

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

QUANTAL RESPONSE EQUILIBRIA FOR EXTENSIVE FORM GAMES

Richard D. McKelvey

Thomas R. Palfrey



SOCIAL SCIENCE WORKING PAPER 947

December 1995

Quantal Response Equilibria for Extensive Form Games

Richard D. McKelvey

Thomas R. Palfrey

Abstract

This paper investigates the use of standard econometric models for quantal choice to study equilibria of extensive form games. Players make choices based on a quantal choice model, and assume other players do so as well. We define an Agent Quantal Response Equilibrium (AQRE), which applies QRE to the agent normal form of an extensive form game and imposes a statistical version of sequential rationality. We also define a parametric specification, called logit-AQRE, in which quantal choice probabilities are given by logit response functions.

AQRE makes predictions that contradict the invariance principle in systematic ways. We show that these predictions match up with some experimental findings by Schotter, Weigelt and Wilson (1993) about the play of games that differ only with respect to inessential transformations of the extensive form. The logit-AQRE also implies a unique selection from the set of subgame perfect equilibria in generic extensive form games. We examine data from signalling game experiments by Banks, Camerer, and Porter (1994) and Brandts and Holt (1993). We find that the logit-AQRE selection applied to these games succeeds in predicting patterns of behavior observed in these experiments, even when our prediction conflicts with more standard equilibrium refinements, such as the intuitive criterion. We also reexamine data from the McKelvey and Palfrey (1992) centipede experiment.

Quantal Response Equilibria for Extensive Form Games

Richard D. McKelvey* Thomas R. Palfrey

December 12, 1995

1 Introduction

In a companion paper (McKelvey and Palfrey, 1995) we introduced the idea of Quantal Response Equilibrium (QRE), and applied this concept to noncooperative games in normal form. Quantal responses are smoothed-out best responses, in the sense that players are more likely to choose better strategies than worse strategies, but do not play a best response with probability one. The idea has its origins in statistical limited dependent variable models such as discrete choice (in economics and psychology) and stimulus/dosage response and bioassay (in biology and medical research), and is an essential ingredient of the work on bounded rationality in games by Rosenthal (1991) and control cost models of equilibrium selection by van Damme (1987). The added complication in applying the concept to game theory, in contrast to individual choice, is that the choice probabilities of the players have an important interactive component, since they are simultaneously determined in equilibrium.

QRE is an internally consistent equilibrium model, in the sense that the quantal response functions are based on the equilibrium probability distribution of the opponents' strategy choices, rather than simply arbitrary beliefs the players could have about those probabilities.¹

An interesting property of QRE is that systematic deviations from Nash equilibrium may be predicted even without introducing systematic features to the error structure. The systematic effects of the statistical disturbances are (indirect) equilibrium phenom-

*This paper has benefited from comments by participants at the Conference on Learning in Games (Texas A&M University, February 1994) and the European Meetings of the Econometric Society (Maastricht, September 1994). We thank the participants for their comments and also thank Peter Coughlan, Mark Fey, Eugene Grayver, and Rob Weber for their research assistance. The financial support of the National Science Foundation (Grant #SBR-9223701) is gratefully acknowledged.

¹See also Chen, Friedman and Thisse (1995) who study learning dynamics in the context of a similar quantal response equilibria model, which they call "boundedly rational Nash equilibria".

ena that arise because of the strategic response of the players to the noisy environment. This leads to what can be viewed as a statistical generalization of Nash equilibrium.

In the earlier paper we show that for normal form games there is a correspondence between QRE and Bayesian equilibria of what Harsanyi (1973) calls randomly disturbed games. Specifically, we show that a particular parametrization of quantal response equilibrium, which we call the logit equilibrium, corresponds to the Bayesian equilibria of the game with a vector of independent log Weibull-distributed disturbances added to the payoffs of the players, with each individual's disturbances being private information to that player. That paper then establishes some properties of QRE in finite n -person games, and re-examines a diverse collection of experimental bimatrix games that have unique logit (and Nash) equilibria. Statistical analysis of that data was carried out using standard maximum likelihood techniques in a structural model that is directly implied by the quantal response model. We estimate the variance of the Weibull disturbances for that data and compare the predicted frequencies of strategy choices with the actual frequencies. We find that the theory tracks the data fairly well, and can account for some of the observed systematic deviations from the predictions of the exact Nash equilibrium model.

This paper reconsiders Nash equilibrium in extensive form games from the statistical point of view, and extends the concept of logit equilibrium to the extensive form representation. A number of interesting theoretical and empirical findings emerge. On the theoretical side, the new formulation predicts systematic violations of what is known in the foundational game theory literature as *invariance*. It has been argued vigorously by some (e.g. Kohlberg and Mertens 1986) that “the reduced normal form captures all the relevant information for decision purposes,” leading to the traditional game theory view that any good theoretical equilibrium concept for noncooperative games must satisfy invariance. In our statistical theory of decisionmaking in games the compelling arguments for invariance fall apart. In fact, it is easy to construct examples where we predict systematic differences in the predicted patterns of play depending on which “equivalent” version of a game is played. Moreover, the intuition behind these differences is very sensible.

Besides discriminating between different versions of a game that have equivalent strategic form representations, our model also makes different predictions depending upon whether the game is played in its “agent” normal form or its more traditional normal form, and makes different predictions between the expected patterns of play in the normal form of the game and the reduced normal form of the same game. Loosely speaking, the intuitive reason our model predicts “representation dependence” is that the error structure introduces private information in particular ways that depend on the exact details of how the game is actually played. Because of the specific statistical predictions that are implied by the dependence, many aspects of the quantal response equilibrium theory can be tested directly by applying standard maximum likelihood techniques to data generated in controlled laboratory experiments. For computational reasons, we focus on a version of QRE which is similar in spirit to the agent model of how an extensive form

game is played, and so we call this the Agent Quantal Response Equilibrium (AQRE).

In addition to predicting systematic violations of the invariance principle, the logit-AQRE model implies a unique selection from the set of Nash equilibria. This selection is defined by the connected component of the logit-AQRE correspondence, as explained in McKelvey and Palfrey (1995). This selection thus generates a unique prediction of a strictly positive probability distribution over play paths for every value of the logit response parameter, for almost any finite extensive form game. As demonstrated in our earlier paper, this selection may be inconsistent with trembling had perfection as applied to the reduced normal form of the game, as well as all other refinements of which we are aware.

In signaling games, there is a well-known multiple equilibrium problem to which a great deal of effort has been directed by modern game theorists, generating a vast literature on the general problem of deductively-based refinements. Several experiments have been conducted to test whether the refinements developed to distinguish these equilibria are useful in helping to predict which equilibria (if any) are more likely to be played. This is a natural setting in which to apply AQRE, since this equilibrium concept does not have the problem of trying to define behavior off the equilibrium path. This equilibrium concept also makes multiple equilibrium predictions, but the logit-AQRE has a natural refinement that generically selects a unique equilibrium in a way that is much different from traditional (deductive) refinement arguments. In addition to predicting the patterns of play “on the Nash equilibrium path,” this equilibrium concept also predicts patterns of play off the Nash equilibrium path as well. Some of the anomalies (vis-a-vis more standard theories) uncovered in these experiments have to do with behavior off the equilibrium path, and our model successfully accounts for these.

In past centipede game experiments conducted by us (McKelvey and Palfrey [1992]) frequent violations of Nash predicted play were observed. This fact and the fact that the observed outcomes were substantial Pareto improvements over Nash play, were rationalized by a model in which some of the players have altruistic preferences. Using arguments based on reputation building (Kreps and Wilson [1982]) it is shown that frequent violations of Nash behavior (similar to what was observed) can be accounted for even if altruistic players are “rare.” While this explanation of the data does account pretty well for most of the salient features of that particular dataset, it is clearly ad hoc. The explanation involves the invention, or assumption, of a “deviant type” who systematically violates Nash behavior in exactly the direction observed in the data. A preferable explanation would be able to account for this data without resorting to such “ad hocery.” The AQRE provides a framework for doing exactly that, and also has the desirable feature of being applicable to arbitrary games without necessitating the invention of systematically deviant types, tailored to the peculiarities of specific games.

Accordingly, we also reexamine the centipede data² in the context of the logit-AQRE. We find that a two-parameter specification of AQRE can account for the same qualitative

²Zauner (1993) has independently conducted a reexamination of the centipede data using a similar model, but with different distributional assumptions about the error structure.

violations of Nash behavior as the five-parameter altruism model, although the exact fit to the data is not as tight. In another paper, we have presented some new data (contradictory to the altruism explanation) from much different centipede-type experiments where all outcomes are Pareto optimal (Fey, McKelvey, and Palfrey 1994). The logit-AQRE model also fits that data better than other competing models.

The remainder of the paper is divided into four sections. The first section sets out the formal structure of the AQRE model for extensive form games. We also introduce a parametric version based on the logit model, that we later use for estimation. The second section presents some examples, based on the chain-store paradox stage game, demonstrating one kind of violations of the invariance principle that are implied by our model, and looks at some recent experimental data reported in Schotter, Weigelt and Wilson (1994) that are consistent with our predictions. The third section explains the refinement for signaling games based on logit-AQRE, and examines data from experimental signaling games conducted by Brandts and Holt (1993) and by Banks, Camerer and Porter (1993). The fourth section reexamines the data from the centipede game, and discusses similar findings by Zauner (1993), and Fey, McKelvey, and Palfrey (1994). We conclude with a few brief remarks about possible extensions of the QRE model.

2 Definitions and Notation

Let N be a finite set of $n + 1$ *players*, one of whom (player “0”) is designated *chance*, X be a finite set of *outcomes*, and let A be a finite set of *actions*. Let $\Gamma = (\Gamma, Q)$ be a *topological tree* with Γ the set of *nodes*, and $Q \subseteq \Gamma \times \Gamma$ a binary relation on Γ representing *branches*, and let Q^* be the transitive closure of Q . If vQv' we say v *immediately follows* v' (or v' *immediately precedes* v). If vQ^*v' we say v *follows* v' (or v' *precedes* v). Write $Q(v)$ and $Q^*(v)$ for the set of immediate followers and the set of followers of $v \in \Gamma$, respectively. We assume that Q is an asymmetric, acyclic binary relation on Γ such that every node $v \in \Gamma$ follows at most one node, and such that there is exactly one node, $v^* \in \Gamma$ which follows no node. The element v^* is called the *root node*. A node $v \in \Gamma$ is *terminal* if $Q(v) = \emptyset$, otherwise it is *non terminal*. Let Γ^t and Γ^0 represent the set of terminal and non terminal nodes of Γ , respectively.

We define an extensive form game $\mathcal{G}(N, X, A, \Gamma)$ by defining: an *outcome function* $\psi : \Gamma^t \rightarrow X$, a *partition function* $P : \Gamma^0 \rightarrow N \times I$ (here I is the integers), an *index function* $\phi : \Gamma \rightarrow A$ which is 1-1 on $Q(v)$ for each $v \in \Gamma^0$, and which satisfies $\phi(Q(v)) = \phi(Q(v'))$ whenever $P(v) = P(v')$, and a *probability function* $\lambda : \Gamma^0 \rightarrow \mathcal{R}$ which is a probability function on $Q(v)$ for each $v \in \Gamma^0$ with $P_1(v) = 0$.

