

On the Empirical Content of Quantal Response Equilibrium

by Philip A. Haile, Ali Hortaçsu and Grigory Kosenok*

September 17, 2007

Abstract

The quantal response equilibrium (QRE) notion of Richard D. McKelvey and Thomas R. Palfrey (1995) has recently attracted considerable attention, due in part to its widely documented ability to rationalize observed behavior in games played by experimental subjects. However, even with strong *a priori* restrictions on unobservables, QRE imposes no falsifiable restrictions: it can rationalize any distribution of behavior in any normal form game. After demonstrating this, we discuss several approaches to testing QRE under additional maintained assumptions. (JEL C72, C52, C90)

The quantal response equilibrium (QRE) notion of McKelvey and Palfrey (1995) can be viewed as an extension of standard random utility models of discrete (“quantal”) choice to strategic settings, or as a generalization of Nash equilibrium that allows noisy optimizing behavior while maintaining the internal consistency of rational expectations. Formally, QRE is based on the introduction of random perturbations to the payoffs associated with each action a player can take.¹ Realizations of these perturbations affect which action is the best response to the equilibrium distribution of opponents’ behavior.

¹We give a more complete discussion in the following section. The literature has considered generalizations of the QRE to extensive form games (McKelvey and Palfrey (1998)) and games with continuous strategy spaces (e.g, Simon P. Anderson, Jacob K. Goeree and Charles A. Holt (2002)). We restrict attention to normal form games for simplicity.

Both interpretations of QRE have strong intuitive appeal, and much recent work has shown that QRE can rationalize behavior in a variety of experimental settings where Nash equilibrium cannot. In particular, when parameters (of the distributions of payoff perturbations) are chosen so that the predicted distributions of outcomes match the data as well as possible, the fit is often very good. McKelvey and Palfrey’s original paper demonstrated the ability of QRE to explain departures from Nash equilibrium behavior in several games. Since then, the success of QRE in matching 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-offer bargaining (Goeree and Holt (2000)), coordination games (Anderson, Goeree and Holt (2001)), and the “traveler’s dilemma” (C. Monica Capra, Goeree, Rosario Gomez, and Holt (1999), Goeree and Holt (2001)).² The quotation below, from Colin F. Camerer, Teck-Hua Ho and Juin Kuan Chong (2004), suggests the impact this evidence has had: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 to which other models are routinely compared.

Given this recent work and its influence, it is natural to ask how informative the ability of QRE to fit the data really is. Our first result provides a strong negative answer to this question for a type of data often considered in the literature: QRE is not falsifiable in any normal form game, even with significant *a priori* restrictions on payoff perturbations. In particular, any behavior can be rationalized by a QRE, even when each player’s payoff perturbations are restricted to be independent across actions or to have identical marginal distributions for all actions. Hence, an evaluation of fit in a single game (no matter the number of replications) is uninformative without strong *a priori* restrictions on distributions of payoff perturbations.

²Martin Dufwenberg, et al. (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 (1999).

This first result implies no critique of the QRE notion, but merely points to the challenge of developing useful approaches to testing the QRE hypothesis. Testing requires maintained hypotheses beyond those of the QRE notion itself. Success in fitting data will therefore be informative only to the degree that the additional maintained assumptions place significant restrictions on the set of outcomes consistent with QRE. On the other hand, failures of QRE to fit the data (e.g., Goeree, Holt, and Palfrey (2003); Ho, Camerer and Chong (2007)) will be informative about the QRE notion to the extent that one has confidence in the auxiliary assumptions that provide falsifiable restrictions. Useful testing approaches must trade off these limitations, and the most appropriate approach may depend on the application. Below we discuss several promising approaches. Each maintains restrictions on how distributions of payoff perturbations can vary across related normal-form games, leading to falsifiable comparative statics predictions.

In the following section we define notation, review the definition of QRE, and discuss common applications of the QRE in the literature. We then present our non-falsifiability result in section II. Section III provides a discussion of the result, leading to our exploration of testing approaches in section IV. We conclude in section V.

