

On the Empirical Content of Quantal Response Equilibrium

by

Philip A. Haile, Ali Hortaçsu, Grigory Kosenok

August 2003

COWLES FOUNDATION DISCUSSION PAPER NO. 1432



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.econ.yale.edu/>

On the Empirical Content of Quantal Response Equilibrium*

Philip A. Haile[†] Ali Hortaçsu[‡] Grigory Kosenok[§]

August 14, 2003

Abstract

The quantal response equilibrium (QRE) notion of McKelvey and Palfrey (1995) has recently attracted considerable attention, due largely to its widely documented ability to rationalize observed behavior in games played by experimental subjects. We show that this ability to fit the data, as typically measured in this literature, is uninformative. Without *a priori* distributional assumptions, a QRE can match any distribution of behavior by each player in any normal form game. We discuss approaches that might be taken to provide valid empirical evaluation of the QRE and discuss its potential value as an approximating empirical structure.

Keywords: quantal response equilibrium, testable restrictions, comparative statics

*We have benefited from comments from Simon Anderson, Jim Andreoni, Steve Berry, Colin Camerer, Jacob Goeree, Charlie Holt, John Kennan, Tom Palfrey, Larry Samuelson, Bill Sandholm and participants in the theory lunch at Wisconsin. Kun Huang provided capable research assistance. Financial support from the National Science Foundation and the Alfred P. Sloan Foundation is gratefully acknowledged.

[†]Department of Economics, Yale University, Box 208624, New Haven, CT 06520, philip.haile@yale.edu.

[‡]Department of Economics, University of Chicago, 1126 E. 59th St., Chicago, IL 60637, hortacsu@uchicago.edu.

[§]New Economic School, Nakhimovsky Prospekt 47, Suite 1721 (11), 117418 Moscow, Russia, gkosenok@nes.ru.

1 Introduction

The quantal response equilibrium (QRE) notion of McKelvey and Palfrey (1995) is a generalization of Nash equilibrium that allows imperfect optimizing behavior while maintaining the internal consistency of rational expectations. Roughly speaking, equilibrium is attained when each player employs a (nondegenerate) mixture of pure strategies, each of which is “close” to being best a response to the mixtures used by other players. The formal notion of “close” is based on realizations of random perturbations to the payoffs associated with each pure strategy a player can follow.¹

The QRE has intuitive appeal as a coherent formal solution concept allowing the possibilities that (a) the specified game imperfectly captures the true relation between strategies and payoffs, and/or (b) players may fail to play best responses, especially when the cost of using a suboptimal strategy is small. Much recent work also suggests that predictions of the QRE can match observed behavior well in a variety of experimental settings. In particular, when parameters of distributions of payoff perturbations are chosen so that the predicted distributions of outcomes fit the data as well as possible, the fit is often very good. McKelvey and Palfrey’s original paper demonstrated the ability of the QRE to explain departures from Nash equilibrium predictions in several games. Since then, the success of the QRE in rationalizing observed behavior has been demonstrated in a variety of experimental settings, including all-pay auctions (Anderson, Goeree and Holt (1998)), first-price auctions (Goeree, Holt and Palfrey (2002)), alternating-offers bargaining (Goeree and Holt (2000)), coordination games (Anderson, Goeree and Holt, forthcoming), and the “traveler’s dilemma” (Capra, Goeree, Gomez and Holt (1999), Goeree and Holt (2001)).² As the quotation below suggests, this success in explaining observed behavior has led many researchers in this area to view the QRE as a new standard:³

Quantal response equilibrium (QRE), a statistical generalization of Nash, almost always explains the direction of deviations from Nash and should replace Nash as the static benchmark that other models are compared to. (Camerer, Ho and Chong, 2001)

Given this recent work and its apparent influence, it is natural to ask how informative the ability of the QRE to fit the data really is. This is the subject of this note.

In the following section we define notation and review the definition of a QRE. In section 3 we describe the way the QRE notion has been applied to experimental data and what researchers

¹We review the formal definition of a QRE in the following section. The literature has considered generalizations to extensive form games (McKelvey and Palfrey, 1998) and games with continuous strategy spaces (e.g, Anderson, Goeree and Holt, forthcoming). We restrict attention to normal form games for simplicity.

²Dufwenberg, Gneezy, Goeree and Nagel (2002) suggest that they find an exception proving the rule, noting “Our results are unusual in that we document a feature of the data that is impossible to reconcile with the [QRE].”

³See also, e.g., the provocatively titled paper of Goeree and Holt (1999b).

have typically meant when they say that the QRE does a good job explaining observed behavior. We then present our main result: in any normal form game, any distribution of play by each player is consistent with a QRE. Any restriction on outcomes obtainable from a QRE comes only from *a priori* restrictions on the distributions of payoff perturbations. Put differently, for any game and any observed behavior, there is a distribution of payoff perturbations that will imply QRE behavior that matches the observed behavior perfectly. Hence, an evaluation of the “fit” of the QRE is uninformative unless one has *a priori* knowledge restricting the class of distributions one should consider. Examining fit with distributional assumptions chosen for convenience enables an evaluation only of the flexibility of the parametric family chosen.

We want to emphasize that this is not a critique of the QRE itself, only a critique of the approach taken to evaluate the predictive value of the QRE in much of the literature. This naturally leads to the question of how one might evaluate the QRE without arbitrary distributional assumptions as maintained hypotheses. This is one of two topics we take up briefly in a concluding section. Examining comparative statics predictions offers one promising approach. Taking this approach requires a different maintained assumption: that the distribution of payoff perturbations is constant or changes in known (*a priori*) ways as a game changes or across different games altogether. However, with such an assumption the QRE can provide testable restrictions. The second issue we raise is the potential value of the QRE as an empirical model—one that exploits the value of theory for providing relations between observables and the primitives of interest, but in a way that may be more robust than standard approaches relying on an assumption of Nash equilibrium.

2 Quantal Response Equilibrium

2.1 Model and Definition

Here we review the definition of a QRE, loosely following McKelvey and Palfrey (1995). We refer readers to their paper for additional detail, including discussion of the relation of the QRE to other solution concepts. Consider a finite n -person normal form game Γ . The set of pure strategies available to player i is denoted by $S_i = \{s_{i1}, \dots, s_{iJ_i}\}$, with $S = \times_i S_i$. Let Δ_i denote the set of all probability measures on S_i , i.e., the set of all functions $p_i : S_i \rightarrow [0, 1]$ satisfying $p(s_i) \geq 0 \forall s_i \in S_i$ and $\sum_{j=1}^{J_i} p_i(s_{ij}) = 1$. Let $\Delta \equiv \times_i \Delta_i$ denote the set of probability measures on S , with elements $p = (p_1, \dots, p_n)$. For simplicity, let p_{ij} represent $p_i(s_{ij})$.

Payoffs of Γ are given by functions $u_i(s_i, s_{-i}) : S_i \times_{j \neq i} S_j \rightarrow \mathbb{R}$. In the usual way, these payoff functions can be extended to the probability domain by letting $u_i(p) = \sum_{s \in S} p(s) u_i(s)$.

Hence, e.g., the argument s_{ij} of the payoff function $u_i(s_{ij}, \hat{s}_{-i})$ is reinterpreted as shorthand for a probability measure in Δ_i that places all mass on strategy s_{ij} . Finally, for every $p_{-i} \in \times_{j \neq i} \Delta_j$ and $p = (p_i, p_{-i})$, define $\bar{u}_{ij}(p) = u_i(s_{ij}, p_{-i})$ and $\bar{u}_i(p) = (\bar{u}_{i1}(p), \dots, \bar{u}_{iJ_i}(p))$.