For any $v \in \Gamma^0$, define $A(v) = \phi(Q(v))$ to be the set of *actions* available at v . For each $v \in \Gamma^0$ define $h(v) = \{v' \in \Gamma^0 : P(v) = P(v')\}$ to be the *information set* containing v . Let $H = \{h : h = h(v) \text{ for some } v \in \Gamma^0\}$ be the set of all information sets. If $h = h(v)$ write $P(h) = P(v)$ and $A(h) = A(v)$. Write $H_i = \{h \in H : P_1(h) = i\}$ to be the set of information sets for player i . If $P(v) = (i, j)$, write $h_i^j = h(v)$. An action $a \in A(h)$ is said

to precede a node $v \in \Gamma$ if there is an $x \in h$ and $y \in Q(x)$ with $v \in Q^*(y)$ and $\phi(y) = a$.

An extensive form game \mathcal{G} is said to have *perfect recall* if for every player $i \in N$, all information sets $h, h' \in H_i$, every action $a \in A(h)$, and all nodes $x, y \in h'$, a precedes x iff a precedes y . In a game of perfect recall, a player remembers what he previously knew and what he previously did. We assume throughout that \mathcal{G} has perfect recall.

A *behavior strategy* for player $i \in N$ is a function $b_i : H_i \rightarrow \mathcal{M}(A)$ satisfying $b_i(h) \in \mathcal{M}(A(h))$ for all $h \in H_i$. Here $\mathcal{M}(C)$ indicates the set of all probability measures over the set C . Again, for $h \in H_i$ we use the shorthand $b_{ih} = b_i(h)$, and $b_{ij} = b_i(h_i^j)$. For each $a \in A(h)$, $b_{iha} = b_{ih}(a)$ denotes the probability of action a . Let B_i denote the set of all behavior strategies for player i , and $B = \prod_{i \in N} B_i$ be the set of behavior strategy n -tuples. If $v \in h \in H_0$ is a chance node, and $v' \in Q(v)$, then we define b by $b_h(v') = \lambda_v(v')$. We use the notation $B_{-i} = \prod_{j \neq i} B_j$, and write elements of B in the form $b = (b_i, b_{-i}) \in B_i \times B_{-i} = B$ when we want to focus on a particular player $i \in N$. A pure strategy is a behavior strategy that assigns a degenerate probability distribution to each information set assigned to a player.

Let B° denote the interior of B . Each behavior strategy n -tuple $b \in B^\circ$ determines a strictly positive *realization probability* $\rho(v|b)$ for each node $v \in \Gamma$, as follows: If $v = v^*$ is the root node, then define $\rho(v|b) = 1$. Otherwise, find v' with $v \in Q(v')$. Since $v' \in \Gamma^0$, $v' \in h \in H_i$ for some $i \in N$, and define $\rho(v|b) = \rho(v')b_h(v')$. For any $h \in H$, define $\rho(h|b) = \sum_{v \in H} \rho(v|b)$. Also, for any $h \in H$, define $Q^*(h) = \cup\{Q^*(v) : v \in h\}$. Then we can define a *conditional realization probability*, $\rho(v|h, b)$ on $Q^*(h)$ by $\rho(v|h, b)\rho(h|b) = \rho(v|b)$. Note that for any $b \in B^\circ$, that $\rho(v|b)$ defines a strictly positive probability measure over the terminal nodes. Similarly, $\rho(v|h, b)$ defines a probability measure over $\Gamma^t \cap Q^*(h)$.

For any $i \in N$ and $b \in B^\circ$, define the payoff function $u_i : B^\circ \rightarrow \mathcal{R}$ by

$$u_i(b) = \sum_{v \in \Gamma^t} \rho(v|b)u_i(v).$$

For any $h \in H$, $i \in N$ and $b \in B^\circ$, define the conditional payoff function $u_i : B^\circ \rightarrow \mathcal{R}$ by

$$u_i(b|h) = \sum_{v \in \Gamma^t \cap Q^*(h)} \rho(v|h, b)u_i(v).$$

For any interior behavior strategy $b \in B^\circ$, any i and any information set $h_i^j \in H_i$, we denote $u_i(a, b|h_i^j)$ the conditional payoff to i of playing action $a \in A(h_i^j)$ at h_i^j with probability one, and playing b_i elsewhere.

Define a *valuation function* $x : H \rightarrow \mathcal{R}^K$, where for each $h \in H$, $x(h) \in \mathcal{R}^{|A(h)|}$. We write $x_{ha} = x(h)(a)$.

Write $X_{ij} = \mathcal{R}^{|A(h_i^j)|}$ to represent the space of possible expected payoffs for actions that player i might adopt at information set $h_i^j \in H_i$, and $X_i = \prod_j X_{ij}$, and $X = \prod_{i \in N} X_i$. We define the function $\bar{u} : B^\circ \rightarrow X$ by

$$\bar{u}(b) = (\bar{u}_1(b), \dots, \bar{u}_n(b)),$$

where

$$\bar{u}_{ija}(b) = u_i(a, b|h_i^j).$$

Thus $\bar{u}(b)$ is a vector whose components are the conditional expected payoff to each available action a of each information set j of each player i .

3 Extensive Form Quantal Response Equilibrium

For each i , h_i^j , $a \in A(h_i^j)$, let ε_{ija} be a random variable, which represents i 's *payoff disturbance* from choosing a at information set h_i^j . For any $b \in B^\circ$, let

$$\hat{u}_{ija} = \bar{u}_{ija}(b) + \varepsilon_{ija}.$$

We assume that ε is private information to the players. So player i observes ε_{ija} but the other players (or the econometrician) do not.

Let ε_{ij} be the vector of i 's errors at h_i^j , let ε_i be the entire error vector across i 's information sets, and let $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$ be the vector of errors for all players. We assume ε is an absolutely continuous random vector (with respect to Lebesgue measure), distributed according to a joint distribution with density function $f(\varepsilon)$. We also assume that the ε_{ij} are statistically independent and that $\mathcal{E}(\varepsilon_{ija})$ exists for all $i, j, a \in A(h_i^j)$. Any probability density function, f , satisfying the above assumptions is called *admissible*.

We assume that at each information set, h_j^i , player i selects an action that maximizes \hat{u}_{ija} . For any action $a \in A(h_j^i)$ and any $\bar{u} \in X$ define

$$R_{ija}(\bar{u}) = \{\varepsilon : \bar{u}_{i \cdot a} + \varepsilon_{ija} \geq \bar{u}_{ij\hat{a}} + \varepsilon_{ij\hat{a}} \quad \forall \hat{a} \in A(h_j^i)\}$$

and

$$\sigma_{ija}(\bar{u}) = \int_{R_{ija}(\bar{u})} f(\varepsilon) d\varepsilon$$

Definition 1 For any extensive form game \mathcal{G} and error structure $f(\varepsilon)$, a behavioral strategy $b^* \in B$ is an **agent quantal response equilibrium (AQRE)** if it is a fixed point in B of $\sigma \circ \bar{u}$. It is a vector $b^* \in B^\circ$ such that for all $i \in N$, $1 \leq j \leq J_i$, $a \in A(h_j^i)$, $b_{ija}^* = \sigma_{ija}(\bar{u}(b^*))$

That is, b^* is a quantal response equilibrium if, when all players are choosing their best response at every information set h_j^i , taking b_{-ij}^* as given, this generates a behavior strategy that is exactly the same as b^* .

This definition assumes the “agent model” of play in an extensive form game, where different information sets of a given player are assumed to be played out by different agents, all of whom share the same payoff function. In our model, each agent ij simply chooses the maximum of \hat{u}_{ija} at information set h_j^i and acts independently of the other agents of the same player. For this reason we refer to the equilibrium defined above as the *agent quantal response equilibrium (AQRE)*.

Theorem 1 *For any admissible f , an AQRE exists.*

Proof: Define $\phi = \sigma \circ \bar{u} : B \mapsto B$. Note that ϕ is continuous. Extend ϕ to have domain B by

$$\phi_{ij}(b) = \text{co}\{\lim_{t \rightarrow \infty} \phi_{ij}(b_t) : \{b_t\} \subseteq B^\circ \text{ and } \lim_{t \rightarrow \infty} b_t = b\}.$$

Then B is compact, convex, and $\phi : B \mapsto B$ is convex valued and upper hemicontinuous. Hence, by the Kakutani fixed point theorem, ϕ has a fixed point, $b^* \in B$. It is easy to show that there can be no fixed point on the boundary. Hence, there is a $b^* \in B^\circ$ with $\sigma \circ \bar{u}(b^*) = b^*$. \blacksquare

The above theorem guarantees existence for AQRE, and the definition guarantees that the equilibrium behavioral strategies place positive probability on every available action in every information set of every player. The equilibrium correspondence has the same “nice” topological properties as the Nash equilibrium correspondence.

The above model can be generalized by dropping the independence assumption, or by specifying an observation function $\mathcal{O}(h_j^i)$, where for each i , h_j^i , $\mathcal{O}(h_j^i)$ specifies a signal i receives about ε at information set h_j^i , rather than just assuming i observes only \bar{u}_{ij} at h_{ij} . While these changes would require some additional notation, the basic ideas and properties of the quantal response model would not change. Some models of $\mathcal{O}(h_j^i)$ other than the one we explore here might be interesting in their own right. We leave these issues as a subject of future work.

4 The Logistic AQRE

In the remainder of the paper, we focus mainly on a specialized version of the model, where each ε_{ija} is independently and identically distributed according to the type I extreme value, (or log Weibull) distribution with cumulative density $F(\varepsilon_{ija}) = e^{-e^{-\lambda \varepsilon_{ija}}}$. This distribution has mean of $\frac{\gamma}{\lambda}$, where γ is the Euler constant (0.577) and variance of $\frac{\pi^2}{6\lambda^2}$. So λ is proportional to the precision of F . This distribution of the disturbances leads to choice probabilities following a multinomial logit distribution (see eg., McFadden, 1975). In particular, at h_j^i , b_j^i (in the AQRE) is given by

$$b_j^i(a) = \frac{e^{\lambda \bar{u}_{ija}(b)}}{\sum_{a' \in A(h_j^i)} e^{\lambda \bar{u}_{ija'}(b)}}.$$

A *logit-AQRE* is any solution to this set of k equations (one equation for each action in each information set of each agent).

For each $\lambda \in [0, \infty)$, define the *logit-AQRE correspondence* $b^* : \mathcal{R}_+ \mapsto B^\circ$ as the set of logit-AQRE behavioral strategies when the precision of F equals λ :

$$b^*(\lambda) = \{b \in B^\circ : b_j^i(a) = \frac{e^{\lambda \bar{u}_{ija}(b)}}{\sum_{a' \in A(h_j^i)} e^{\lambda \bar{u}_{ija'}(b)}}\}.$$

We establish several properties of the logit-AQRE correspondence as a function of λ . We already established from Theorem 1 of the previous section that $b^*(\lambda) \neq \emptyset, \forall \lambda$. That is, at least one logit-AQRE exists for every value of λ . A second property is that limit points of $b^*(\cdot)$ as λ goes to 0 are behavioral strategy Nash equilibria of the game. Both of these properties follow from the same arguments as in McKelvey and Palfrey (1993).

A third property is that limit points of $b^*(\lambda)$ are not only Nash equilibrium behavior strategy profiles, but are sequential equilibrium strategy profiles:

Theorem 2 *For every finite extensive form game, every limit point of a sequence of logit-AQREs with λ going to infinity corresponds to the strategy of a sequential equilibrium assessment of the game.*

Proof: Let b_∞^* be a limit point of $b^*(\lambda)$. Then there exists consistent beliefs μ_∞^* (assignments of a probability distribution over the nodes at each information set that satisfy the Kreps-Wilson (1982) consistency condition) such that, under those beliefs, b_∞^* specifies optimal behavior for every continuation game. That is b_∞^* is sequentially rational given μ_∞^* and μ_∞^* is consistent with b_∞^* . This is proved below.