I. Quantal Response Equilibrium

A. 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 QRE to other solution concepts. Consider a finite n -person normal form game Γ . The set of pure strategies (actions) 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 . 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 $u_i(s_{ij}, s_{-i})$ is reinterpreted as short-

hand for a probability measure in Δ_i placing 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 notion is based on the introduction of payoff perturbations associated with each action 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 . 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 action s_{ij} if and only if⁴

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

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

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

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 . 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 shocks 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 of the game Γ , a distribution of play by i 's opponents, and a joint distribution of i 's payoff perturbations, σ_i describes the probabilities

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 $\epsilon_{i'}$ are independent for $i' \neq i$; (2) the "baseline" payoff functions $u_i(s_i, s_{-i})$ and densities f_i are common knowledge; and (3) for each player i 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 game of incomplete information. Note that in this case, given the correctly anticipated equilibrium behavior of opponents, each player faces a standard polychotomous choice problem with additive random expected utilities. This observation is useful in estimation, since the distribution of equilibrium play by opponents will typically be directly observable to the researcher. It is also used in the proof of Theorem 1 below.

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 $j \in 1, \dots, J_i$, $\pi_{ij} = \sigma_{ij}(\bar{u}_i(\pi))$.

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 merely “usually close” to optimal rather than “always fully” 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. One may also view the perturbations as a device for “smoothing out” best response functions in the hope of obtaining more robust and/or plausible predictions (cf. Robert W. Rosenthal, 1989). 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. Alternatively, ϵ_{ij} might reflect unmodeled determinants of i 's utility from using strategy j . This interpretation is appealing in many applications since a fully specified theoretical 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 $u_i(s_{ij}, p_{-i})$ 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 on p_{-i} (as in the QRE) might be useful for predicting behavior or as an empirical model.⁶

⁵See also Hsuan-Chi Chen, Jed Friedman and Jacques-Francois Thisse (1997). Interpretations mirror those for random utility models in the discrete choice literature.

⁶Examples of empirical applications of QRE include Curtis S. Signorino (1999), Katja Seim (2002),

B. Application and Evaluation

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

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

With p observable, and all $\bar{u}_{ij}(\cdot)$ known, the unknown parameter λ is then easily estimated by maximum-likelihood.⁷

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 Although formal testing is uncommon, visual inspection often suggests a very good fit. Since a QRE would simply be a Nash equilibrium if perturbations were degenerate, the fit *must* improve when one adds the freedom to choose the best fitting member of a parametric family. In fact, however, the fit is often greatly improved. 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 this 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

Goeree, Holt and Palfrey (2002), Patrick Bajari and Hortaçsu (2005), Andrew Sweeting (2004), and Angelique Augereau, Shane Greenstein and Marc Rysman (forthcoming). See also Bajari (1998).

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 $\lambda \rightarrow \infty$.

⁸See, e.g., Michael R. Baye and John Morgan (2004), Timothy N. Cason and Stanley S. Reynolds (2005), Goeree, Holt and Palfrey (2002), Serena Guarnaschelli, McKelvey and Palfrey (2000), McKelvey and Palfrey (1995, 1998), McKelvey, Palfrey and Roberto A. Weber (2000), Mark Fey, McKelvey and Palfrey (1996).

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.

II. How Informative is Fit?

One might expect the QRE notion to impose considerable structure on the types of behavior consistent with equilibrium. Choice probabilities forming a QRE are a fixed point of the composite mapping $\sigma \circ \bar{u}$, and experience suggests that fixed points are special. However, the freedom to choose the joint densities f_i to fit the data gives considerable flexibility to QRE, particularly if one is unwilling to assume *a priori* that payoff perturbations are i.i.d..