The QRE is based on the introduction of payoff perturbations associated with each pure strategy of each player. For player i let

$$\hat{u}_{ij}(p) = \bar{u}_{ij}(p) + \epsilon_{ij}$$

where the vector of perturbations $\epsilon_i \equiv (\epsilon_{i1}, \dots, \epsilon_{iJ_i})$ is drawn from a joint density $f_i(\epsilon_{i1}, \dots, \epsilon_{iJ_i})$. For all i and j , ϵ_{ij} is assumed to have the same mean, which may be normalized to zero. Each player i is then assumed to use strategy s_{ij} if and only if

$$\hat{u}_{ij}(p) \geq \hat{u}_{ik}(p) \quad \forall k = 1, \dots, J_i. \quad (1)$$

Given a vector $u'_i = (u'_{i1}, \dots, u'_{iJ_i}) \in \mathbb{R}^{J_i}$, let

$$R_{ij}(u'_i) = \{\epsilon_i \in \mathbb{R}^{J_i} : u'_{ij} + \epsilon_{ij} \geq u'_{ik} + \epsilon_{ik} \quad \forall k = 1, \dots, J_i\}. \quad (2)$$

Conditional on the distribution p_{-i} characterizing the behavior of i 's opponents, $R_{ij}(\bar{u}_i(p))$ is the set of realizations of the vector ϵ_i that would lead i to choose action j (ignoring ties, which occur with probability zero).

Let

$$\sigma_{ij}(u'_i) = \int_{R_{ij}(u'_i)} f_i(\epsilon_i) d\epsilon_i$$

denote the probability of realizing a vector of perturbations in $R_{ij}(u'_i)$ and let $\sigma_i = (\sigma_{i1}, \dots, \sigma_{iJ_i})$. McKelvey and Palfrey (1995) call σ_i player i 's *statistical best response function* or *quantal response function*. Given the ‘‘baseline’’ payoffs $u_j(\cdot) \forall j$, a distribution of play by i 's opponents, and a joint distribution of i 's payoff perturbations, σ_i describes the probabilities with which each of i 's strategies will be chosen by i . A *quantal response equilibrium* is attained when the distribution of behavior of all players is consistent with their statistical best response functions. More precisely, letting $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\bar{u} = (\bar{u}_1, \dots, \bar{u}_n)$, a QRE is a fixed point of the composite function $\sigma \circ \bar{u} : \Delta \rightarrow \Delta$, which maps joint distributions over all players' pure strategies into statistical best responses for all players.

Definition 1 A quantal response equilibrium (QRE) is any $\pi \in \Delta$ such that for all $i \in 1, \dots, n$ and all $j \in 1, \dots, J_i$, $\pi_{ij} = \sigma_{ij}(\bar{u}_i(\pi))$.

⁴This rule is consistent with rational choice by i given the payoff function \hat{u}_{ij} if the following assumptions are added: (1) ϵ_i and ϵ_j are independent for $j \neq i$; (2) the ‘‘baseline’’ payoff functions $u_j(s_j, s_{-j})$ and densities f_j are common knowledge; and (3) the vector ϵ_i is i 's private information. As McKelvey and Palfrey (1995) show for a particular distribution of perturbations, under these assumptions a QRE is a Bayesian Nash equilibrium of the resulting perturbed game of incomplete information.

There are several possible interpretations of the QRE notion. One need not take the payoff perturbations literally. The idea that players use strategies that are only “close” to optimal has natural appeal, and the QRE offers a coherent formalization of this idea—one that closes the model of error-prone decisions with the assumption of rational expectations about opponents’ behavior. However, as McKelvey and Palfrey (1995) suggest, the payoff perturbations can have natural economic interpretations as well.⁵ Each ϵ_{ij} could reflect the error made by player i in calculating his expected utility from strategy j , due perhaps to unmodeled costs of information processing.⁶ Equilibrium then reflects the intuitive idea that players, while perhaps not always choosing payoff-maximizing strategies, are at least more responsive to larger differences in payoffs; i.e., deviations from optimizing behavior will be less likely when the costs of such deviations are large. Alternatively, ϵ_{ij} might reflect unmodeled determinants of i ’s utility from using strategy j . This interpretation is appealing in many applications since a model can, of course, only approximate a real economic environment. Furthermore, any true payoff function $\tilde{u}_i(s_{ij}, p_{-i})$ can be represented as the sum of an arbitrary baseline payoff $u_i(s_i, p_{-i})$ and a correction $\epsilon_{ij}(p_{-i}) = \tilde{u}_i(s_{ij}, p_{-i}) - u_i(s_{ij}, p_{-i})$. If the game underlying the baseline payoffs $u_i(s_{ij}, p_{-i})$ provides a good approximation to the truth, representing $\epsilon_{ij}(p_{-i})$ by a random variable that does not depend p_{-i} (as in the QRE) could provide a useful approximation.

2.2 Application and Evaluation

Following McKelvey and Palfrey (1995), application of the QRE to data from experiments has proceeded by first specifying the joint densities f_i (up to a finite-dimensional parameter) for all players. In every application we are aware of, it has been assumed that ϵ_{ij} is independently and identically distributed (iid) across all i and j . In most applications, it is assumed that each ϵ_{ij} is an independent draw from an extreme value distribution, yielding the familiar logit choice probabilities

$$p_{ij} = \frac{e^{\lambda \bar{u}_{ij}(p)}}{\sum_{k=1}^{J_i} e^{\lambda \bar{u}_{ik}(p)}}. \quad (3)$$

⁵See also Chen, Friedman and Thisse (1997) and the related discussion of random utility models in Manski (1977).

⁶There is, however, an issue if one wishes to view the behavioral assumption (1) as reflecting rational choice conditional on misperceptions of payoffs. Player i must observe only $\hat{u}_{ij}(p)$ for each pure strategy j , not the individual components $\bar{u}_{ij}(p)$ and ϵ_{ij} (otherwise he could just ignore ϵ_{ij}). However, i must know the distribution p_{-i} , and the correct anticipation of this equilibrium distribution requires the unnatural assumption that i know j ’s true payoffs $\hat{u}_{jk}(p)$ but not his own. In practice, however, p_{-i} might be learned from experience. Chen, Friedman and Thisse (1997) and Hofbauer and Sandholm (2002) explore this possibility for several classes of games.

With p observable, the unknown parameter λ is then easily estimated by maximum-likelihood,⁷ i.e., λ is chosen so that the implied QRE behavior matches actual behavior as closely as possible.

Typically the ability of the QRE to rationalize the data is then assessed based on the match between the observed probabilities on each pure strategy and those predicted by the QRE at the estimated parameter value(s).⁸ In some cases particular moments of the strategy distributions are compared; e.g., Goeree, Holt and Palfrey (2002) compare predicted and actual mean bids at each valuation for their first-price auction experiments. While formal testing is uncommon, visual inspection of histograms usually suggests an extremely good fit. Of course, since the QRE with a degenerate distribution of perturbations reduces to a Nash equilibrium, the fit *must* improve when one adds the freedom to choose the best fitting member of a parametric family like the extreme value. In fact, however, the fit is often greatly improved relative to Nash equilibrium. The following excerpt from Fey, McKelvey and Palfrey (1996, p. 286–287), which relies on this type of comparison in centipede games, is typical of the conclusions drawn from such an evaluation of fit:

Among the models we evaluate, the Quantal Response Equilibrium model best explains the data. It offers a better fit than the Learning model and, as it is an equilibrium model, is internally consistent. It also accounts for the pattern of increasing take probabilities within a match. These facts lend strong support to the Quantal Response Equilibrium model.