Take any sequence $\{\lambda_k\}_{k=1}^\infty$ and $\{b_k^*\}_{k=1}^\infty$ such that $\lim_{k \rightarrow \infty} \lambda_k = \infty$ and $\lim_{k \rightarrow \infty} b_k^* = b_\infty^*$ and b_k^* is a logit equilibrium for λ_k for each k . First note that for all k and for all information sets h , $b^*(a) > 0 \forall a \in A(h)$. Consequently every node is reached with positive probability. Therefore, by Bayes rule, for each k , b_k^* uniquely defines a set of beliefs, μ_k^* , over the nodes of every information set. Moreover, since μ varies continuously with b , there is a unique limit μ_∞^* . Therefore μ_∞^* are consistent beliefs with respect to b_∞^* . What remains to be shown is that $b_\infty^*(h)$ is optimal given $\mu_\infty^*(h)$ for all h . (If so, then $(b_\infty^*, \mu_\infty^*)$ is a sequential equilibrium, so that b_∞^* is a sequential equilibrium strategy.) Suppose not, then there exists $\varepsilon > 0$ such that for some h_j^i there is some pair of actions $\{a, a'\} \in A(h_j^i)$ such that $b_{j\infty}^{*i}(a) > 0$ but $\bar{u}_{ija'}(b_\infty^*) - \bar{u}_{ija}(b_\infty^*) > \varepsilon$. But then there must exist $K > 0$ such that

$$\bar{u}_{ija'}(b_k^*) - \bar{u}_{ija}(b_k^*) \geq \frac{\varepsilon}{2} \quad \forall k \geq K.$$

Therefore, for all $k \geq K$,

$$\frac{b_{jk}^{*i}(a)}{b_{jk}^{*i}(a')} \leq e^{-\lambda_k \frac{\varepsilon}{2}} \quad \forall k \geq K.$$

So,

$$\begin{aligned} b_{j\infty}^{*i}(a) &\equiv \lim_{k \rightarrow \infty} b_{jk}^{*i}(a) \leq \lim_{k \rightarrow \infty} b_{jk}^{*i}(a') e^{-\lambda_k \frac{\varepsilon}{2}} \\ &= 0, \end{aligned}$$

which contradicts $b_{j\infty}^{*i}(a) > 0$. ■

The next result establishes uniqueness of the AQRE for games of perfect information.

Theorem 3 *For every finite extensive form game of perfect information, and all $0 \leq \lambda \leq \infty$, $b^*(\lambda)$ is unique.*

Proof: This is proved by construction of $b^*(\lambda)$. Since the game has perfect information the predecessors of each terminal node all lie in different (singleton) information sets. Consider one such information set, $h, i(h) = i, A(h) = \{a_1, \dots, a_K\}$ and for each $a \in A(h)$, \exists unique outcome $\nu_a \in \Gamma^t$ such that $\rho(\nu_a|h, b_a) = 1$, where b_a is the behavior strategy profile b , with $b_i(h)$ replaced by a , so $u_i(a|h) = u_i(\nu_a)$. This *uniquely* determines the action probabilities of logit-AQRE at any node immediately preceding a terminal node by

$$b_i^*(a_j|h) = \frac{e^{\lambda u_j(\nu_{a_j})}}{\sum_{k=1}^K e^{\lambda u_i(\nu_{a_k})}}$$

Then the expected utility to an agent who has reached h but has not chosen an action at h is $\sum_{a_k} b_i^*(a_k, h) u_i(\nu(a_k, h)) = u_i(b|h)$. Notice that $u_i(b|h)$ is therefore uniquely determined. Next consider the immediate predecessors of the immediate predecessors of Γ^t . Repeating the above argument yields a unique $b^*(h')$ for all such h' such that $Q|(Q(V)) \subseteq \Gamma^t \forall v \in h'$. Since the game is finite, we can iterate this argument until b^* is defined (uniquely) at each information set. ■

A fifth property we currently offer only as a conjecture: that for generic games there exists a unique connected path in the graph of this logit-AQRE correspondence which includes b_0^* (i.e., the equilibrium when $\lambda = 0$). This unique connected path defines a solution \bar{b}_λ^* for all values of $\lambda \geq 0$, and furthermore there is a unique limit point $\bar{b}^* = \lim_{\lambda \rightarrow \infty} \bar{b}_\lambda^*$. We call \bar{b}^* the *logit-AQRE solution* of the game.

Conjecture 1 *Almost all finite extensive form games have a unique connected path selection from the logit-AQRE correspondence.*

Note that example 1 in McKelvey and Palfrey (1994) demonstrates that the converse of Theorem 2 is not true. That is, there are sequential equilibria that are not approachable by a sequence of logit-AQREs. Thus, the set of limit points logit-AQRE provides a refinement of sequential equilibrium. Whenever there is a unique connected path selection of the logit-AQRE correspondence, this gives a refinement of the AQRE-approachable sequential equilibria which is an even stronger refinement of sequential equilibrium. As we will see below, it offers a quite different refinement from standard deductive refinements

(such as the intuitive criterion)³, and predicts much better than the intuitive refinement in signalling games that have been conducted as laboratory experiments.

Another nice feature of AQRE is that it makes predictions about the relative likelihood of *all* different play paths. In contrast, other refinements tend to predict that only a subset of the play paths can occur. Because all play paths can occur with positive probability in AQRE, we don't have to worry about the arbitrary specification of player beliefs "off the equilibrium path."

In the next section we demonstrate how this model of equilibrium behavior in extensive form games contradicts the invariance principle, that all strategically relevant information in an extensive form game is captured in the normal form. We illustrate this in the context of a particular game, the one-shot chain-store paradox. The way the game is actually played (i.e., as a game where both players simultaneously choose strategies *vs.* a two-stage game in which the incumbent first observes whether an entrant has entered before choosing his strategy) will systematically affect the predicted pattern of play using the logit-AQRE. Furthermore the differences implied theoretically by logit-AQRE mirror experimental findings reported recently by Schotter, Weigelt, and Wilson [1994].

5 Invariance

Consider the well-known chainstore paradox stage game,⁴ with the extensive form illustrated in Figure 1(a). The normal form representation of this game is given in Table 1.

		2	
		<i>F</i>	<i>A</i>
1	<i>N</i>	2,2	2,2
	<i>E</i>	0,0	3,1

Table 1.
Normal form of Chainstore stage game.

³That is, there exists games in which there is a unique intuitive equilibrium which is different from the unique logit-AQRE selection.

⁴This game can also be interpreted as a discrete (binary) version of an ultimatum game in which the offerer may offer either a 50/50 split or a 75/25 split of the pie. The unfair split (75/25) may be either accepted or rejected by the other player. The fair split is automatically accepted. Gale, Binmore, and Samuelson [1993] call it the "ultimatum minigame," and study the dynamic stability of its Nash equilibria under the replicator dynamic.

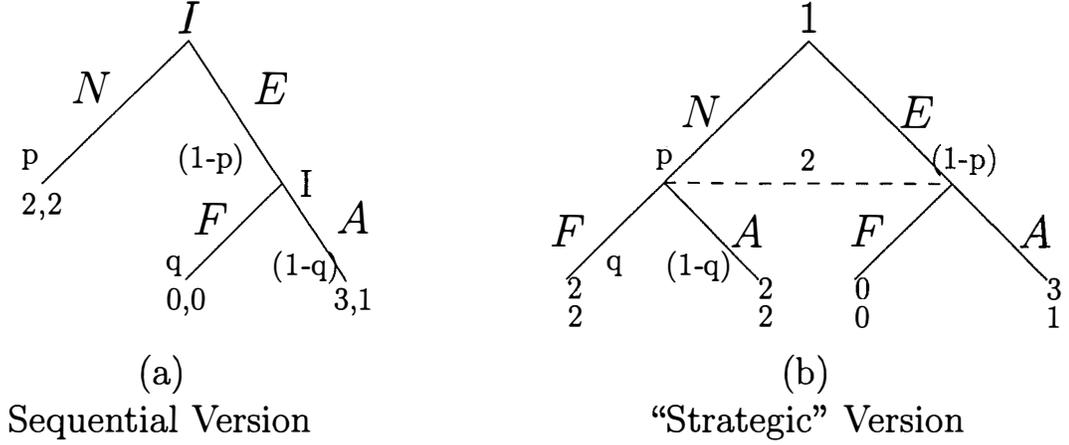


Figure 1: Chainstore paradox: Player 1 is the entrant, player 2 is the incumbent (E = Enter, N = Not enter, F = Fight, A = Acquiesce)

The logit equilibrium (McKelvey and Palfrey (1994)) for the normal form of Table 1 is equivalent to the logit AQRE of the extensive form game of Figure 1(b), which is a strategically invariant transformation of the Figure 1(a) game.

We wish to compare the quantal response equilibrium as it applies to the game in Figure 1(a) to the quantal response equilibrium as it applies to the game in Figure 1(b). To do this, we compare the logit-AQRE as a function of λ in the first case to the logit equilibrium [McKelvey and Palfrey, 1993] in the second case.

The equilibrium conditions for the logit-AQRE of the extensive form game in Figure 1(a) are:

$$p = \frac{1}{1 + e^{(1-3q)\lambda}} \quad (1)$$

$$q = \frac{1}{1 + e^\lambda} \quad (2)$$

where $p = \text{prob} \{1 \text{ moves } N\}$ and
 $q = \text{prob} \{2 \text{ moves } F\}$

The equilibrium conditions for the extensive form game in Figure 1(b) are:

$$p' = \frac{e^{2\lambda}}{e^{2\lambda} + e^{3(1-q')\lambda}} = \frac{1}{1 + e^{(1-3q')\lambda}} \quad (3)$$

$$q' = \frac{e^{2p'\lambda}}{e^{(2p'+(1-p'))\lambda} + e^{2p'\lambda}} = \frac{1}{1 + e^{(1-p')\lambda}} \quad (4)$$

when $D > 0$. Similarly, $q > q' > \frac{1}{2}$ if $D < 0$. In general, $|q - \frac{1}{2}| > |q' - \frac{1}{2}|$, which implies that deviations from the subgame perfect Nash equilibrium prediction should always be *greater* when the game is played in the strategy version. How this affects the moves by player 1 is more subtle, and depends on whether C and D have the same or different signs. E is a relatively more attractive option in the strategy version if and only if $C > 0$ and $q' > q$ or $C < 0$ and $q' < q$. This requires $DC < 0$. Similarly E is a relatively less attractive option if and only if $DC > 0$. Therefore $p > p' \Leftrightarrow CD < 0$.

The qualitative predictions of AQRE match up with some experimental findings presented by Schotter, Weigelt, and Wilson [1994] (SWW). They conducted an experiment with $A = 4$, $B = 3$, $C = 6$, and $D = 2$, where some subjects played the Figure 1 game and other subjects played the Figure 2 game, and still other subjects played the Table 1 game. Pooling the results for the Figure 2 games and the Table 1 games, they found that entry was deterred more often with the strategy version of the game than with the sequential version of the game. Furthermore, the second mover chose Fight more often in the strategy version of the game.⁵

		n	f_i	Perfect		
				Rand	Nash	QRE
Sequential Version	\hat{p}	7	.089	.500	0.000	.066
	$1 - \hat{p}$	79	.911	.500	1.000	.934
	\hat{q}	2	.025	.500	0.000	.045
	$1 - \hat{q}$	77	.975	.500	1.000	.955
λ				0	∞	1.531
λ_{lo}				0		1.36
λ_{hi}				0		1.86
$-\mathcal{L}^*$					∞	34.16
Strategic Version	\hat{p}'	46	.575	.500	0.000	.472
	$1 - \hat{p}'$	34	.425	.500	1.000	.528
	\hat{q}'	16	.200	.500	0.000	.309
	$1 - \hat{q}'$	64	.800	.500	1.000	.691
λ				0	∞	.762
λ_{lo}						.514
λ_{hi}						.932
$-\mathcal{L}^*$				110.9	∞	98.69

Table 2
SWW experimental results (36 subjects)

The logit-AQRE correspondence for both the sequential and strategy versions of the game are given in Figure 4. Superimposed on those graphs are the observed data points

⁵The authors also found that these differences were magnified by presenting the strategy version to the subjects in its matrix form. We are not proposing an explanation for this.