To see this, consider relaxing the assumption of i.i.d. perturbations across each 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 imply independence of ϵ_{ij} across strategies j , without requiring that they be identically distributed. Joint densities f_i in \mathcal{S}_J allow dependence of ϵ_{ij} and ϵ_{ik} , $k \neq j$, but require ϵ_{ij} to be identically distributed for all j .

The following result shows that even when payoff perturbations are restricted to come from densities in one of these fairly restrictive classes, 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 a QRE to match the distribution of behavior of each player perfectly.⁹

⁹As the proof makes clear, the results apply to the "1-player" case of an additive random utility discrete choice model. For that paradigm, Steven Berry (1994) has shown that if utilities for each choice j are given by $\bar{u}_j + \epsilon_{ij}$ and an arbitrary joint distribution of the perturbations $\{\epsilon_{ij}\}_{j=1}^J$ is given, there exists a (unique) vector of mean utilities $(\bar{u}_1, \dots, \bar{u}_J)$ that will rationalize arbitrary probabilities on choices $\{1, \dots, J\}$. This contrasts with our result where, in the discrete choice case, arbitrary mean utilities are given and we choose a

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

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

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

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

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

For simplicity, for both part (i) and part (ii), we will consider here the case of a game in which every player has two pure strategies. An Appendix shows how to generalize these results to the case of an arbitrary number of strategies for each player.

(i) Take player 1 and let p_{1j} be the (given) probability that player 1 chooses strategy s_{1j} . Let $(\epsilon_{11}, \epsilon_{12})$ be independent draws from two-point distributions such that

$$\epsilon_{1j} = \begin{cases} \alpha_j & \text{with prob. } q_j \\ -\frac{q_j}{1-q_j} \alpha_j & \text{with prob. } 1 - q_j \end{cases}$$

distribution of mean-zero disturbances to match arbitrary data. John K. Dagsvik (1995), Daniel L. McFadden and Kenneth Train (2000), and Harry Joe (2001) consider related problems of choosing distributions from particular families to approximate choice probabilities (they also consider variation in the set of choices and/or choice characteristics). See also McFadden (1978). As in Berry (1994), all of these allow mean utilities to be chosen to fit the data. None of these results implies the others. Ours is more relevant to experimental settings, where mean payoffs are given.




Figure 1. Illustration for part (i).

for some $\alpha_j > 0$ and $q_j \in (0, 1)$ to be determined.¹⁰ By construction, each ε_{1j} has mean zero. Suppose $H_1^{12}(p_{-1}) > 0$ (the complementary case is analogous). Figure 1 illustrates. Realizations of $(\varepsilon_{11}, \varepsilon_{12})$ in the shaded region lead to strategy s_{11} being chosen over s_{12} . To conform with Figure 1, choose α_1 to satisfy $H_1^{12}(p_{-1}) < \alpha_1 < 2H_1^{12}(p_{-1})$ and let $\alpha_2 = \gamma[\alpha_1 - H_1^{12}(p_{-1})]$ for some $\gamma \in [0, 1)$. Let $q_2 = 1/2$. We can then match p_{11} exactly by setting $q_1 = p_{11}$. Repeating the argument for each player then proves the result.

Part (ii). Begin with player i and suppose $H_i^{12}(p_{-i}) > 0$ (the case $H_i^{12}(p_{-i}) \leq 0$ is analogous). Choose any $\delta_1 > \bar{u}_{i2}(p) - \bar{u}_{i1}(p)$. Let $\varepsilon_{i2} = \xi$ be uniformly distributed on $[-\kappa, \kappa]$, where $\kappa > \frac{\delta_1}{2}$ will be chosen below. Let

$$\varepsilon_{i1} = \begin{cases} \xi + \delta_1 & \xi + \delta_1 \leq \kappa \\ \xi + \delta_1 - 2\kappa & \xi + \delta_1 > \kappa. \end{cases}$$

¹⁰The two-point support is used only to provide a simple construction. Our prior working paper, Haile, Hortaçsu and Kosenok. (2003), showed that the mixtures of univariate normal densities (replacing the mixtures of Dirac-delta functions here) can be used to obtain the same result with continuously distributed perturbations.