3 How Informative is Fit?

3.1 Result

One might naturally expect the QRE notion to impose considerable structure on the behavior consistent with equilibrium. As Goeree, Holt and Palfrey (2002) have suggested, the QRE requires a “consistency condition that the probabilities which determine expected utility . . . match the choice probabilities . . . that result from probabilistic choice.” Put differently, only probabilities that form a fixed point of the composite mapping $\sigma \circ \bar{u}$ can form a QRE, and experience suggests that fixed points are special.

⁷In the applications that have avoided the logit formulation, a power function specification has been used, but the approach is the same. In the logit specification, $1/\lambda$ is proportional to the variance of the payoff perturbations, with equilibrium behavior converging to a Nash equilibrium as $\lambda \rightarrow \infty$. In practice, estimates of λ often increase as players gain experience with the game (see, e.g., McKelvey and Palfrey (1995) or Fey, McKelvey and Palfrey (1996), although McKelvey, Palfrey and Weber (2000) find weaker evidence for this).

⁸See, e.g., Anderson, Goeree and Holt (1998a, 1998b), Baye and Morgan (2002), Camerer, Ho and Chong (2001), Capra et al (1999, 2002), Cason and Reynolds (2003), Fischbacher and Thoni (2002), Goeree and Holt (2000, 2001), Goeree, Holt and Palfrey (2002), McKelvey and Palfrey (1995, 1998), McKelvey, Palfrey and Weber (2000).

However, the freedom to choose the joint densities f_i to fit the data gives considerable flexibility to the QRE, particularly if one is unwilling to assume *a priori* that all payoff perturbations are iid. To see this, consider maintaining the assumption that perturbations are independent across players but relaxing the assumption of iid perturbations across a player's strategies in one of two ways. Let

$$\mathcal{I}_J = \{\text{joint pdfs for } J \text{ independent, mean-zero random variables}\}$$

$$\mathcal{S}_J = \{\text{joint pdfs for } J \text{ mean-zero random variables with identical marginal distributions}\}.$$

Joint densities f_i in the set \mathcal{I}_{J_i} satisfy an assumption that the ϵ_{ij} are independent across j , but need not imply identical marginal distributions across for each ϵ_{ij} . Joint densities f_i in \mathcal{S}_{J_i} allow dependence of ϵ_{ij} and ϵ_{ik} , $k \neq j$, but require them to have the same marginal distribution. The following result shows that when the assumption of iid perturbations is relaxed in either of these ways, the QRE imposes no restriction on behavior. For *any* game and *any* distribution of observed behavior on the interior of the J_i -dimensional simplex for each i , there exist densities from $\mathcal{I}_{J_i} \forall i$, as well as densities from $\mathcal{S}_{J_i} \forall i$, any of which will enable the QRE to match the distribution of behavior of each player *perfectly*.

Theorem 1 *Take any finite n -player normal form game Γ with $j = 1, \dots, J_i$ pure strategies for each player i . For any p on the interior of Δ ,*

- (i) *there exist joint probability density functions $f_i \in \mathcal{I}_{J_i} \forall i$ such that p forms a QRE of Γ .*
- (ii) *there exist joint probability density functions $f_i \in \mathcal{S}_{J_i} \forall i$ such that p forms a QRE of Γ .*

Proof: Given p_{-i} , the probability that player i plays action j in a QRE is given by

$$\sigma_{ij}(\bar{u}(p)) = \Pr \{ \epsilon_{ij} \geq \epsilon_{ik} + \bar{u}_{ik}(p) - \bar{u}_{ij}(p) \quad \forall k = 1, \dots, J_i \}.$$

Noting that $\bar{u}_{ij}(p)$ and $\bar{u}_{ik}(p)$ depend only on p_{-i} , let

$$H_i^{jk}(p_{-i}) = \bar{u}_{ik}(p) - \bar{u}_{ij}(p).$$

Part (i) [part (ii)] will then be proven if we can show that for each player i and any given $(p_{i1}, \dots, p_{iJ_i}) \in (0, 1)^{J_i}$, a density $f_i \in \mathcal{I}_{J_i}$ [$f_i \in \mathcal{S}_{J_i}$] can be found that implies

$$\Pr \left\{ \epsilon_{ij} \geq \epsilon_{ik} + H_i^{jk}(p_{-i}) \quad \forall k = 1, \dots, J_i \right\} = p_{ij} \quad j = 1, \dots, J_i \quad (4)$$

i.e., that the probabilities p_{ij} are in fact best responses given p_{-i} .

- (i) Suppose initially that all ϵ_{ij} are independent draws from two-point distributions such that

$$\epsilon_{ij} = \begin{cases} \alpha_j & \text{w.p. } q_j \\ -\frac{q_j}{1-q_j} \alpha_j & \text{w.p. } 1 - q_j \end{cases}$$

for some $\alpha_j > 0$ and $q_j \in (0, 1)$ to be determined below. By construction, each ϵ_{ij} has expectation zero. The event $\left\{ \epsilon_{ij} \geq \epsilon_{ik} + H_i^{jk}(p_{-i}) \right\}$ occurs with probability

$$\begin{aligned}
A_{jk} &= q_j q_k \mathbb{1} \left\{ \alpha_j - \alpha_k > H_i^{jk} \right\} \\
&+ q_j(1 - q_k) \mathbb{1} \left\{ \alpha_j + \frac{q_k}{1 - q_k} \alpha_k > H_i^{jk} \right\} \\
&+ (1 - q_j)q_k \mathbb{1} \left\{ -\frac{q_j}{1 - q_j} \alpha_j - \alpha_k > H_i^{jk} \right\} \\
&+ (1 - q_j)(1 - q_k) \mathbb{1} \left\{ -\frac{q_j}{1 - q_j} \alpha_j + \frac{q_k}{1 - q_k} \alpha_k > H_i^{jk} \right\}
\end{aligned} \tag{5}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and we have suppressed the dependence of H_i^{jk} on p_{-i} . Now begin by fixing $\alpha_{J_i} > 0$ and $q_{J_i} \in (0, 1)$ at arbitrary values. For any $q_{J_{i-1}} \in (0, 1)$ and all sufficiently large $\alpha_{J_{i-1}}$ we have

$$\begin{aligned}
\mathbb{1} \left\{ \alpha_{J_{i-1}} - \alpha_{J_i} > H_i^{(J_{i-1})J_i} \right\} &= 1 \\
\mathbb{1} \left\{ \alpha_{J_{i-1}} + \frac{q_{J_i}}{1 - q_{J_i}} \alpha_{J_i} > H_i^{(J_{i-1})J_i} \right\} &= 1 \\
\mathbb{1} \left\{ -\frac{q_{J_{i-1}}}{1 - q_{J_{i-1}}} \alpha_{J_{i-1}} - \alpha_{J_i} > H_i^{(J_{i-1})J_i} \right\} &= 0 \\
\mathbb{1} \left\{ -\frac{q_{J_{i-1}}}{1 - q_{J_{i-1}}} \alpha_{J_{i-1}} + \frac{q_{J_i}}{1 - q_{J_i}} \alpha_{J_i} > H_i^{(J_{i-1})J_i} \right\} &= 0
\end{aligned}$$

so that

$$A_{(J_{i-1})J_i} = q_{J_{i-1}}q_{J_i} + q_{J_{i-1}}(1 - q_{J_i}) = q_{J_{i-1}}.$$