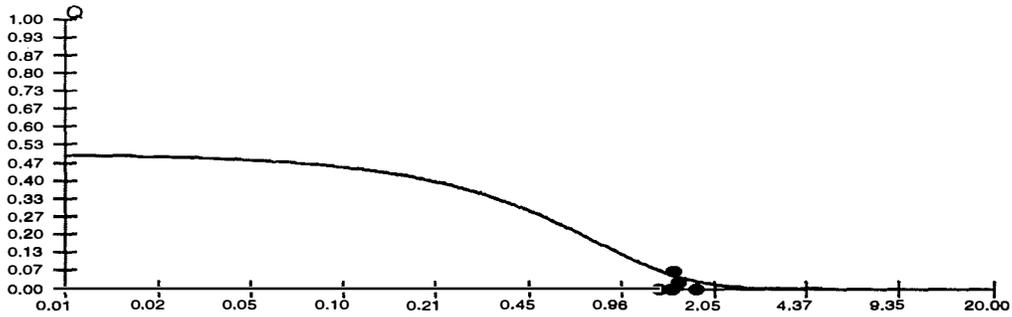
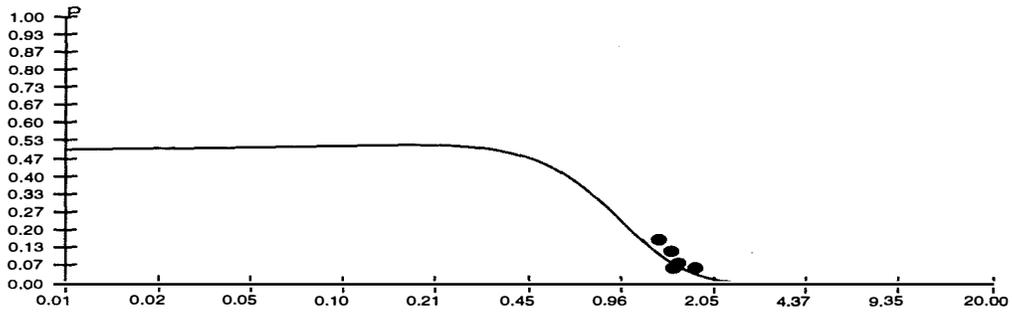


Figure 3: SWW (1994) Sequential version

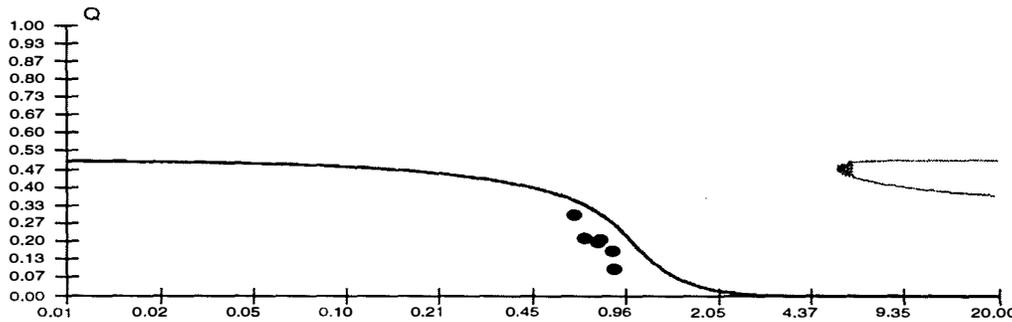
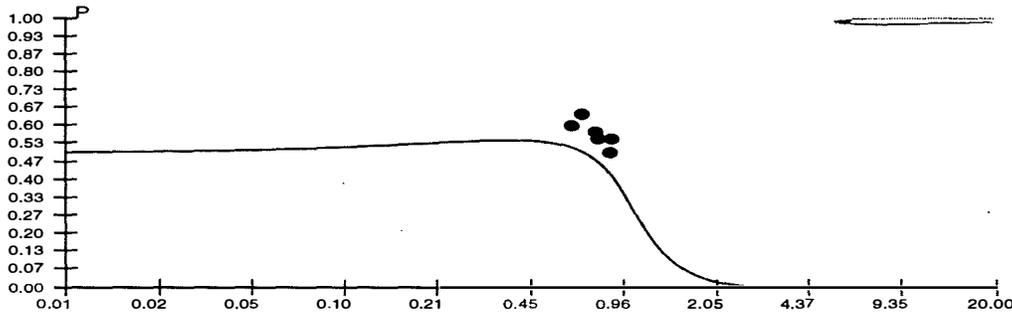


Figure 4: SWW (1994) Strategy version

plotted at the maximum likelihood values of the logit response parameter. Table 2 gives the predicted and actual data aggregated across all of their experiments. In this table (as well as subsequent tables), λ represents the maximum likelihood estimate of λ , and $(\lambda_{lo}, \lambda_{hi})$ is a 95% confidence interval for λ . The estimation here, as well as the estimations presented later in this paper are performed under the restriction that all subjects (both row and column players) have the same λ .⁶ There are multiple points on each graph, with each point representing a different experience level of the subjects. As is evident from the graphs, experience seems to have little effect in their data.⁷ One can readily see that the logit-AQRE fits quite well in both versions of the game, although it underpredicts the frequency of (N,A) outcomes somewhat in the strategy version of the game. This underprediction partly results from our having included data from their matrix version of the game as well as the extensive form version.

There are several other games examined by SWW, but none of these others were games of perfect information, and all have multiple AQREs, even in the sequential versions. We do not analyze them here.

6 Equilibrium Selection

As a second application of AQRE, we look at simple sender-receiver signaling games. Several experimental studies have been conducted using sender-receiver games, notably by Banks, Camerer and Porter (1994), Brandts and Holt (1992, 1993), and Partow and Schotter (1993). In many ways this is an ideal class of games to study with AQRE.

First, the games typically have multiple sequential equilibria, and logit-AQRE generically selects a unique one as the limit of the unique connected path in the logit-AQRE graph. The earlier figures showing the logit-AQRE correspondence for the SWW experiments illustrates this selection. Notice that in the strategy version of their game there are multiple logit-AQRE for high values λ . However, there is a unique connected path starting at the unique equilibrium at $\lambda = 0$. The unique connected path selection is the darker curve in that figure. The proof is similar for extensive form games. For reasons of space, we do not reexamine the results of all the experimental signaling games that have been conducted to date. We focus on a subset of the experimental games reported by Banks, Camerer and Porter (1994) (BCP) and Brandts and Holt (1993) (BH).

Banks, Camerer, Porter, [1994]

The BCP experiment consisted of a series of two person signaling games, with two sender types, three messages, and three responses. Each game was designed to have two Nash equilibria, one of which was always further up the chain of deductive refinements

⁶Obviously, the fit to the data will improve dramatically if we allow for differences in λ across the different players. We briefly discuss some alternative ways to incorporate heterogeneity later in the paper.

⁷This may be partly due to the fact that subjects played the games at most six times, so none of the experience levels are very high.

than the other. The experiments were intended to give data on the conditions under which subjects play the more refined equilibrium. We analyze games 2, 3, and 4 of BCP. These games are given in Table 3. Each of these games has two Nash equilibrium outcomes, with one being a more refined equilibrium than the other.⁸

The logit-AQRE graphs are displayed in figures 9–11 for these three games. The estimation of λ is reported in Table 4 for each game. We see that in each of the three cases, the QRE selects the more refined equilibrium, which corresponds to what was observed in the data.

Moreover, the QRE predictions pick up quite well which strategies are relatively more overplayed (or underplayed) compared to the refined equilibrium predictions. In BCP #2, the predicted equilibrium has both types sending message 3, with the receiver choosing action 1 in response to message 3, and choosing action 2 in response to messages 1 and 2. For both types of senders, AQRE predicts that nearly all message errors should involve sending message 1. Indeed in all 18 deviations from this equilibrium, senders use message 1, and never use message 2.⁹ Response errors to the equilibrium message are predicted by AQRE to be more than 10 times as likely to be action 3 than to be action 2, and indeed action 3 errors are observed in all 10 of the responders' deviations from equilibrium. Also, QRE predicts that type 1 senders are less likely to deviate from the sequential equilibrium than type 2 senders, which was also observed in the data.

In BCP #3 the predicted (i.e., unique intuitive) equilibrium has both types sending message 1, with the receiver choosing action 2 in response to message 1, and choosing action 1 in response to both nonequilibrium messages. Of the 33 deviations from the intuitive equilibrium predicted strategies, most were deviations by type 1 senders, who are predicted by the AQRE model to send the intuitive message approximately half the time. (Indeed, type 1 senders were observed to send the intuitive message exactly half the time!)

In BCP #4, the predicted (i.e., divine) equilibrium has both agents sending message 3, with the receiver choosing action 1 in response to the equilibrium message 3, and action 2 otherwise. In the quantal response equilibrium, both sender types are predicted to send the divine equilibrium message only about half the time. Aggregating both types of senders, this is what we observe in the data, where message 3 is sent 28 times out of 60 chances; message 1 is predicted to be sent slightly more than a third of the time, and we observe it sent 23 out of 60 times. The frequency of actions in response to each message is predicted even more accurately (see Table 4).

⁸In game 2, one equilibrium is sequential and the other is not. In game 3, both are sequential, but only one is intuitive. In game 4, both are intuitive, but only one is divine.

⁹Neither m_1 nor m_2 are dominated strategies.

BCP # 2 (Nash vs. Sequential)

m_1	a_1	a_2^n	a_3	m_2	a_1	a_2	a_3	m_3	a_1^s	a_2	a_3
t_1	1, 2	2, 2	0, 3	t_1	1, 2	1, 1	2, 1	t_1	3, 1	0, 0	2, 1
t_2	2, 2	1, 4	3, 2	t_2	2, 2	0, 4	3, 1	t_2	2, 2	0, 0	2, 1

Nash: (m_1, a_2, a_2, a_2) Seq: (m_3, a_2, a_2, a_1)

BCP # 3 (Sequential vs. Intuitive)

m_1	a_1	a_2^I	a_3	m_2	a_1	a_2	r_3^s	m_3	a_1^s	a_2	a_3
t_1	0, 3	2, 2	2, 1	t_1	1, 2	2, 1	3, 0	t_1	1, 6	4, 1	2, 0
t_2	1, 0	3, 2	2, 1	t_2	0, 1	3, 1	2, 6	t_2	0, 0	4, 1	0, 6

Sequential: (m_3, a_1, a_3, a_1) Intuitive: (m_1, a_2, a_1, a_1)

BCP # 4 (Intuitive vs. Divine)

m_1	a_1	a_2	a_3	m_2	a_1	a_2	a_3^I	m_3	a_1^D	a_2	a_3
t_1	4, 0	0, 3	0, 4	t_1	2, 0	0, 3	3, 2	t_1	2, 3	1, 0	1, 2
t_2	3, 4	3, 3	1, 0	t_2	0, 3	0, 0	2, 2	t_2	4, 3	0, 4	3, 0

Intuitive: (m_2, a_3, a_3, a_2) Divine: (m_3, a_2, a_2, a_1)

