only hold if $F_1(x^*) > F_2(x^*)$, contradicting the original condition. Case (b) is proved with a horizontal lens argument based on Figure 2. This proposition also applies to the case of more than two players, since the effects of others' densities and distributions affect both $\pi_1^{e'}(F_2, f_2, x^*, \alpha_1) \geq \pi_2^{e'}(F_1, f_1, x^*, \alpha_2)$ in the same manner, when evaluated at the same value of x . Q.E.D.

Appendix D: Single Peakedness

Proposition 7. (Single Peakedness) If the logit equilibrium for a game satisfying the local payoff property is symmetric and the expected payoff derivative, $\pi^{e'}(F, f, x, \alpha)$, non-increasing in x and strictly decreasing in the common F function, then the equilibrium density that solves (2) will be single peaked.

Proof. The assumptions, together with Proposition 3, imply that the equilibrium is symmetric across players, so we will drop the player subscripts from the notation that follows. Since the density in (1) is proportional to an exponential function of expected payoffs, we need to show that the expected payoff function is concave in x . To do this, consider the second derivative of expected payoff with respect to x , i.e. the derivative of $\pi^{e'}(F(x), f(x), x, \alpha)$ with respect to x , taking into account the direct and indirect effects through arguments in the density and distribution functions. This derivative is:

$$(D1) \quad d\pi^{e'}/dx = \partial\pi^{e'}/\partial F \cdot f(x) + \partial\pi^{e'}/\partial f \cdot f'(x) + \partial\pi^{e'}/\partial x.$$

The first and third terms on the right side of (D1) are negative by assumption, with the first term being strictly negative, and the logit differential equation (2) implies that the second term is zero at any stationary point with $\pi^{e'} = 0$. It follows that the right side of (D1) is negative at any stationary point of the expected payoff function, and therefore, that any stationary point will be a local maximum. Q.E.D.

References

- Akerlof and Yellin (1985) "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?," *American Economic Review*, 75, 708-720.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998a) "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction," *Journal of Political Economy*, 106(4), August 1998, 828-853.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998b) "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games," *Journal of Public Economics*, 70, 297-323.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1999) "Stochastic Game Theory: Adjustment and Equilibrium with Bounded Rationality," working paper, University of Virginia.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2000) "Minimum-Effort Coordination Games: An Equilibrium Analysis of Bounded Rationality," *Games and Economic Behavior*, forthcoming.
- Basu, Kaushik (1994) "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory," *American Economic Review*, 84(2), 391-395.
- Baye, Michael R. and John Morgan (1999) "Bounded Rationality in Homogeneous Product Pricing Games," Discussion Paper, Indiana University.
- Binmore, Ken and Larry Samuelson (1999) "Evolutionary Drift and Equilibrium Selection," *Review of Economic Studies*, 66, 363-393.
- Bolton, Gary E. and Axel Ockenfels (2000) "A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90(1), 166-193.
- Camerer, Colin and Teck-Hua Ho (1999) "Experience Weighted Attraction Learning in Normal-Form Games," *Econometrica*, 67, 827-874.
- 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), June, 678-690.
- Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (2000) "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition,"

- International Economic Review*, forthcoming.
- Carlsson, H. and E. van Damme (1993). "Global Games and Equilibrium Selection," *Econometrica*, 61(5) September, 989-1018.
- Chen, Hsiao-Chi, James W. Friedman, and Jacques-François Thisse (1996) "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, 18, 32-54.
- Crawford, Vince and Miguel Costa-Gomez (2000) "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, forthcoming.
- Dufwenberg, Martin and Uri Gneezy (2000) "Price Competition and Market Concentration: An Experimental Study," *International Journal of Industrial Organization*, 18(1), 7-22.
- Erev, Ido and Alvin E. Roth (1998) "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88(4), 848-881.
- Fehr, Ernst and Klaus Schmidt (1999) "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, August, 114(3), 817-868.
- Fudenberg, Drew, and David K. Levine (1998) *Learning in Games*, Cambridge, MIT Press.
- Goeree, Jacob K., Simon P. Anderson, and Charles A. Holt (1998) "The War of Attrition with Noisy Players," in *Advances in Applied Microeconomics, Volume 7*, M.R. Baye, ed., Greenwich, Conn.: JAI Press, 15-29.
- Goeree, Jacob K. and Charles A. Holt (1999a) "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, September, 10564-10567.
- Goeree, Jacob K. and Charles A. Holt (1999b) "An Experimental Study of Costly Coordination," Discussion Paper, University of Virginia.
- Goeree, Jacob K. and Charles A. Holt (2000a) "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," *European Economic Review*, 44, 1079-1089.
- Goeree, Jacob K. and Charles A. Holt (2000b) "Models of Noisy Introspection," Discussion Paper, University of Virginia.