Fix $\alpha_{J_{i-1}}$ at one such value, $\alpha_{J_{i-1}}^*$. Because the matrix of elements $H_i^{jk} \forall i, j$ is antisymmetric, we then also have $A_{J_i(J_{i-1})} = 1 - q_{J_{i-1}}$. Now consider selection of $\alpha_{J_{i-2}}$. As before, for any $q_{J_{i-2}} \in (0, 1)$, there exists sufficiently large $\alpha_{J_{i-2}}$ such that

$$\begin{aligned}
A_{(J_{i-2})(J_{i-1})} &= q_{J_{i-2}} \\
A_{(J_{i-2})J_i} &= q_{J_{i-2}} \\
A_{J_i(J_{i-2})} &= 1 - q_{J_{i-2}} \\
A_{(J_{i-1})(J_{i-2})} &= 1 - q_{J_{i-2}}.
\end{aligned}$$

Fix $\alpha_{J_{i-2}}$ at one such value $\alpha_{J_{i-2}}^*$. Proceeding in this fashion, given any $q_j \in (0, 1) \forall j$, we can choose each α_j so that

$$A_{jk} = \begin{cases} q_j & \text{if } j < k \\ 1 - q_k & \text{if } j > k. \end{cases} \tag{6}$$

This construction introduces a particular second-order stochastic dominance ordering of the random variables ε_{ij} . With this ordering, the event

$$\left\{ \varepsilon_{ij} \geq \varepsilon_{ik} + H_i^{jk} (p_{-i}) \quad \forall k = 1, \dots, J_i \right\}$$

is equivalent to the event $\{\varepsilon_{ij} > 0, \varepsilon_{ik} < 0 \quad \forall k < j\}$ when $j < J_i$, and to the event $\{\varepsilon_{ik} < 0 \quad \forall k < j\}$ when $j = J_i$ (realizations of ε_{ik} for $k > j$ do not matter). Because all ε_{ij} are independent, these events have probability $q_j \prod_{k < j} (1 - q_k)$ for $j < J_i$ and probability $\prod_{k < J_i} (1 - q_k)$ for $j = J_i$. So to satisfy (4), for each $j < J_i$ we set

$$q_j = \frac{p_{ij}}{1 - \sum_{k < j} p_{ik}}$$

(recall that the values of each q_j above were arbitrary and that q_{J_i} has been set to an arbitrary value). Note that $q_j \in (0, 1) \quad \forall j$ because $p_{ij} \in (0, 1) \quad \forall j$ and $\sum_{j=1}^{J_i} p_{ij} = 1$. Repeating this argument for each player i then shows that we can construct distributions for each ε_{ij} that yield any desired probabilities as a QRE if we ignore the fact that the definition of a QRE assumed continuously distributed perturbations.⁹ However, the mixtures of Dirac-delta functions used as densities here can be replaced with mixtures of univariate normal densities (with small variances) to obtain the same result. We show this in the appendix.¹⁰

(ii) Let ξ be uniformly distributed on $[-\kappa, \kappa]$, for some $\kappa > 0$, to be chosen below. For $j = 1, \dots, J_i$ define

$$\varepsilon_{ij} = \begin{cases} \xi + \delta_j & \text{if } \xi + \delta_j < \kappa \\ \xi + \delta_j - 2\kappa & \text{if } \xi + \delta_j > \kappa \end{cases} \quad (7)$$

where each δ_j is a distinct value in the interval $[0, 2\kappa]$ to be determined below. Each ε_{ij} is then uniformly distributed on $[-\kappa, \kappa]$. Fix δ_{J_i} at zero and, without loss of generality, impose $\delta_1 > \delta_2 > \dots > \delta_{J_i}$. Now suppose for the moment that $H_i^{jk} = 0$ for all j and k . Then for each j

$$\Pr\{\varepsilon_{ij} > \varepsilon_{ik}, \quad k = 1, \dots, J_i\} = \frac{\delta_{j-1} - \delta_j}{2\kappa}$$

where we define $\delta_0 = 2\kappa$. Setting these probabilities equal to the given values p_{i1}, \dots, p_{iJ_i} , we obtain a solution

$$\delta_j = \left(1 - \sum_{k=1}^j p_{ik} \right) 2\kappa \quad j = 1, \dots, J_i - 1. \quad (8)$$

⁹There are infinitely many other constructions since there are infinitely many ways to choose the parameters α_j (e.g., varying the starting value α_{J_i} in the proof, selecting different values of each α_j^* , or introducing the second-order stochastic dominance for any other ordering of the pure strategies).

¹⁰It is intuitive that mixtures of normals could approximate the two-point distributions above arbitrarily well. The appendix shows, however, that we can match the probabilities p_{ij} exactly.

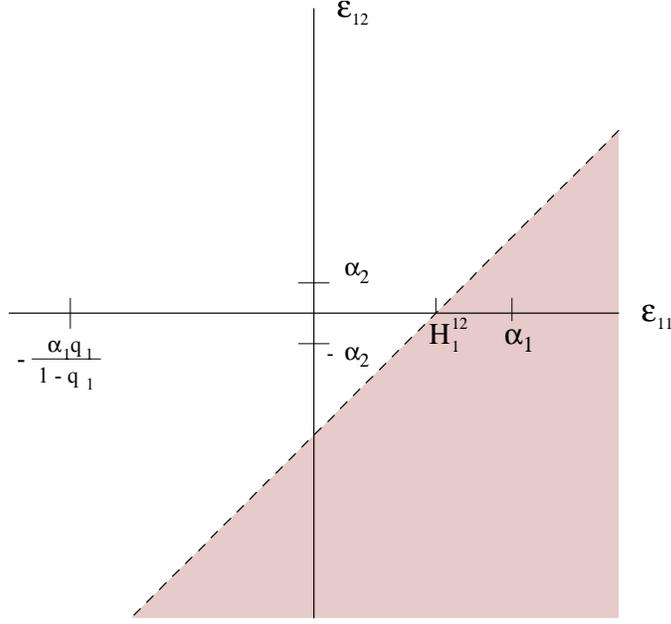


Figure 1: Example for part (i).

We now drop the assumption that each $H_i^{jk} = 0$. Note that when (8) holds, $|\epsilon_{ij} - \epsilon_{ik}| \geq 2\kappa (\min_{j=1, \dots, J_i} p_{ij})$ for all $j \neq k$.¹¹ Hence, by choosing

$$\kappa > \frac{\max_{k,j} |H_i^{jk}|}{2 \min_{j=1, \dots, J_i} p_{ij}}$$

(8) still gives

$$\Pr\{\epsilon_{ij} > \epsilon_{ik} + H_i^{jk} : k = 1, \dots, J_i\} = p_{ij} \quad \forall j \neq k.$$

Repeating this construction for every player completes the proof. \square

Figure 1 illustrates the example used to prove part (i) for the case of a game with two pure strategies. Here we have set $q_2 = 1/2$. Realizations of $(\epsilon_{i1}, \epsilon_{i2})$ in the shaded region (i.e., (α_1, α_2) or $(\alpha_1, -\alpha_2)$) lead to strategy s_{i1} being chosen over s_{i2} . This occurs with probability q_1 , which we are free to set equal to p_{i1} . Probability $(1 - q_1)$ is then put on the “balancing” point $\frac{-\alpha_1 q_1}{1 - q_1}$ to make $E[\epsilon_{i1}] = 0$.