Table 3
BCP (1994) Signaling Games

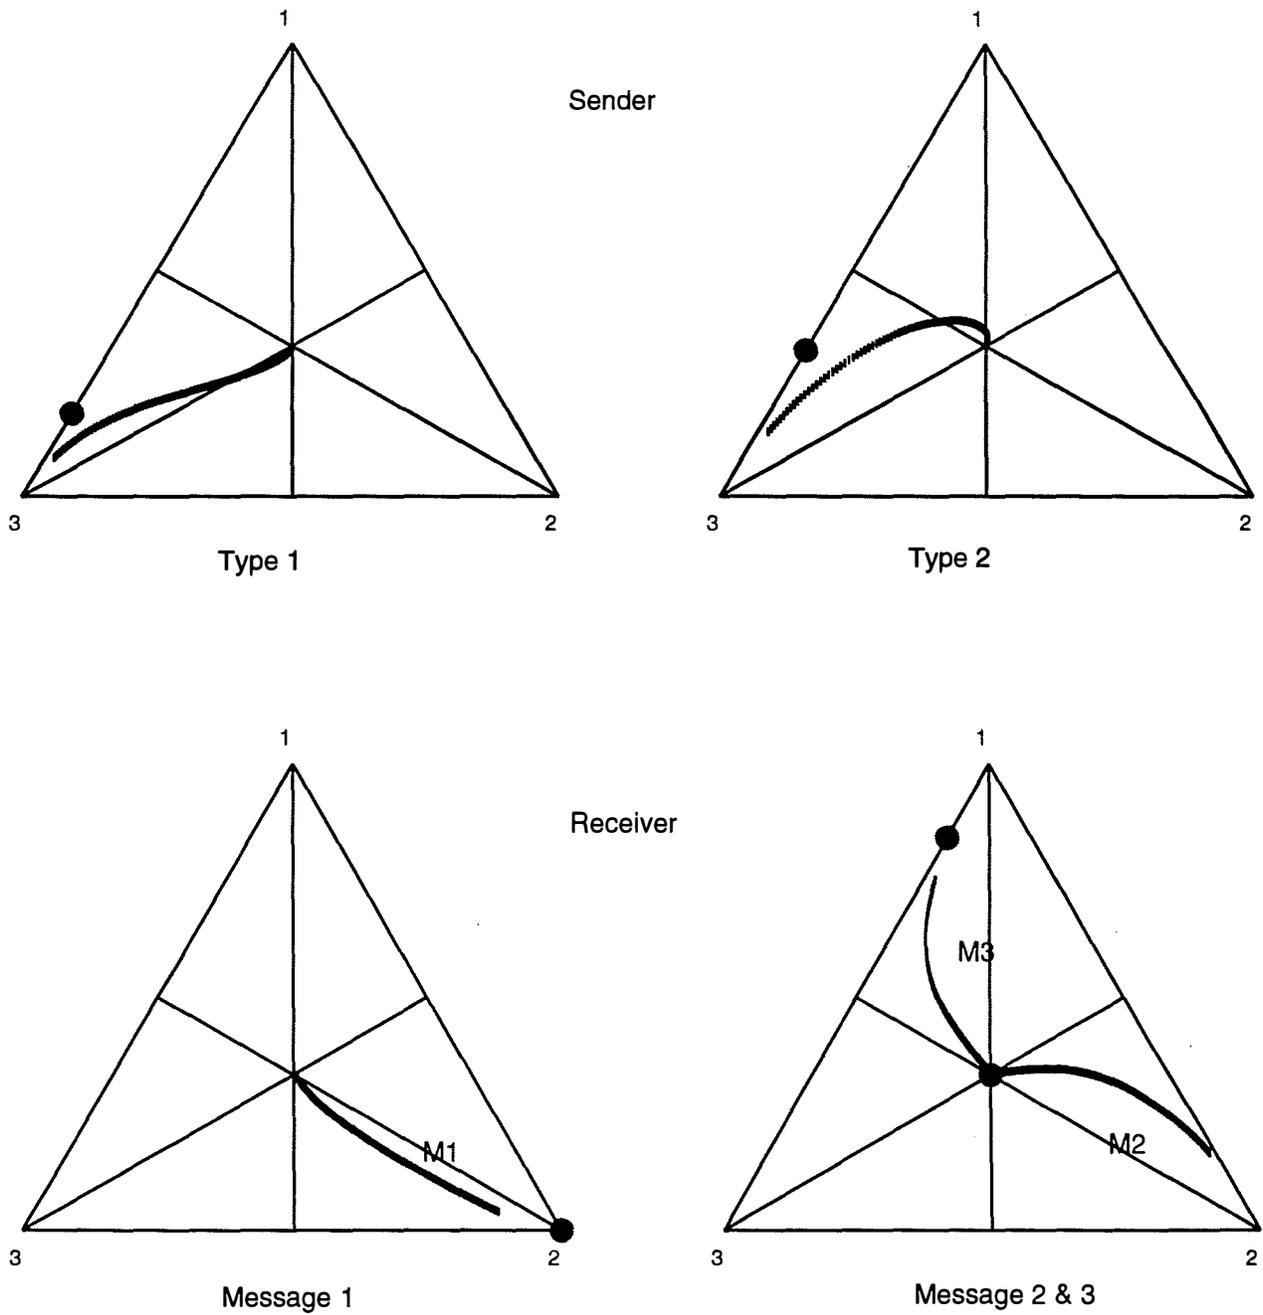
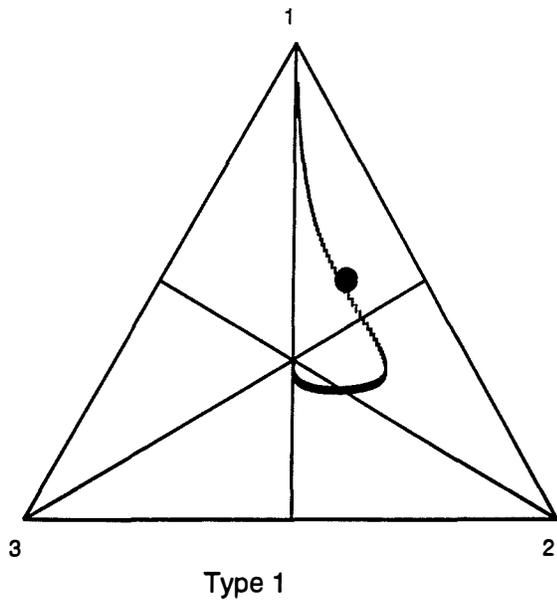
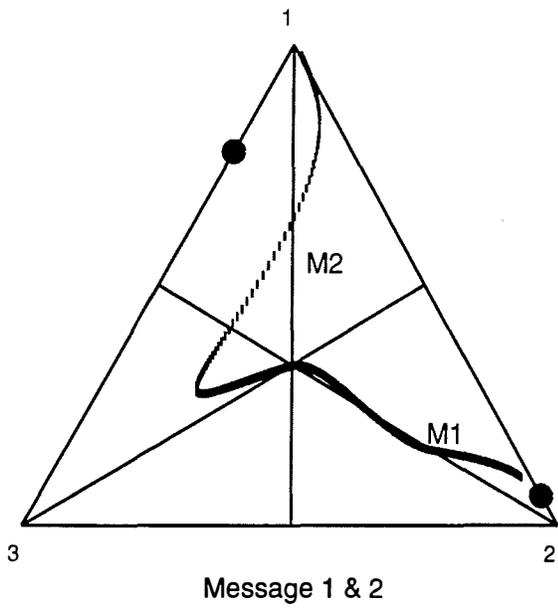
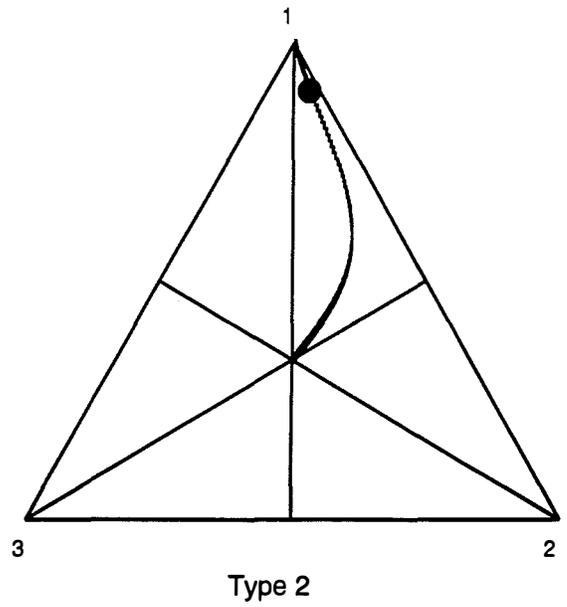


Figure 5: BCP (1994) Game 2. Sequential vs. Nash



Sender



Receiver

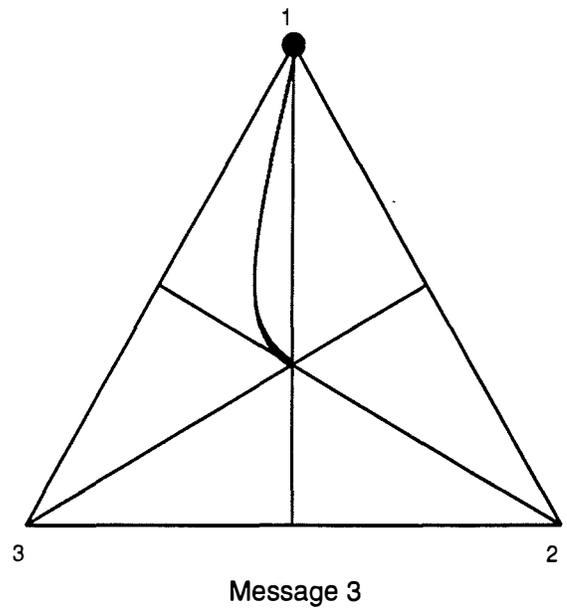
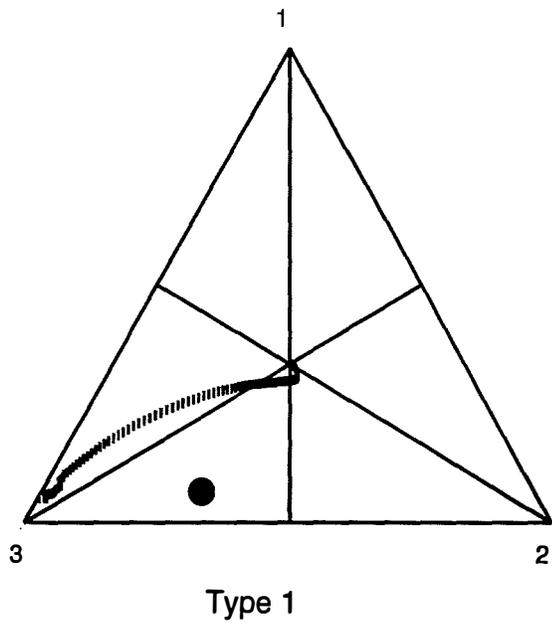
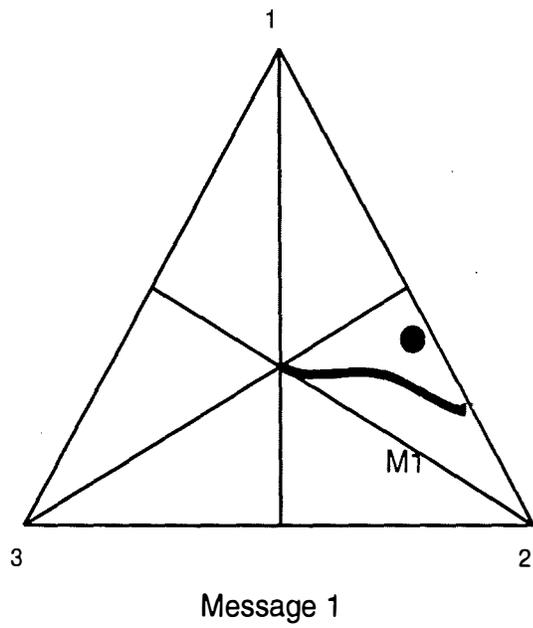
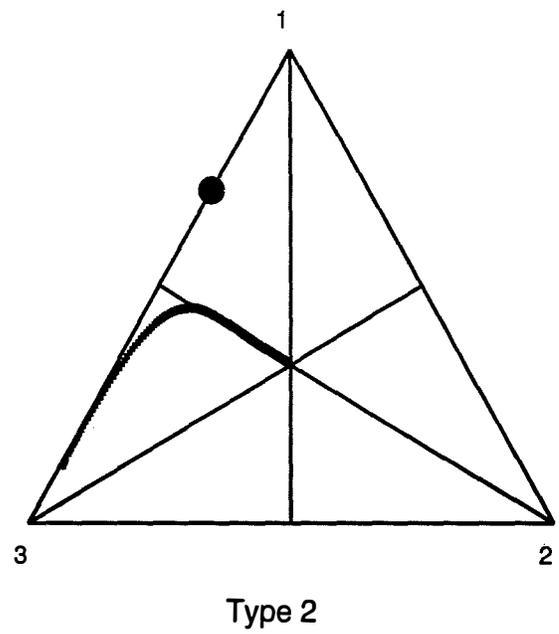