or, letting \oplus represent addition on the circle $[-\kappa, \kappa]$ (see Figure 2),

$$\epsilon_{i1} = \xi \oplus \delta_1.$$

The marginal distributions of ϵ_{i1} and ϵ_{i2} are then both uniform on $[-\kappa, \kappa]$. In Figure 2, the bold arc of the circle shows 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, since $\delta_1 > \bar{u}_{i2}(p) - \bar{u}_{i1}(p)$, is also the probability that choice 1 is preferred to choice 2. Then because $\epsilon_{i1} > \epsilon_{i2}$ if and only if $\epsilon_{i2} \leq \kappa - \delta_1$, we have

$$\begin{aligned} p_{i1} &= \Pr(\epsilon_{i1} > \bar{u}_{i2} - \bar{u}_{i1} + \epsilon_{i2}) \\ &= \Pr(\epsilon_{i2} \leq \kappa - \delta_1) \\ &= 1 - \frac{\delta_1}{2\kappa}. \end{aligned}$$

Because we are free to choose any $\kappa > \frac{\delta_1}{2}$, any $p_{i1} \in (0, 1)$ can be matched. Repeating the argument for each player then proves the result. \square

III. Discussion

Theorem 1 shows that when the assumption of i.i.d. payoff perturbations is partially relaxed, any distribution of behavior by each player is consistent with a QRE. Hence, any falsifiable implication of QRE must be derived from additional maintained hypotheses on payoff perturbations. Even if one views the perturbations only as a device for smoothing out best response functions, one must be concerned about whether the way this is done is important. Theorem 1 shows that this choice completely determines what can and cannot be rationalized by QRE. This raises at least three important questions.

One question is how relevant our result is, given the literature's focus on i.i.d. perturbations from particular parametric families (typically logit). Those assumptions do imply testable restrictions. For example, McKelvey and Palfrey (1995, proof of Theorem 3) have shown that generally the set of probabilities that can form a logit QRE is a one-dimensional manifold, i.e., a set of curves, each of which implicitly defines all probabilities in terms of




Figure 2. Illustration for part (ii).

just one.¹¹ This can be a very strong restriction, although this is something worth checking in each application. For example, in a symmetric 2×2 game it may have limited bite, at least under the usual assumptions of symmetric equilibrium with identical distributions of perturbations for each player: in that case, the adding up constraint already forces the 2 choice probabilities to lie on a line, and the logit-QRE may not rule out much more. However, in games with more than 2 strategies per player, this can starkly limit the outcomes that a logit QRE can rationalize.¹² With only the i.i.d. assumption, the set of probabilities consistent with QRE can become considerably larger, but still place useful limits on the

¹¹For example, the behavior of a player with 3 available actions is characterized by 2 probabilities. In a logit QRE these probabilities must lie on a set of curves (often one curve) in $[0, 1]^2$. If one requires a single logit parameter to rationalize the behavior of both players (each with 3 pure strategies), equilibrium is characterized by a set of curves in $[0, 1]^4$.

¹²The same is true for asymmetric 2×2 games if one requires the same logit parameter to rationalize the behavior of each player. See, e.g., McKelvey, Palfrey and Weber (2000). They find that with this restriction the logit QRE fails to fit the data and thus propose allowing different logit parameters for each player.

behavior a QRE can explain.¹³

Because the restrictions implied by the i.i.d. (or i.i.d. logit) assumption vary with the game in question, it would be useful in practice to simulate the range of QRE outcomes possible given the particular game and distributional assumptions being considered. While informal, this could provide a sense of how to interpret the success or failure of the particular specification to rationalize the data. As an illustration, consider the following game studied by McKelvey and Palfrey (1995)