Figure 2 illustrates the construction used for part (ii), again for the case $J_i = 2$. Here we have used the notation \oplus to represent addition on the circle running from $-\kappa$ to κ . The bold arc of this

¹¹When $\xi + \delta_i$ and $\xi + \delta_j$ both exceed κ or are both smaller than κ , this is immediate from (8). When $\xi + \delta_i > \kappa > \xi + \delta_j$, $|\epsilon_i - \epsilon_j| = |\delta_j - \delta_i + 2\kappa|$, and the claim then follows from (8).

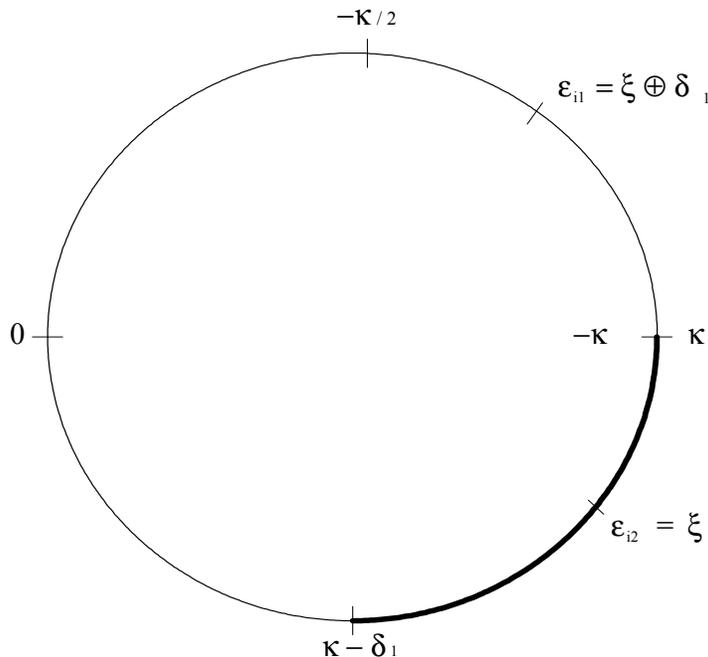


Figure 2: Example for part (ii).

circle indicates the set of realizations of ξ that yield $\epsilon_{i2} > \epsilon_{i1}$ (one such realization is shown). The length of this arc (divided by 2κ) determines the probability of this event which, for sufficiently large κ , is also the probability that s_{i2} is chosen over s_{i1} .

3.2 Discussion

Theorem 1 shows that if the assumption of iid payoff perturbations is relaxed, any distribution of behavior by each player is consistent with a QRE. Hence we pause to ask whether the assumption of iid perturbations is a natural *a priori* restriction on economic grounds. Little that is concrete can be said here given the ignorance of the true underlying structure that is implicit in representing payoffs with random shocks. This leads us to be cautious about placing any *a priori* restriction on the distribution of these shocks. Nonetheless, since the iid assumption does impose testable restrictions on outcomes, we briefly consider the plausibility of iid perturbations.

As McKelvey and Palfrey (1995) have observed, the iid assumption has the intuitive implication that better strategies (conditional on p_{-i}) are played with higher probabilities (cf. Rosenthal (1989)). However the plausibility of the iid assumption itself is not clear. The assumption of *identically distributed* perturbations might be a natural starting point; however, if perturbations

are taken to reflect players' misperceptions of true payoffs, it also seems natural that the variance of the perturbation to $\bar{u}_{ij}(p)$ might depend on the magnitude of $\bar{u}_{ij}(p)$. The plausibility of identically distributed perturbations is unclear to us if perturbations are viewed as corrections for the economist's misspecification of the game (due, e.g., to unmodeled costs/tastes for using different strategies). However, this interpretation alone seems inconsistent with widely documented evidence that behavior gets closer to Nash equilibrium as experimental subjects gain experience playing a game (cf., e.g., McKelvey and Palfrey (1992) and the references in footnote 7).

The assumption of *independent* perturbations seems much more questionable. An implication of independence that is unnatural for many games is that there is no sense in which payoffs from "similar" strategies (like contributing \$1 to a public good and contributing \$1.10) are subject to similar errors. Indeed, if the support of the iid perturbations is unbounded (as in the logit QRE), with positive probability players sometimes perceive *arbitrarily* large differences in payoffs between any pair of strategies, no matter how similar. Related to this, modifying a game by duplicating one strategy of one player will change the "real" outcomes of the game, much as in McFadden's well-known "red bus/blue bus" example. In fact, the IIA property itself carries over to the logit QRE. Such properties have, of course, long been a concern in the discrete choice literature, where considerable effort has been directed at developing tractable random utility models that relax the iid assumption. In the strategic context of the QRE, motivations for relaxing the iid assumption are similar. Hence, we conclude that on economic grounds the iid assumption is a questionable *a priori* restriction for many games.

4 Evaluating the Empirical Value of the QRE

The theorem above indicates that without an economic foundation for *a priori* restrictions on the distribution of payoff perturbations (or, equivalently, directly on quantal response functions), the widely documented ability of the QRE to fit the data is uninformative. The common practice of examining fit using a particular specification of the QRE can reveal whether a particular parametric family (e.g., the extreme value) is sufficiently flexible to allow a good approximation of the data, but can reveal nothing about the value of the QRE notion itself. This obviously does not mean that the QRE is useless, nor even that it is without empirical content. However, it does raise several questions that are the subject of this concluding section, which we hope will stimulate further work:

- (1) Are there approaches for meaningful empirical evaluation of the QRE hypothesis?
- (2) If so, what evidence do we have?
- (3) If not, or if the evidence is not supportive, can the QRE be empirically useful nonetheless?

4.1 Comparative Statics

One possible approach for evaluating the predictive value of the QRE involves examining changes in behavior as payoffs or other elements of a game change, i.e., testing comparative statics predictions. It should immediately be emphasized that doing this requires maintaining an assumption that the distributions of perturbations are fixed (or change in known ways) as a game changes. If one is free to choose a new distribution of perturbations for each game, Theorem 1 ensures that the QRE can match behavior perfectly in every game. A maintained assumption of a fixed distribution of perturbations may be more difficult to justify in some applications than others. For example, contrary to this assumption, in some applications one might expect errors made in assessing payoffs to have variances that proportional to (or at least increasing in) the payoffs themselves. Alternatively, variation in the complexity of games might suggest different distributions of “errors” for different games. Without a clear economic foundation for the QRE perturbations,¹² however, this question is not easily resolved.

With this important caveat, comparative statics predictions offer falsifiable restrictions of the QRE that can be tested empirically. Little attention has been given to such testing thus far. While many papers have examined the fit of the QRE in different treatments (e.g., varying payoffs) or in different games altogether, with few exceptions a new value of the distributional parameter(s) is estimated each time, rendering the fit in comparative statics uninformative. Two notable exceptions are Capra, Goeree, Gomez and Holt (1999) and Goeree and Holt (1999) which demonstrate that the QRE with a single distribution of perturbations can rationalize observed comparative statics in the “traveller’s dilemma” game and a coordination game, respectively.