Figure 6: BCP (1994) Game 3. Intuitive vs. Sequential



Sender



Receiver

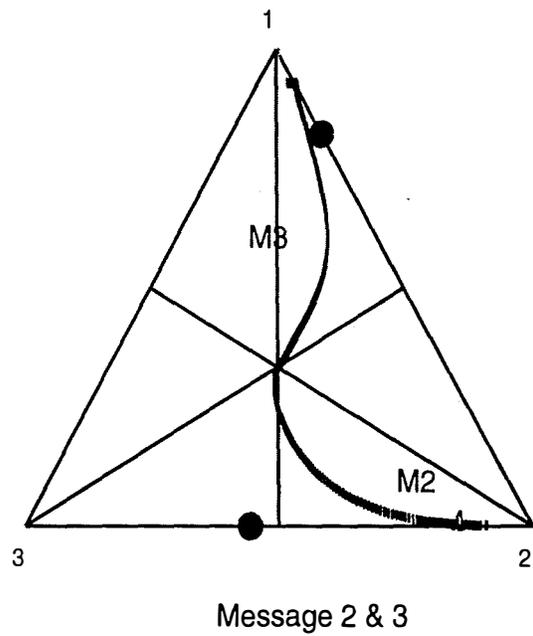


Figure 7: BCP (1994) Game 4. Divine vs. Intuitive

		BCP #2			BCP #3			BCP #4		
		n	f_i	QRE	n	f_i	QRE	n	f_i	QRE
t_1	m_1	7	.180	.103	20	.500	.497	2	.070	.274
	m_2	0	.000	.021	14	.350	.339	9	.300	.221
	m_3	31	.820	.875	6	.150	.165	19	.630	.504
t_2	m_1	11	.380	.176	45	.900	.927	21	.700	.422
	m_2	0	.000	.026	4	.080	.057	0	.000	.032
	m_3	23	.620	.798	1	.020	.016	9	.300	.546
m_1	a_1	0	.000	.049	4	.062	.153	9	.391	.298
	a_2	18	1.000	.837	61	.938	.704	13	.565	.593
	a_3	0	.000	.113	0	.000	.142	1	.043	.109
m_2	a_1	0		.184	14	.778	.686	0	.000	.044
	a_2	0		.797	0	.000	.174	4	.444	.647
	a_3	0		.019	4	.222	.140	5	.556	.308
m_3	a_1	53	.841	.726	7	1.000	.999	23	.821	.704
	a_2	0	.000	.026	0	.000	.001	5	.179	.235
	a_3	10	.159	.248	0	.000	.000	0	.000	.062
$\lambda (\sigma^2)$				2.249			1.598			1.193
λ_{lo}				1.901			1.441			0.989
λ_{hi}				2.715			1.864			1.470
$-\mathcal{L}^*$				83.39			101.05			95.65

Table 4
BCP experimental results

Brandts and Holt, [1994]

As a follow-up to the BCP experiment Brandts and Holt (1994) proposed two alternatives to Game 3 (which differentiated the intuitive equilibrium from a non-intuitive sequential equilibrium), which they ran using similar¹⁰ procedures. The games, which differ slightly from BCP because the sender has only two available messages, are given in Table 5. The two games each have a sequential and an intuitive equilibrium at the same strategies. The data along with the QRE estimates are given in Table 6 and Figures 8-9.

The point of the experiment was to attempt to construct a game in which play might conceivably converge to the *less* refined equilibrium. Brandts and Holt proposed a descriptive story of how a particular dynamic of play in early rounds could potentially lead to the unintuitive equilibrium, if the payoffs in BCP #3 were changed slightly. This is an ideal setting for putting the logit AQRE selection to the test, since the rough idea behind our selection is that if subjects begin an experiment at a QRE with relatively high error rates (i.e., low value of λ) and “follow” the equilibrium selection as λ increases

¹⁰There were some minor procedural differences that are probably inconsequential. BH successfully replicated the results of BCP Game 3.

with experience, then they should drift toward the sequential equilibrium implied by our selection.

In fact, BH #4 provides exactly this kind of “critical test,” since the logit-AQRE selection in the *nonintuitive* equilibrium, (S,D,C). This is exactly what was observed in the data. Table 6 presents the results of the maximum likelihood estimation of λ , and the observed frequencies, pooling all experience levels. In figures 8 and 9, we present graphically the results of estimation with a finer breakdown of experience levels. In those graphs, higher numbers correspond to higher levels of experience (in the same way as in BH (1993, Figure 5, p. 442)). The curves in the graphs represent the predicted move frequencies for the logit-AQRE selection, as a function of λ . While the estimated values of λ are not strictly increasing in experience levels, there is an overall trend in the right direction.

BH # 3 (Sequential vs. Intuitive)

$m = I$	C	D^n	E	$m = S$	C	D	E
A	45,30	15,0	30,15	A	30,90	0,15	45,15
B	30,30	0,45	30,15	B	45,0	15,30	30,15

Sequential: (S, D, C) Intuitive: (I, C, D)

BH # 4 (Sequential vs. Intuitive)

$m = I$	C	D^n	E	$m = S$	C	D	E
A	30,30	0,0	50,35	A	45,90	15,15	100,30
B	30,30	30,45	30,0	B	45,0	0,30	0,15

Sequential: (S, D, C) Intuitive: (I, C, D)

Table 5

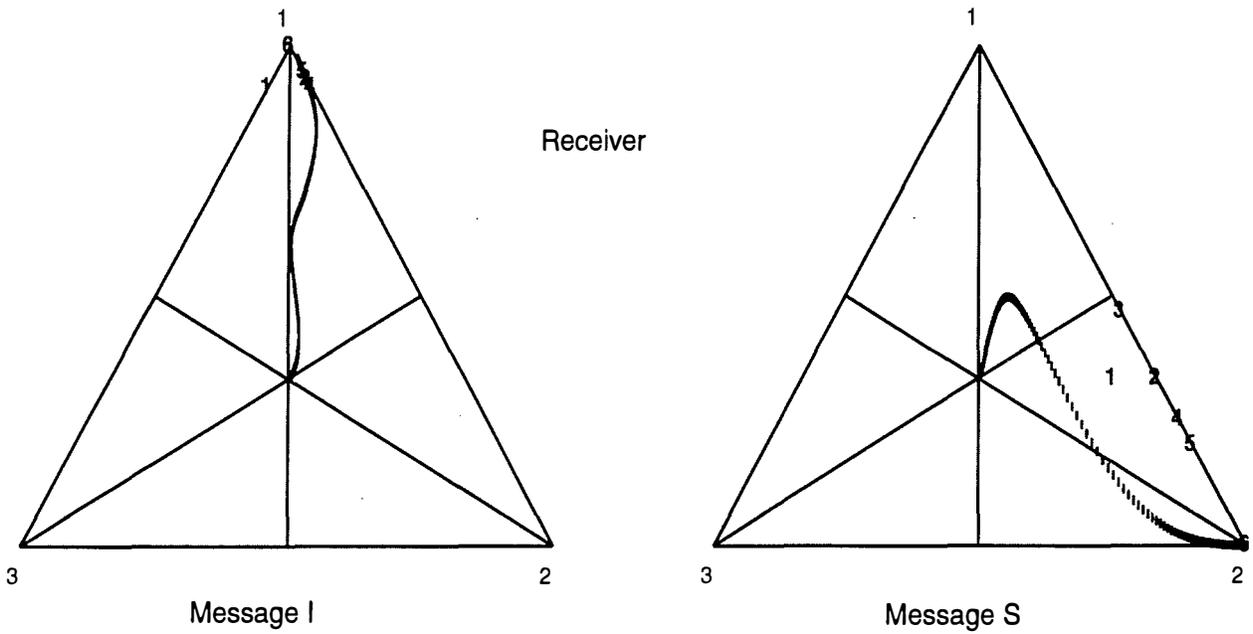
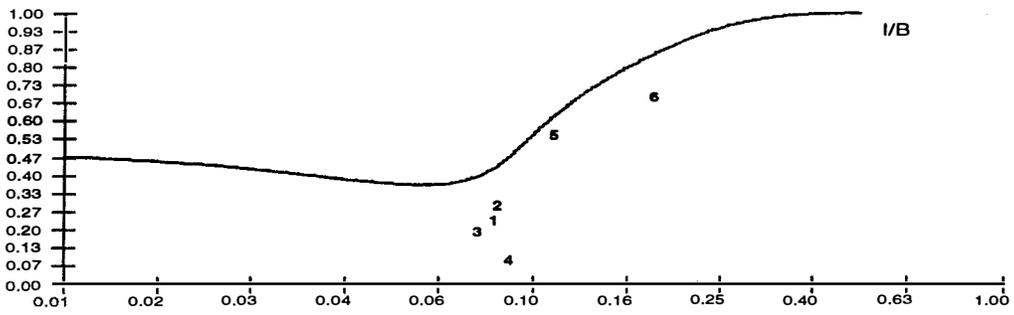
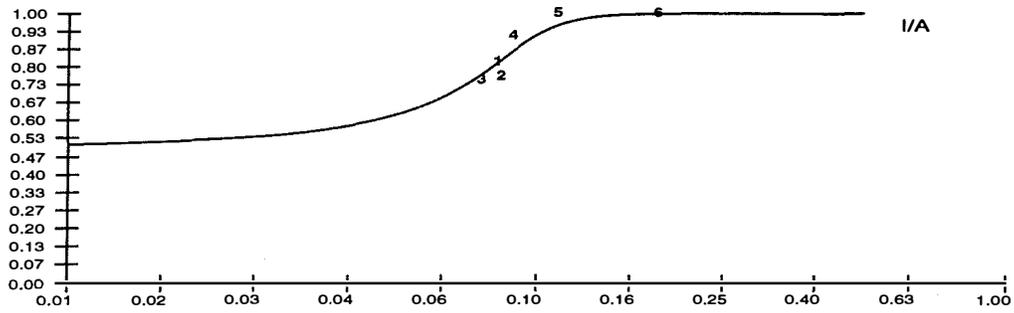