	B_1	B_2	B_3
A_1	(15, -15)	(0, 0)	(-2, 2)
A_2	(0, 0)	(-15, 15)	(-1, 1)
A_3	(1, -1)	(2, -2)	(0, 0)

with unique Nash equilibrium at (A_3, B_3) . A feature of this game is that action A_2 becomes unattractive for player 1 if there is a nontrivial chance that 2 plays B_2 . A symmetric argument applies to 2's action B_1 . When payoff perturbations are introduced to form a QRE, all actions are played with positive probability. Except when the variance of the payoff perturbations is so large that they swamp the baseline payoffs (so that a QRE puts nearly identical probabilities on all actions), this makes the actions A_2 and B_1 undesirable. As computations by McKelvey and Palfrey (1995, Table III and Figure 3) show, logit QRE probabilities for these actions are nearly zero for all values of the logit parameter above a certain threshold (given the symmetry of the game, they assume the same logit parameter for all players). Beyond this threshold, there is really only one probability to match in a symmetric equilibrium (since the third probability must add to 1) with the choice of the one logit parameter. McKelvey and Palfrey show that there is a continuum of quantal response equilibria of the following form (with symmetric properties for player 2): probability of nearly zero on A_2 and essentially any division of the remaining probability between A_1 and A_3 satisfying $\Pr(A_3) \geq \Pr(A_1)$ —a condition implied by a weaker assumption of exchangeable

¹³For example, the Monotonicity property described in the following section must hold. See also Goeree, Holt and Palfrey (2003), who describe an additional restriction of the (implicit) assumption of i.i.d. perturbations in an asymmetric matching pennies game. They reject the assumption of QRE with i.i.d. perturbations but are able to fit the data by introducing a risk aversion parameter.

perturbations (discussed below). Hence, if actual play puts probability close to zero on A_2 and B_1 (something we might expect in a QRE or under other notions of how games are played), there is no significant restriction of the i.i.d. logit assumption beyond what is implied by much weaker restrictions. Although this game is somewhat special, it points to the importance of paying attention to what types of outcomes can be rejected given the game and maintained assumptions one wishes to consider.

Of course, some restrictions implied by i.i.d (or i.i.d. logit) perturbations may be undesirable. In the discrete choice context, objections to these restrictions have been based both on *a priori* considerations and on the implications for comparative statics and counterfactual predictions. For example, one might expect larger payoffs to have perturbations with larger variances, or strategies that are similar to have similar perturbations. As is well known, the independence assumption has unnatural implications, including the IIA property of Robert D. Luce (1958) in the case of the logit (cf. Gerard Debreu, 1960). Considerable effort has been directed at developing tractable models of random utility discrete choice that relax the i.i.d. assumption. In the strategic context of QRE, motivations for relaxing the i.i.d. assumption are the same: these distributions will often be too restrictive to fit a rich data set, to fit observed comparative statics,¹⁴ or to lead to reasonable out-of-sample predictions. Testable restrictions that rely on weaker assumptions may enable more meaningful evaluation of the QRE hypothesis itself.

A second question is the relationship of our result on the *falsifiability* of the QRE to *identifiability* of empirical models based on the QRE. Falsifiability and identifiability are

¹⁴Although many papers have examined the fit of the logit QRE in different treatments (e.g., varying payoffs), typically a new value of the logit is estimated for each treatment. Capra, et al. (1999) demonstrate the ability of QRE with a single distribution of perturbations to rationalize observed comparative statics in the “traveller’s dilemma” game. Goeree, Holt and Palfrey (2002) formally test the assumption of a fixed distribution and reject. Camerer, Ho and Chong (2004) report that a fixed distribution does poorly in predicting outcomes across the different games they analyze. Like Timothy N. Cason and Stanley S. Reynolds (2005), they suggest that the distribution of perturbations required to rationalize the data varies with the scale of the payoffs. Other studies that re-estimate the distribution for each treatment (e.g., McKelvey and Palfrey (1995), Fey, McKelvey and Palfrey (1996)) report that the distribution that best explains behavior varies as players gain experience.