However, the suggestion that the QRE with a fixed distribution of perturbations may have good predictive power across a variety of experimental settings is not widely supported. First, as pointed in footnote 7 above, the variance of the perturbations that rationalizes the data often appears to decline over time. Capra, Goeree, Gomez and Holt (1999), for example, use only data from the last three periods of their experimental treatments to estimate the error precision, presumably for this reason. Second, even if one views the fixed-distribution QRE as a theory of “steady state” behavior, an examination of the experimental results suggests that in fact very different distributions are estimated for different games. McKelvey and Palfrey (1995), for example, estimated the logit model (recall (3)) separately for data from a number of different experiments and obtained estimates of λ varying from 0.25 to 4.64. Wide ranges of estimates are obtained in

¹²Aside from analytical convenience, there is little justification offered in the literature for the *a priori* restriction that payoff perturbations are independent and/or identically distributed or to belong to a particular parametric family. An exception is the work of Anderson, Goeree, and Holt (1999), which develops a theoretical foundation for the logit specification of the QRE as the limit of a noisy directional learning process.

subsequent work as well (e.g., Fey, McKelvey and Palfrey (1996), McKelvey, Palfrey and Weber (2000), Camerer, Ho and Chong (2002)). In a recent paper, Goeree and Holt (2002) formally test and reject the hypothesis that the distribution of perturbations is constant across auctions with different distributions of valuations. Considering a different type of comparative static, Dufwenberg, Gneezy, Goeree and Nagel (2002) find that changes in behavior observed when a price floor is imposed in a pricing game cannot be explained by their QRE specification.

While the evidence thus far is mixed, evaluating comparative statics predictions of the QRE (fixing the perturbation distribution) has not actually been a focus of the literature, perhaps because the importance of this type of evaluation was not fully appreciated. Additional formal testing is needed to better understand whether the QRE can predict outcomes for a range of games with a fixed distribution of payoff perturbations. Rejecting this hypothesis, of course, would not imply a rejection of the QRE, since tests of comparative statics predictions will necessarily join the QRE hypothesis with hypotheses about how the distribution of perturbations varies as the game changes. It is also worth noting that even if comparative statics predictions of the QRE are formally rejected, the QRE may nonetheless serve better for out-of-sample prediction than alternative solution concepts. Investigating this possibility in games for which there is a clear motivation for such out-of-sample prediction is another potentially useful direction for further research.¹³

4.2 QRE as an Empirical Structure

Consider for a moment the standard additive random utility model (ARUM)¹⁴ in which consumer i 's utility from good j (with characteristics x_j) is given by

$$u_{ij} = g(x_j, \beta) + \epsilon_{ij}$$

for some function g . This formulation is analogous to that of the QRE for a “one player game.”

Corollary 1 *(i) Given any fixed set of choices $j = 1, \dots, J$, any function g , and any parameter vector β , there are no observed choice probabilities inconsistent with the additive random utility model. (ii) the ARUM is nonparametrically unidentified from individual choices from a fixed choice set.*

¹³A testable restriction we have not mentioned is that with perturbations that are independent across players, variation in strategies chosen should be independent across players as well. Of course, independent perturbations are not required by the QRE and may be unnatural in some cases, particularly if we interpret payoff perturbations as corrections for unmodeled elements of the game actually being played. Further, the independence restriction can hold even if a particular specification of the QRE poorly captures actual behavior.

¹⁴See, e.g., Anderson, DePalma and Thisse (1992). This framework includes the standard multinomial logit and probit as well as richer models like the nested logit (e.g., McFadden (1978), Cardell (1997)) or random coefficients probit/logit (e.g., Hausman and Wise (1978), Boyd and Mellman (1980), Cardell and Dunbar (1980), Berry, Levinshohn and Pakes (1995)).

Proof: Part (i) follows directly from the proof of Theorem 1, with the mean utility $g(x_j, \beta)$ playing the role of the conditional (on p_{-i}) mean payoff $\bar{u}_{ij}(p)$ in the strategic setting. Part (ii) then follows, since for every $g(\cdot, \beta)$ there is some joint distribution of the $\epsilon_{ij} \forall j$ leading to a perfect fit.¹⁵ \square

This corollary does not imply any critique of the extensive empirical literature using the ARUM.¹⁶ Indeed, this is our point. Despite our results regarding empirical *testing* of the QRE, the QRE may be useful in the same way that the ARUM is, i.e., as an empirical structure for uncovering features of payoffs from field data. Such a structure would exploit theory to obtain relationships between observables and primitives of interest, but in a way that might be more robust than standard structural approaches based on an assumption of Nash equilibrium.

Surprisingly, little attention has been given to this possibility.¹⁷ Examining the usefulness of such an approach seems a natural and valuable direction for work in experimental economics. In particular, estimates of payoff parameters obtained by interpreting observed behavior through the QRE structure could be compared to the known true underlying parameters. As in the discrete choice literature, one may find that although models with iid shocks can have undesirable properties, richer models analogous to, e.g., the nested logit or random coefficients probit/logit can be quite useful. Note that the analog of variation in choice sets here (cf. footnote 16) is variation in the underlying game—e.g., variation in payoffs, the strategies available, or the number of players. Just as examining comparative statics is important to valid direct evaluation of the QRE, examination of data from different games will be important if one is to obtain estimates through the QRE structure that are not merely artifacts of *a priori* distributional assumptions.¹⁸ This is an area we hope to explore in future work.

¹⁵Manski (1988) states a version of this result for the case of binary choice.

¹⁶We are not aware of attempts to evaluate the ARUM based on its ability to fit the data. Part (ii), however, does suggest that one should view with caution estimates of these models that do not exploit variation in choice sets—variation in, e.g., which choices are available, the characteristics of different choices, or the prices of different choices. Such variation is analogous to the variation in games necessary to test comparative statics predictions of the QRE.

¹⁷Bajari (1999), Signorino (1999), Bajari and Hortaçsu (2001), Goeree, Holt and Palfrey (2002), and Seim (2002) are the exceptions we are aware of. McKelvey and Palfrey (1995, p. 7) mention using the QRE for estimation, but apparently only meant fitting a parametric specification of the perturbation distribution to data.

¹⁸In parametric empirical applications of the QRE, the scale of the perturbation distribution generally will not be identified (just as in discrete choice models). A single scale normalization imposed in estimation using data from different games is analogous to the assumption of a fixed perturbation distribution across games needed to provide a meaningful evaluation of comparative statics predictions (cf. section 4.1).

Appendix

Here we show that a variation on the example constructed in the proof of part (i) of Theorem 1 can deliver the same result with continuously distributed perturbations. We begin with the values of α_j constructed in the text. Now, however, let each ϵ_{ij} be drawn independently from a mixture of two normal distributions, with mixing weights q_j and $(1 - q_j)$, means α_j and $\frac{q_j}{1 - q_j}\alpha_j$, and common variance σ^2 . Letting ϕ denoting the standard normal density, the density of ϵ_{ij} is then

$$f(\epsilon; q_j, \sigma) = q_j \phi\left(\frac{\epsilon - \alpha_j}{|\sigma|}\right) + (1 - q_j) \phi\left(\frac{\epsilon - \frac{q_j}{1 - q_j}\alpha_j}{|\sigma|}\right).$$

The parameters σ and $q_j \forall j$, will be determined below. Note that for any q_j and $\sigma > 0$, ϵ_{ij} has mean zero.