Figure 8: BH (1993) Game 3. Sequential vs. Intuitive

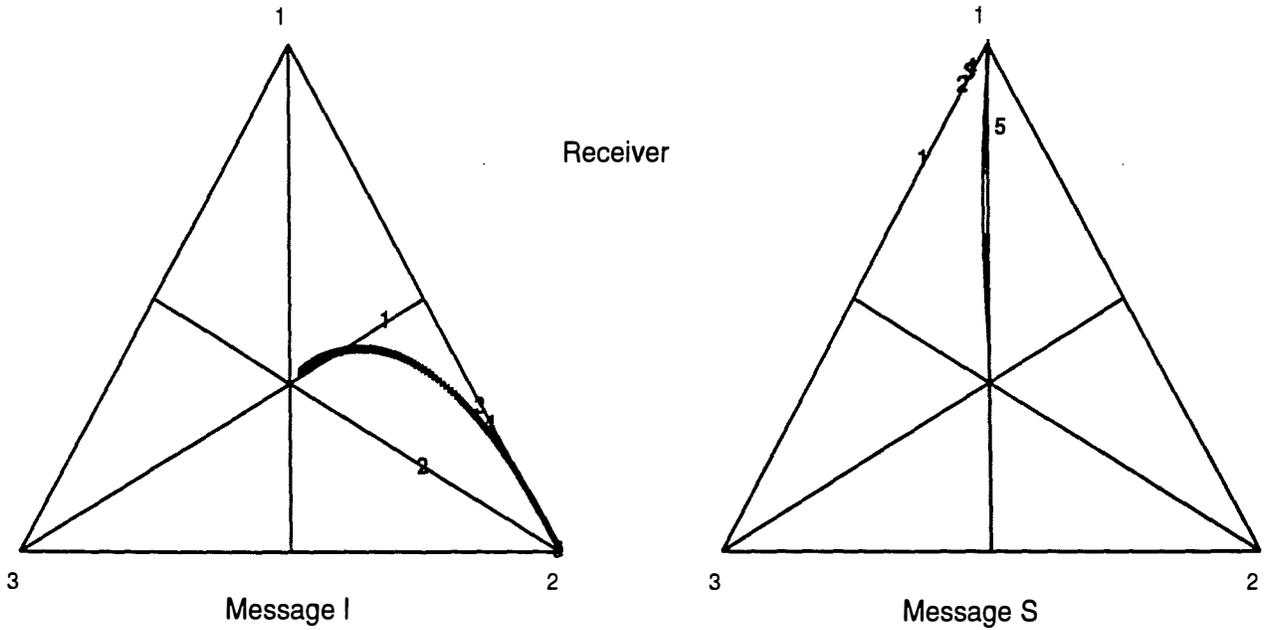
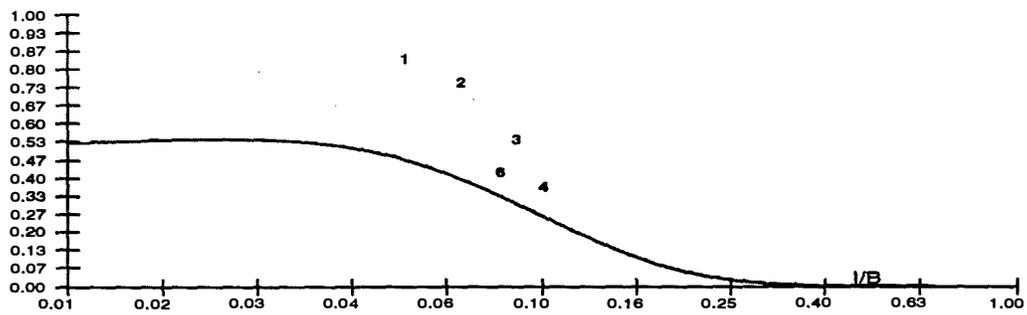
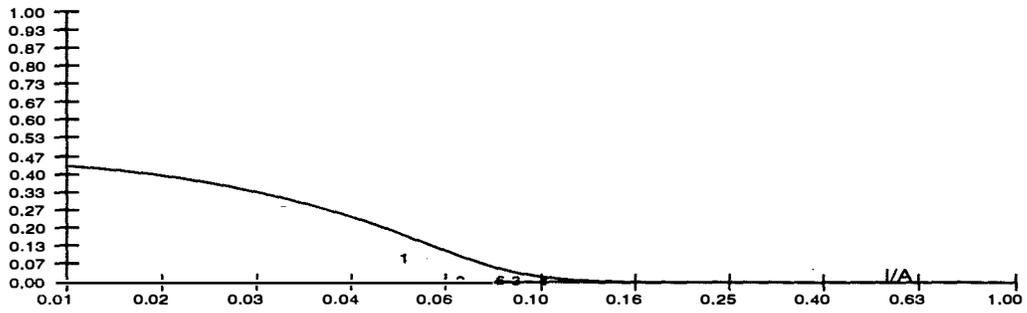


Figure 9: BH (1993) Game 4. Sequential vs. Intuitive

		BH #3			BH #4		
		<i>n</i>	<i>f_i</i>	QRE	<i>n</i>	<i>f_i</i>	QRE
<i>t = A</i>	<i>m = I</i>	32	.970	.950	0	.000	.026
	<i>m = S</i>	1	.030	.050	38	1.000	.974
<i>t = B</i>	<i>m = I</i>	18	.462	.610	15	.441	.280
	<i>m = S</i>	21	.538	.390	19	.559	.720
<i>m = I</i>	<i>a = C</i>	48	.960	.689	1	.067	.252
	<i>a = D</i>	2	.040	.178	13	.867	.729
	<i>a = E</i>	0	.000	.134	1	.067	.019
<i>m = S</i>	<i>a = C</i>	4	.182	.101	52	.912	.888
	<i>a = D</i>	18	.818	.728	2	.035	.050
	<i>a = E</i>	0	.000	.171	3	.053	.062
$\lambda (\sigma^2)$.109			.095
λ_{lo}				.095			.076
λ_{hi}				.128			.117
$-\mathcal{L}^*$				69.53			56.27

Table 6
Data and estimates for BH experiments.

7 Centipede Game Experiments

McKelvey and Palfrey (1992) studied experimentally two versions of the centipede game (Rosenthal 1982, Binmore 1990, Kirps 1990), displayed in Figures 10 and 11, respectively. In that paper they proposed a theory of play based on incomplete information, altruism, and heterogeneous prior beliefs and fit a 5-parameter model to the data. The model captured the essential qualitative features of the data, and fit the exact move frequencies of the data fairly well.

In this section we re-examine the centipede data using the logit EFQRE model. We describe the details of the model as it applies to the extensive form-model versions of Figure 10. The computations for the six-move version are similar. Denote player 1's behavior strategies as (p_1, p_2) and player 2's behavior strategies as (q_1, q_2) , where

$$\begin{aligned}
 p_1 &= \text{prob \{1 chooses } T \text{ at first move\}} \\
 p_2 &= \text{prob \{1 chooses } T \text{ at third move (if reached)\}} \\
 q_1 &= \text{prob \{2 chooses } T \text{ at second move (if reached)\}} \\
 q_2 &= \text{prob \{2 chooses } T \text{ at fourth move (if reached)\}}
 \end{aligned}$$

At the first move, player 1 estimates the payoff of T and P by:

$$\begin{aligned}
 U_{11}(T) &= .4 + \varepsilon_{1T} \\
 U_{11}(P) &= .2q_1 + (1 - q_1)[1.6p_2 + (1 - p_2)(.8q_2 + 6.4(1 - q_2))] + \varepsilon_{1P}
 \end{aligned}$$

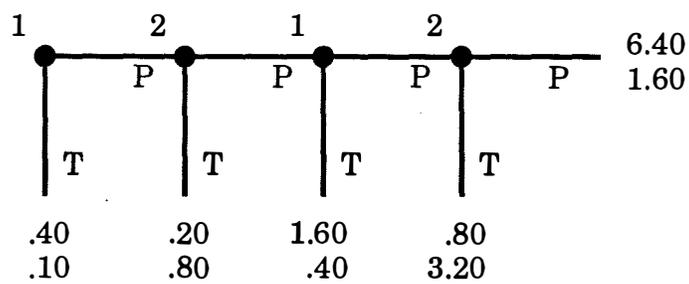


Figure 10: The Four Move Centipede Game.

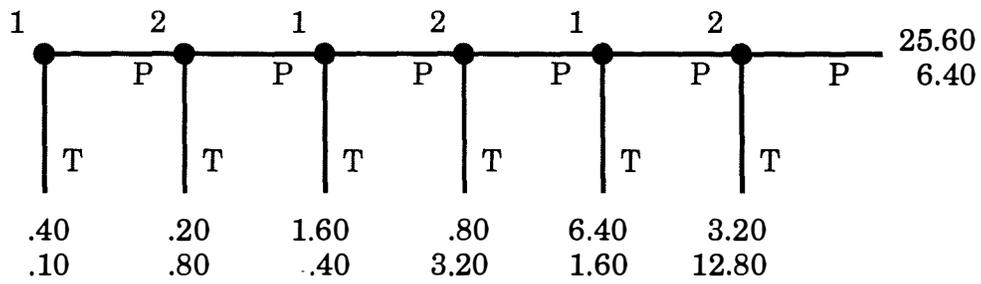


Figure 11: The Six Move Centipede Game.

McKelvey-Palfrey (1992) centipede experiments

where ε_{1T} and ε_{1P} are independent random variables with a Weibull distribution with parameter λ . The logit formula then implies

$$p_1 = \frac{1}{1 + e^{\lambda[.2q_1 + (1-q_1)[1.6p_2 + (1+p_2)[.8q_2 + 6.4(1-q_2)]] - .4}}$$

Similarly, we obtain expressions for p_2 , q_1 , and q_2 , given below,

$$\begin{aligned} q_1 &= \frac{1}{1 + e^{\lambda[.4p_2 + (1-p_2)[3.2q_2 + 1.6(1-q_2)] - .8}} \\ p_2 &= \frac{1}{1 + e^{\lambda[.8q_2 + 6.4(1-q_2) - 1.6]}} \\ q_2 &= \frac{1}{1 + e^{\lambda[1.6 - 3.2]}} \end{aligned}$$

This system of four equations and four unknowns is solved recursively (q_2 , then p_2 , then q_1 , then p_1) and has a unique solution. Figure 12 graphs the solutions of (p_1, p_2, q_1, q_2) as a function of λ .

There are several interesting features of the equilibrium graph. First, for most values of λ we have $p_1 < q_1 < p_2 < q_2$. This is consistent with the experimental findings, where $p_1 = .07$ $q_1 = .38$ $p_2 = .65$ $q_2 = .75$

Second, for *very* low values of λ , p_1 , q_1 , and p_2 are all decreasing in λ . That is, the equilibrium does not converge monotonically toward the Nash equilibrium as a function of λ .

We estimate the best fitting value of λ for this dataset, using standard maximum likelihood techniques¹¹ The results are reported in Tables 7 and 8.

That table also reports the estimates based on the (reduced) normal form QRE model of play. Since the game was not played in this manner, we would expect a poorer fit than the AQRE model, which we find to be true in both the 4-move and 6-move games. The graph of the normal form QRE is given in figure 13. Notice that even though the normal form model does not fit as well, it admits much lower take probabilities for intermediate values of λ , compared to the AQRE model.