related but distinct. In general neither implies the other. One implication of Theorem 1 is immediate, however. If knowing the payoffs of the underlying game places no restriction on outcomes, observed outcomes cannot place any restriction on (much less identify) payoffs when they are unknown. Hence, Theorem 1 implies that observing the distribution of behavior in a single game could reveal nothing about latent expected payoffs, even when the perturbations are restricted to be draws from distributions in the set I_{J_i} or S_{J_i} . This is a more negative result than the failure of (point) identification: the data contain no information whatsoever about underlying payoffs.¹⁵ However, this implication is of little relevance in the experimental literature on QRE, since expected payoffs are easily calculated from the payoffs in the underlying game and the observed behavior of opponents.¹⁶

A third important question raised by Theorem 1 is whether there are other ways to evaluate the QRE hypothesis.¹⁷ We devote the following section to discussion of a few promising approaches.

¹⁵This result, like Theorem 1, extends immediately to the “one player” discrete choice environment. It should be pointed out, however, that the empirical literature on discrete choice has rarely considered estimation using data only on the characteristics and choice probabilities for a single choice set. The econometrics and empirical literatures on discrete choice generally rely on variation in covariates and/or the choice set itself, maintaining assumptions about how the distribution of perturbations varies with these changes. We explore similar ideas to develop testable restrictions of QRE below.

¹⁶An open question is whether there are useful conditions under which each f_i could be identified using data from experiments. See the discussion in footnote 19 below. Interest in identification of f_i may be more limited in experimental settings than in applications to field data. However, if f_i were known and believed to be invariant across a set of related games, this would provide a means of making point predictions that could be tested. We discuss tests based on this “invariance” assumption below without requiring identification of f_i .

¹⁷If one takes the incomplete information interpretation of QRE, which requires an assumption of independence of payoff perturbations across players (see footnote 4), this additional assumption could be tested by testing independence of players’ actions. This requires looking at players separately, contrary to the common practice of examining symmetric games and pooling observations over players to maximize the number of observations.

IV. Testing the QRE Hypothesis

The most promising testing approaches, in our view, come from observation of behavior in different games. This alone is not sufficient to deliver testable restrictions, since Theorem 1 implies that any behavior can be matched by appropriately selecting distributions of perturbations for each game. However, combining variation in the game with limits on how distribution functions can vary can enable several testing approaches, some of which build on well-known results from the discrete choice literature. The discussion here cannot be exhaustive; rather, our aim is to focus on a few approaches that may be particularly promising in practice. Two definitions will be useful for what follows:

DEFINITION 2 (Exchangeability): *The random variables $(\epsilon_{i1}, \dots, \epsilon_{iJ_i})$ are **exchangeable** if $f_i(\epsilon_{i1}, \dots, \epsilon_{iJ_i}) = f_i(\epsilon_{i\rho(1)}, \dots, \epsilon_{i\rho(J_i)})$ for every permutation operator ρ on the set $\{1, \dots, J_i\}$.*

DEFINITION 3 (Invariance): *The joint distribution of $(\epsilon_{i1}, \dots, \epsilon_{iJ_i})$ is **invariant** if $F_i(\epsilon_{i1}, \dots, \epsilon_{iJ_i} | u_i(\cdot)) = F_i(\epsilon_{i1}, \dots, \epsilon_{iJ_i})$ for all $(\epsilon_{i1}, \dots, \epsilon_{iJ_i})$ and all payoff functions $u_i(\cdot)$.*

Exchangeability (a.k.a. “interchangeability”) is a strong form of symmetry, requiring not only identical marginal distributions for each ϵ_{ij} , but also identical covariances, conditional moments, etc. In the more familiar and closely related discrete choice literature, exchangeability holds for the conditional logit model and also for the multinomial probit model under the restrictions $E[\epsilon_{ij}^2] = \sigma^2 \forall i, j$ and $E[\epsilon_{ij}\epsilon_{ij'}] = \rho \forall j, j' \neq j$. Nested logit and mixed logit (or probit) models, on the other hand, violate exchangeability by design. *Invariance* is a property requiring a similar lack of sensitivity to variations in payoffs, but allows the possibility of asymmetry—for example, the possibility that perturbations are larger for some actions than others. While strong, this assumption has often been used in the econometrics literature on discrete choice models. Invariance is typically maintained in applications of the conditional logit, nested logit, and multinomial probit, for example, but relaxed in mixed logit/probit (random coefficients) models.