Given the values of H_i^{jk} and the distribution p_{-i} , equation (2) defines a region $R_j \subset \mathbb{R}^{J_i}$ of realizations of the vector ϵ_i that would lead i to choose strategy j . For $\sigma \neq 0$ let

$$G(\sigma, q) = \begin{bmatrix} G_1(\sigma, q) \\ \vdots \\ G_{J_i-1}(\sigma, q) \end{bmatrix} = \begin{bmatrix} p_1 - \int_{R_1} f(\epsilon_1; q_1, \sigma) \dots f(\epsilon_{J_i}; q_{J_i}, \sigma) d\epsilon_1 \dots d\epsilon_{J_i} \\ \vdots \\ p_{J_i-1} - \int_{R_{J_i-1}} f(\epsilon_1; q_1, \sigma) \dots f(\epsilon_{J_i}; q_{J_i}, \sigma) d\epsilon_1 \dots d\epsilon_{J_i} \end{bmatrix}. \quad (9)$$

For $\sigma = 0$, we let the normal densities collapse and define

$$G(0, q) = \begin{bmatrix} p_1 - q_1 \\ \vdots \\ p_j - q_j \prod_{k < j} (1 - q_k) \\ \vdots \\ p_{J_i-1} - q_{J_i-1} \prod_{k < J_i-1} (1 - q_k) \end{bmatrix}. \quad (10)$$

To match arbitrary probabilities (p_1, \dots, p_{J_i}) in the interior of Δ using the normal mixtures, we will show that for small σ , we can choose $q = (q_1, \dots, q_{J_i-1})$ to solve the system

$$G(\sigma, q) = 0_{J_i-1} \quad (11)$$

where 0_{J_i-1} denotes a $(J_i - 1)$ -vector of zeros.¹⁹ The example in the text showed that there is a solution q^0 to (11) when $\sigma = 0$. We show that for small $\sigma > 0$ there is still a solution. This follows immediately from three lemmas.²⁰

¹⁹Note that we use the identities $p_{J_i} = 1 - \sum_{j < J_i} p_j$ and $\int_{R_{J_i}} f(\epsilon_1; q_1, \sigma) \dots f(\epsilon_{J_i}; q_{J_i}, \sigma) d\epsilon_1 \dots d\epsilon_{J_i} = 1 - \sum_{j < J_i} \int_{R_j} f(\epsilon_1; q_1, \sigma) \dots f(\epsilon_{J_i}; q_{J_i}, \sigma) d\epsilon_1 \dots d\epsilon_{J_i}$ to obtain a $(J_i - 1) \times (J_i - 1)$ system.

²⁰Lemmas 2 and 3 are related to standard implicit function theorems. Here, however, we are not interested in ensuring existence of implicit *functions*, but only implicit *solutions*. Nor are we interested in differentiability of these solutions. Because of this we are able to prove the results under weaker assumptions; in particular, we make no assumption about differentiability with respect to the parameter x . Proving Lemma 3, then, requires a different approach from that taken to prove existence of a solution in standard multi-dimensional implicit function theorems.

Lemma 1 *i. G is continuous in a neighborhood of $(0, q^0)$.*

ii. The matrix of partial derivatives

$$\nabla_q G(0, q) = \begin{bmatrix} \frac{\partial}{\partial q_1} G_1(0, q) & \cdots & \frac{\partial}{\partial q_{J_i-1}} G_1(0, q) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial q_1} G_{J_i-1}(0, q) & \cdots & \frac{\partial}{\partial q_{J_i-1}} G_{J_i-1}(0, q) \end{bmatrix}$$

is continuous in q .

iii. $|\nabla_q G(0, q^0)| \neq 0$.

Proof. (i) Continuity follows from well known properties of normal random variables.

(ii) From (10),

$$\left. \frac{\partial G_j(0, q)}{\partial q_i} \right| = \begin{cases} q_j \prod_{k < j, k \neq i} (1 - q_k) & i < j \\ -1 & i = j = 1 \\ -\prod_{k < j} (1 - q_k) & i = j \neq 1 \\ 0 & i > j. \end{cases} \quad (12)$$

Continuity in each q_ℓ is immediate.

(iii) From (12), the matrix $\nabla_q G(0, q^0)$ is lower triangular with nonzero diagonal elements, implying a nonzero determinant. \square

Lemma 2 *Consider any continuous function $F(x, y) : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$, with \mathbb{X} and \mathbb{Y} open subsets of \mathbb{R}^m and \mathbb{R} , respectively. Suppose that $F(x^0, y^0) = 0$ for some $(x^0, y^0) \in \mathbb{X} \times \mathbb{Y}$ and that $\left. \frac{\partial F(x^0, y)}{\partial y} \right|_{y=y^0} \neq 0$. Then there exists $\delta > 0$ such that for all $x \in B_\delta(x^0)$ there is some $y(x) \in \mathbb{Y}$ such that $F(x, y(x)) = 0$.*

Proof. Suppose that $\left. \frac{\partial F(x^0, y)}{\partial y} \right|_{y=y^0} > 0$. Then for sufficiently small $\epsilon > 0$,

$$F(x^0, y^0 - \epsilon) < 0 < F(x^0, y^0 + \epsilon).$$

By continuity, there exists $\delta > 0$ such that for all $x \in B_\delta(x^0)$,

$$F(x, y^0 - \epsilon) < 0 < F(x, y^0 + \epsilon).$$

By continuity, then, there is some $y(x) \in (y^0 - \epsilon, y^0 + \epsilon)$ satisfying $F(x, y(x)) = 0$. The argument is analogous if $\left. \frac{\partial F(x^0, y)}{\partial y} \right|_{y=y^0} < 0$. \square

Lemma 3 Consider any continuous function $F(x, y) : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}^J$, with \mathbb{X} and \mathbb{Y} open subsets of \mathbb{R}^m and \mathbb{R}^J , respectively. Suppose that $F(x^0, y^0) = 0_J$ for some $(x^0, y^0) \in \mathbb{X} \times \mathbb{Y}$, $|\nabla_y F(x^0, y^0)| \neq 0$, and that $\nabla_y F(x^0, y)$ is continuous in $y \in B_{\epsilon_0}(y^0)$ for some $\epsilon_0 > 0$. Then there exists $\delta > 0$ such that for all $x \in B_\delta(x^0)$ there is some $y(x) \in \mathbb{Y}$ such that $F(x, y(x)) = 0$.

Proof. The proof proceeds by induction. For $J = 1$, the result is given in Lemma 2. So suppose the claim is true for $J = k - 1$ and that the hypotheses above hold for $J = k$. Since $|\nabla_y F(x^0, y^0)| \neq 0$, there is some $c \in \{1, \dots, k\}$ such that the minor $\nabla_y^{k,c} F(x^0, y^0)$ obtained by dropping row k and column c from $\nabla_y F(x^0, y^0)$ is also nonsingular. Without loss of generality, let $c = k$. Now consider the following $(k - 1) \times (k - 1)$ system in the variables y_1, \dots, y_{k-1} :

$$\begin{bmatrix} F_1(x, y_1, \dots, y_{k-1}, y_k) \\ \vdots \\ F_{k-1}(x, y_1, \dots, y_{k-1}, y_k) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}.$$

We know that for $j = 1, \dots, k$, $F_j(x^0, y_1^0, \dots, y_{k-1}^0, y_k^0) = 0$. Further, by the induction hypothesis (replacing x with the pair (x, y_k)) there is some $\delta_k > 0$ such that for all $(x, y_k) \in B_{\delta_k}(x^0, y_k^0)$ there exist $y_j(x, y_k)$, $j = 1, \dots, k - 1$, which solve

$$F_j(x, y_1(x, y_k), \dots, y_{k-1}(x, y_k), y_k) = 0 \quad j = 1, \dots, k - 1. \quad (13)$$

The proof will then be completed if we can show that for all x in a neighborhood $B_\delta(x^0)$ there is some y_k such that $(x, y_k) \in B_{\delta_k}(x^0, y_k^0)$ and

$$\Phi(x, y_k) \equiv F_k(x, y_1(x, y_k), \dots, y_{k-1}(x, y_k), y_k) = 0. \quad (14)$$