Our findings mirror results reported in Zauner (1993), using a related model. That paper presents a structural model of the error process by adding independently distributed random payoff disturbances, as in Harsanyi (1973), and estimates that variance (see last column of Table 5 and 6.). That paper also adopts an ‘‘agent model’’ approach in that a player does not observe all disturbances to payoffs at the start of the game, but instead only the disturbance to his own current ‘‘take’’ payoff is observed. This is different in spirit to the Harsanyi (1973) model which assumes each player observes all his own payoff disturbances at the beginning of play. Zauner’s model fits better than ours, as measured

¹¹ A similar analysis is reported in Fey, McKelvey and Palfrey (1993) using an experiment based on a constant sum version of the centipede game. The details of the maximum likelihood procedures are explained in that paper.

by the likelihood score, but both models share two similar properties: they capture the feature in the data that take probabilities increase as the end of the game is approached; and they overestimate the take probabilities on the first and last moves. But as Zauner (1993) points out, neither model includes a heterogeneity parameter, which produced a dramatic improvement in fit in the models estimated by McKelvey and Palfrey (1992). One might hope that the inclusion of a heterogeneity parameter would also significantly improve the fits of these two models.

8 Conclusions

There are other ways these models might be improved on. As discussed above, these models assume that a player cannot predict his own future play (the “agent” approach). Alternatively, one might assume that each player observes all his payoff disturbances at the beginning of the game. This “player” approach is more difficult to estimate because the equilibrium can no longer be solved recursively from the last move of the game. Another possible improvement would be to incorporate persistence in the error structure so that a player who passed at the beginning is also likely to pass later on. This would mimic the role of altruism.

We consider here a third possibility, which is to incorporate a second source of error, consistent our previous analysis of this data. Namely, we assume that there is some small probability that players are “altruistic” (and hence choose P at every opportunity.) This gives rise to a game of incomplete information (as described formally in McKelvey, Palfrey, page 827) where with probability q , a player is “rational”, and with probability $(1 - q)$, they are “altruistic”. We then apply the AQRE directly to this extensive form game, obtaining a two parameter model in which we estimate both λ and q . The results of this analysis appear in the last column of Tables 7 and 8. As one can see, there is a significant improvement in fit over the one parameter model, at any reasonable significance level. While the log likelihood is still several orders of magnitude worse than the 5-parameter model of McKelvey and Palfrey (1992), one must keep in mind that that model permitted heterogeneity and time trends, in addition to altruism.

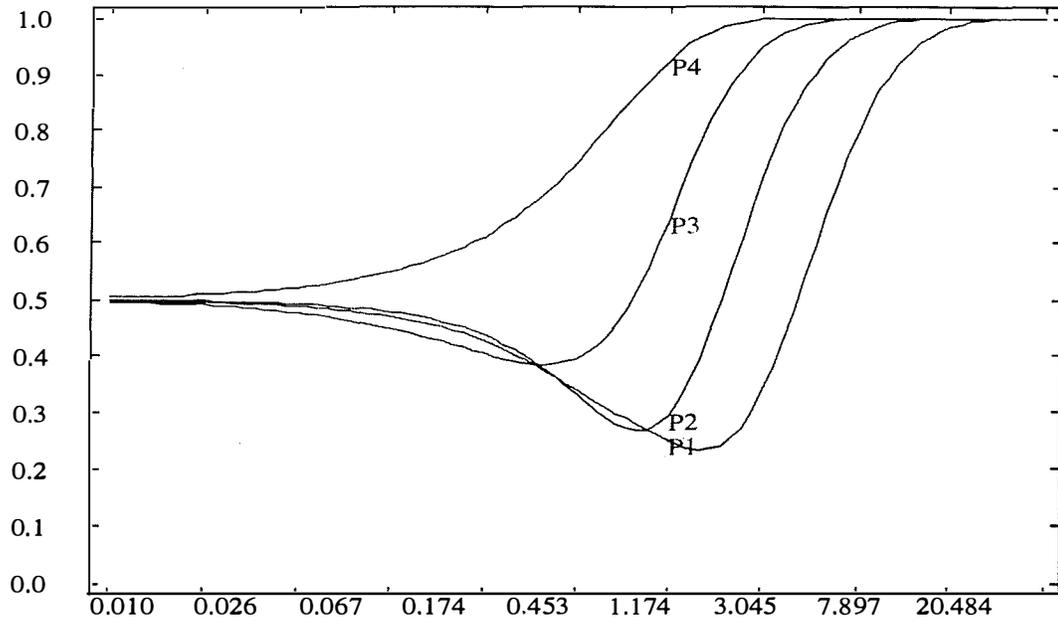


Figure 12: logit-AQRE graph for 4-move centipede game. Curve $P - i$ shows the conditional take probability at mode i .

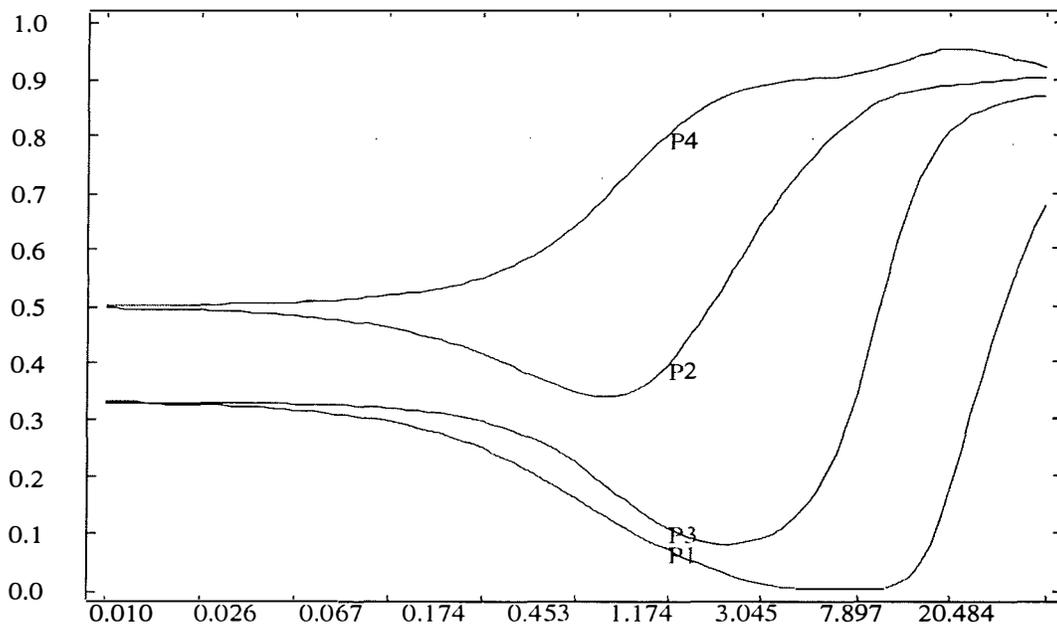


Figure 13: Normal form QRE graph for 4-move centipede game

		n	f_i	Normal Form	Agent model	Zauner model	2 Param model
Outcomes	T	20	.071	.203	.247	.214	.083
	PT	100	.356	.210	.235	.237	.293
	PPT	104	.370	.220	.351	.372	.488
	PPPT	43	.153	.216	.156	.167	.127
	PPPP	14	.050	.150	.011	.010	.009
Take Probs	T	281	.071	.203	.247	.214	.083
	T/P	261	.383	.264	.312	.301	.320
	T/PP	161	.644	.375	.677	.677	.718
	T/PPPP	57	.754	.595	.936	.939	.637
λ				.211	1.677	—	1.655
q				—	—	—	.96
$-\mathcal{L}^*$				437.5	425.0	418.3	402.5

Table 7

		n	f_i	Normal Form	Agent model	Zauner model	2 Param model
Outcomes	T	2	.007	.136	.237	.186	.029
	PT	18	.064	.137	.192	.173	.055
	PPT	56	.199	.135	.136	.138	.082
	PPPT	108	.384	.146	.205	.219	.446
	PPPPPT	71	.253	.162	.192	.231	.339
	PPPPPT	22	.078	.175	.036	.052	.047
	PPPPPP	4	.014	.109	.001	.001	.003
Take Probs	T	281	.007	.136	.237	.186	.029
	T/P	279	.065	.159	.252	.213	.056
	T/PP	261	.215	.185	.239	.215	.085
	T/PPP	205	.527	.246	.472	.436	.504
	T/PPPP	97	.732	.363	.838	.814	.772
	T/PPPPP	26	.846	.615	.979	.980	.409
λ				.08	.60		.666
q		0	0				.97
$-\mathcal{L}^*$				536.6	533.9	506.4	454.3

Table 8

REFERENCES

- Banks, J., C. Camerer, and D. Porter, "Experimental tests of Nash refinements in signaling games," *Games and Economic Behavior*, 6 (1994): 1-31.
- Beja, A., "Imperfect equilibrium," *Games and Economic Behavior*, 4 (1992): 18-36.
- Brandts, J., and C. A. Holt, "Forward Induction: Experimental evidence from two-stage games with complete information," mimeo, University of Virginia, 1990.
- Brandts, J., and C. A. Holt, "An experimental test of equilibrium dominance in signaling games," *American Economic Review*, 82 (1992): 1350-65.
- Brandts, J., and C. A. Holt, "Adjustment patterns and equilibrium selection in experimental signaling games," *International Journal of Game Theory*, 22 (1993): 279-302.
- Chen, Friedman and Thisse.
- Cooper, R., D. DeJong, R. Forsythe, and T. Ross, "Communication in the battle of the sexes game: some experimental results," *RAND Journal of Economics*, 20 (1989): 568-87.
- Cooper, R., D. DeJong, R. Forsythe, and T. Ross, "Selection criteria in coordination games: some experimental results," *The American Economic Review*, 80 (1990): 218-33.
- Fey, M., R. D. McKelvey, and T. R. Palfrey, "Experiments on the constant-sum centipede game," Social Science Working Paper No. 877, California Institute of Technology, (1993), forthcoming: *International Journal of Game Theory*.
- Gale, J., K. Binmore, and L. Samuelson, "Learning to be imperfect: The ultimatum game," SSRI Working Paper No. 9323, University of Wisconsin, (1993).
- Guth, W., P. Ockenfels, and M. Wendel, "Efficiency and trust in fairness? Multiperiod ultimatum bargaining experiments with an increasing cake," mimeo, Universitat Frankfurt/Main, (1991).
- Harsanyi, J., "Games with randomly disturbed payoffs," *International Journal of Game Theory*, 2 (1973): 1-23.
- Kohlberg, E. and J.-F. Mertens, "On the strategic stability of equilibria," *Econometrica*, 54 (1986): 1003-37.
- Kreps, D. and R. Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, 27 (1982): 253-79.

- McCabe, K., S. Rassenti and V. Smith, "Cooperation and the repeat interactions of anonymous pairings," mimeo, 1991 University of Minnesota.
- McFadden, D., "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, P. Zarembka, ed., (New York: Academic Press, 1974).
- McKelvey, R. D., and T. R. Palfrey, "An experimental study of the centipede game," *Econometrica*, 60 (1992): 803–36.
- McKelvey, R. D., and T. R. Palfrey, "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 7 (1995).
- Partow, J. and A. Schotter, "Does game theory predict well for the wrong reasons? An experimental investigation," C.V. Starr Center for Applied Economics Working Paper #93-46, New York University (1993).
- Rapoport, A., "Order of play in strategically equivalent games in extensive form," Department of Management and Policy, University of Arizona, (1993).
- Rosenthal, R., "A bounded-rationality approach to the study of noncooperative games," *International Journal of Game Theory*, 18 (1989): 273–92.
- Schotter, A., K. Weigelt, and C. Wilson, "A Laboratory Investigation of Multiperson Rationality and Presentation Effects," *Games and Economic Behavior*, 6 (1994): 445–68.
- Van Damme, E., *Stability and Perfection of Nash Equilibria*, Berlin: Springer-Verlag, 1987.
- Zauner, K., "Bubbles, speculations and a reconsideration of the centipede game experiment," mimeo, UC San Diego, (1993) (revised as "A Reconsideration of the Centipede Game Experiments," 1994).