A. Approach 1: “Regular” QRE

Responding to an earlier draft of this paper, Goeree, Holt and Palfrey (2005) have proposed a refinement of QRE, “regular” QRE. Suppose $p \in \Delta$ characterizes behavior in a game. As

before, let $\sigma_{ij}(\bar{u}_i(p))$ represent the element of p corresponding to $\Pr(i \text{ plays } j)$. Rather than restricting payoff disturbances explicitly, they define a regular QRE by restricting quantal response functions to satisfy the following axioms.¹⁸

1. **Interiority.** $\sigma_{ij}(\bar{u}_i(p)) > 0$ for all $i, j = 1, \dots, J_i$
2. **Continuity:** $\sigma_{ij}(\bar{u}_i(p))$ is a continuous and differentiable function of $\bar{u}_i(p)$ for all $\bar{u}_i(p) \in R^{J_i}$
3. **Responsiveness:** $\frac{\partial \sigma_{ij}(\bar{u}_i(p))}{\partial \bar{u}_{ij}(p)} > 0$ for all $j = 1, \dots, J_i$
4. **Monotonicity:** $\bar{u}_{ij}(p) > \bar{u}_{ik}(p) \Rightarrow \sigma_{ij}(\bar{u}_i(p)) > \sigma_{ik}(\bar{u}_i(p))$ for all $j = 1, \dots, J_i$

Goeree et al. argue that these axioms are economically and intuitively compelling. Interiority and Continuity are natural technical properties. Responsiveness and Monotonicity are stronger properties with significant economic content. Responsiveness restricts the ways that a player’s behavior can change in response to a *ceteris paribus* change in the equilibrium expected payoff from one action. Monotonicity restricts probabilistic behavior *within* a game, requiring that actions with higher expected payoffs to be played more often.

Clearly these axioms imply testable restrictions. Monotonicity can be checked directly in any game. Responsiveness concerns changes in behavior across games and is more subtle. Raising i ’s payoffs $u_i(s_{ij}, s_{-i})$ from action j in the baseline game will make j more attractive, *ceteris paribus*, but may ultimately lead to a change in the play of i ’s opponents. This will typically change $\bar{u}_{ik}(p)$ for each k , and could even cause $\bar{u}_{ij}(p)$ to be lower than in the original

¹⁸Goeree et al. describe this as a “reduced form” definition of equilibrium and contrast it with the original “structural” definition based on explicit modelling of payoff perturbations. They suggest that the structural QRE implies restrictions on behavior across games (Proposition 2), and that there are outcomes consistent with the (reduced form) regular QRE that cannot be rationalized by a structural QRE (Proposition 6). However, an assumption of Invariance is implicit in their analysis of the structural QRE, while not imposed on the regular QRE. It is immediate from our Theorem 1 that without some restriction on how the joint densities f_i change when a game changes, none of the restrictions described in their Proposition 2 hold, negating the conclusion of their Proposition 6 as well. Note that the use of the terms “structural” and “reduced form” here describe the way equilibrium is defined, not econometric modelling approaches.

equilibrium.¹⁹ However, Proposition 4 of Goeree et al. demonstrates that Responsiveness is nonetheless sufficient in some games to obtain testable restrictions based on changes in the payoffs $u_i(s_i, s_{-i})$ of the baseline game. Although they discuss only one example, testable comparative statics predictions can be derived in other environments in which a change in payoffs that, *ceteris paribus*, makes action j more attractive to player i induces play by i 's opponents that also make action j (weakly) more attractive. We discuss some simple classes of such games below.