- Goeree, Jacob K. and Charles A. Holt (2000c) "An Explanation of Anomalous Behavior in Binary-Choice Games: Entry, Voting, Public Goods, and the Volunteers' Dilemma," Discussion Paper, University of Virginia.
- Goeree, Jacob K. and Charles A. Holt (2000d) "Stochastic Learning Equilibrium," Discussion Paper presented at the June 2000 Economic Science Association Meetings in New York.
- Goeree, Jacob K. and Charles A. Holt (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, forthcoming.
- Goeree, Jacob K., Charles A. Holt, and Susan K. Laury (1999) "Altruism and Noisy Behavior in One-Shot Public Goods Experiments," Discussion Paper.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (1999) "Quantal Response Equilibrium and Overbidding in Private Value Auctions," Discussion Paper, Caltech.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2000) "Risk Averse Behavior in Asymmetric Matching Pennies Games," Discussion Paper, University of Virginia.
- Green, Donald P. and Ian Shapiro (1994) *Pathologies of Rational Choice Theory*, New Haven: Yale University Press.
- Griffel, D. H. (1985) *Applied Functional Analysis*, Chichester: Ellis Horwood.
- Harrison, Glenn W. and Jack Hirshleifer (1989) "An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods," *Journal of Political Economy*, 97(1), 201-225.
- Kagel, John and Alvin Roth, eds. (1995) *Handbook of Experimental Economics*, Princeton: Princeton University Press.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- McKelvey, Richard D. and Thomas R. Palfrey (1998) "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1(1), 9-41.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber (2000) "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2×2 Games with Unique Mixed Strategy Equilibria," *Journal of Economic Behavior and Organization*, 42, 523-548.
- Mookherjee, Dilip, and Barry Sopher (1997) "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 19, April, 97-132.

- Ochs, Jack (1995) "Coordination Problems," in J. Kagel and A. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press, 1995, 195-249.
- Ostrom, Elinor (1998) "A Behavioral Approach to Rational Choice Theory of Collective Action," *American Political Science Review*, 92, 1-22.
- Radner, Roy (1980) "Collusive Behavior in Oligopolies with Long But Finite Lives," *Journal of Economic Theory*, 22, 136-156.
- Reynolds, Stanley S. (1999) "Sequential Bargaining with Asymmetric Information: The Role of Quantal Response Equilibrium in Explaining Experimental Results," Discussion Paper, University of Arizona.
- Rosenthal, Robert W. (1989) "A Bounded Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273-292.
- Roth, Alvin, and Ido Erev (1995) "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164-212.
- Selten, R. (1975) "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25-55.
- Selten, R. and J. Buchta (1994) "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," in Budescu, I. E. D. and Zwick, R., eds., *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, Hillside, N.J.
- Simon, Herbert A. (1957) *Models of Man*, New York: John Wiley & Sons.
- Smith, Vernon L. and James M. Walker (1993) "Rewards, Experience, and Decision Costs in First-Price Auctions," *Economic Inquiry*, 31, 237-245.
- Stahl, Dale O. and Paul W. Wilson (1995) "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218-254.
- 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.
- von Neumann, John, and Oscar Morgenstern (1944) *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.