Given the continuity of F_k and the fact that δ may be chosen arbitrarily small, this result will follow from Lemma 2 if we can show that $\frac{\partial \Phi(x^0, y_k)}{\partial y_k}$ is nonzero at $y_k = y_k^0$. Differentiating (14) gives

$$\begin{aligned} \frac{\partial \Phi(x^0, y_k)}{\partial y_k} &= \sum_{j=1}^{k-1} \frac{\partial}{\partial y_j} F_k(x^0, y_1(x^0, y_k), \dots, y_{k-1}(x^0, y_k), y_k) \frac{\partial y_j(x^0, y_k)}{\partial y_k} \\ &\quad + \frac{\partial}{\partial y_k} F_k(x^0, y_1(x^0, y_k), \dots, y_{k-1}(x^0, y_k), y_k) \end{aligned} \quad (15)$$

in a neighborhood of $y_k = y_k^0$. Since $\nabla_y F(x^0, y)$ is continuous in $y \in B_{\epsilon_0}(y^0)$, a standard implicit function theorem ensures that each $y_j(x^0, y_k)$ is differentiable with respect to y_k in a neighborhood of y_k^0 , implying that the derivative in (15) exists. To see that it must be nonzero, note that we may differentiate each side of (13) with respect to y_k at $y_k = y_k^0$ to obtain

$$\begin{aligned} \sum_{\ell=1}^{k-1} \frac{\partial}{\partial y_\ell} F_j(x^0, y_1(x^0, y_k^0), \dots, y_{k-1}(x^0, y_k^0), y_k^0) \frac{\partial y_\ell(x^0, y_k^0)}{\partial y_k} \\ + \frac{\partial}{\partial y_k} F_j(x^0, y_1(x^0, y_k^0), \dots, y_{k-1}(x^0, y_k^0), y_k^0) = 0 \quad j = 1, \dots, k - 1. \end{aligned} \quad (16)$$

If $\frac{\partial \Phi(x^0, y_k)}{\partial y_k} = 0$ at $y_k = y_k^0$, (15) and (16) would imply that we could express last column of the matrix $\nabla_y F(x^0, y^0)$ as linear combination of first $K - 1$ columns, which is a contradiction. \square

References

- Anderson, S., DePalma, A. and Thisse, J. *Discrete Choice Theory of Product Differentiation*, Cambridge: MIT Press, 1992.
- Anderson, S., Goeree, J. and Holt, C. "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction," *Journal of Political Economy*, 106(4), 828-853, August 1998.
- Anderson, S., Goeree, J. and Holt, C. "Minimum Effort Coordination Games: Stochastic Potential and the logit Equilibrium" *Games and Economic Behavior*, forthcoming.
- Anderson S., Goeree J. and Holt C. "The logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," *Southern Economic Journal*, 69(1), 21-47, July 2002.
- Anderson, S., Goeree, J. and Holt, C. "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games," *Journal of Public Economics*, 70, 297-322, November 1998.
- Anderson, S., Goeree, J. and Holt, C. "Stochastic Game Theory: Adjustment to Equilibrium Under Noisy Directional Learning," working paper, 1999.
- Bajari, P. and Hortacısu, A. "Auction Models When Bidders Make Small Mistakes: Consequences for Theory and Estimation," working paper, August 2001.
- Bajari, P. "Econometrics of Sealed-bid Auctions," *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*, 1998.
- Baye, M. and Morgan, J. "Price Dispersion in the Lab and on the Internet: Theory and Evidence," working paper, 2002.
- Berry, S., Levinsohn, J. and Pakes, A. "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841-890, July 1995.
- Boyd, J. and Mellman, J. "The Effect of Fuel Economy Standards on the U.S. Automotive Market: A Hedonic Demand Analysis," *Transportation Research*, 14A(4-5), 367-378, 1980.
- Camerer, C. F., T. Ho, and J. K. Chong, "Behavioral Game Theory: Thinking, Learning, and Teaching," working paper, Caltech, 2001.
- Camerer, C. F., T. Ho, and J. K. Chong, "A Cognitive Hierarchy Theory of One-Shot Games: Some Preliminary Results," working paper, Caltech, 2002.
- Capra, M., Goeree, J., Gomez, R., and Holt, C. "Anomalous Behavior in a Travelers Dilemma?" *American Economic Review*, 89(3), 678-690, September 1999.
- Capra, M., Gomez, R., Goeree, J. and Holt, C. "Learning and Noisy Equilibrium Behavior in a Study of Imperfect Price Competition," *International Economic Review*, 43(3), 613-636, August 2002.
- Cardell, N.S. "Variance Components Structures for the Extreme Value and Logistic Distributions, with Applications to Models of Heterogeneity," *Econometric Theory*, 13(2), 185-213, April 1997.
- Cardell, N.S. and Dunbar, F. "Measuring the Societal Impacts of Automobile Downsizing," *Transportation Research*, 14A(5-6), 423-434, 1980.

- Cason, T.N. and Reynolds, S. "Bounded Rationality in Laboratory Bargaining with Asymmetric Information," working paper, 2003.
- Chen, H., Friedman, J. and Thisse, J. "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, 18, 32–54, 1997.
- Dufwenberg, M., Gneezy, U., Goeree, J. and Nagel, R. "Price Floors and Competition," working paper, 2002.
- Fey, M., McKelvey, R.D. and Palfrey, T.R. "An Experimental Study of Constant-Sum Centipede Games," *International Journal of Game Theory*, 25(3), 269–287, 1996.
- Fischbacher, U. and Thoni, C. "Excess Entry in an Experimental Winner-Take-All-Market", working paper, 2002.
- Goeree, J. and Holt, C. "An Experimental Study of Costly Coordination," working paper, 1999a.
- Goeree, J. and Holt, C. "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," *European Economic Review*, 44, 1079-1089, 2000.
- Goeree, J. and Holt, C. "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, 10564-10567, September 1999b.
- Goeree, J. and Holt, C., "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402-1422, December 2001.
- Goeree, J., Holt, C. and Palfrey, T., "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), 247-272, May 2002.
- Hausman, J.A. and Wise, D. A. "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences," *Econometrica*, 46(2), 403–426, March 1978.
- Hofbauer, J. and Sandholm, W. "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 70(6), 2265–2294, November 2002.
- Manski, C.F. "The Structure of Random Utility Models," *Theory and Decision*, 8, 229–254, 1977.
- Manski, C.F. "Identification of Binary Response Models," *Journal of the American Statistical Association*, 83, 729–738, 1988.
- McFadden, D. "Modelling the Choice of Residential Location," in A. Karlqvist et al., eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North Holland, 1978.
- McKelvey, R.D. and Palfrey, T.R. "An Experimental Study of the Centipede Game," *Econometrica*, 60(4), 803–836, July 1992.
- McKelvey, R. D. and Palfrey, T. R. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6-38, 1995.
- McKelvey, R.D. and Palfrey, T.R. "Quantal Response Equilibrium in Extensive Form Games," *Experimental Economics*, 1, 9–41, 1998.
- McKelvey, R. D. and Palfrey, T. R. and Weber, R. A. "The Effect 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, 2000.

Rosenthal, R. "A Bounded-Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273–291, 1989.

Seim, K. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," working paper, Stanford, 2002.

Signorino, C. "Strategic Interaction and the Statistical Analysis of International Conflict," *American Political Science Review*, 93(2), 279–297, June 1999.