A natural question is what assumptions on the underlying model imply the axioms, particularly the Monotonicity and Responsiveness conditions. Goeree et al. show that a sufficient condition for Monotonicity is Exchangeability. They point out that Exchangeability and the implied Monotonicity restriction can fail under additive error structures that seem reasonable *a priori*.²⁰ They suggest that in practice such violations are rarely observed. A sufficient condition for Responsiveness is Invariance.²¹

B. Approach 2: Rank-Cumulative Probabilities

Start with any game Γ with payoffs given by the functions $u_i \forall i$. Now modify the game by changing the payoffs, leaving the strategy space for each player fixed. For each player i let $u'_i : \Delta \rightarrow \mathbb{R}$ give the new payoffs. We will say that these are two games that *differ*

¹⁹This feature of quantal choice in a strategic setting also suggests challenges for obtaining identification results that build on those for additive random utility models of discrete choice (e.g., Charles F. Manski (1988), Rosa L. Matzkin (1992)) since these rely the ability to “trace out” the distribution of perturbations through sufficient variation in mean payoffs (utilities) while the distribution of perturbations is held fixed. In some games it may be difficult to generate sufficient variation in expected payoffs through manipulation of the payoffs of the underlying normal form game. Even in an experimental setting, where the game can be modified freely and expected payoffs can be treated as known, identification of f_i may be a challenge.

²⁰In the multinomial choice setting, nested logit, mixed logit/probit models lead to violations of Monotonicity.

²¹We show this below. Goeree, Holt and Palfrey (2005, Propositions 5) argue that, under the usual “admissibility” conditions for the QRE, Exchangeability is sufficient for all their axioms. Their analysis maintains an implicit assumption of Invariance, and their claim is correct under this assumption.

only in payoffs. Having observed repeated play of these two games, let p and p' characterize the observed behavior in each. Define the mean expected payoffs $\mu_i = \frac{1}{J_i} \sum_{j=1}^{J_i} u_i(s_{ij}, p_{-i})$ and $\mu'_i = \frac{1}{J_i} \sum_{j=1}^{J_i} u'_i(s_{ij}, p'_{-i})$. Then for all i and j let $\tilde{u}_{ij} = u_i(s_{ij}, p_{-i}) - \mu_i$ and $\tilde{u}'_{ij} = u'_i(s_{ij}, p'_{-i}) - \mu'_i$, normalizing each player's expected payoff from each action by his mean. Let

$$\begin{aligned} \dot{u}_{ij} &= \tilde{u}'_{ij} - \tilde{u}_{ij} \\ d_{i(jk)} &= \tilde{u}_{ij} - \tilde{u}_{ik} \\ d'_{i(jk)} &= \tilde{u}'_{ij} - \tilde{u}'_{ik}. \end{aligned}$$

Without loss of generality, re-index i 's actions so that $\dot{u}_{i1} \geq \dot{u}_{i2} \geq \dot{u}_{i3} \geq \dots$, where some inequality must be strict except in the trivial case with $\dot{u}_{ij} = 0 \forall j$, i.e., when the games do not differ except possibly in the scaling of all payoffs.²² By our choice of indexing, $d'_{i(jk)} \geq d_{i(jk)} \forall j, k > j$. Furthermore, at least one of these inequalities strict except in the trivial case. Define the *rank-cumulative probabilities*,

$$\rho_{ik} = \sum_{l=1}^k p_{il}, \quad \rho'_{ik} = \sum_{l=1}^k p'_{il}.$$

For example, in the original game, ρ_{ik} gives the probability that i uses a strategy in the set $\{s_{i1}, \dots, s_{ik}\}$.

THEOREM 2: *Consider two games that differ only in payoffs. Under the Invariance assumption, behavior consistent with QRE must produce increasing rank-cumulative probabilities, i.e., $\rho'_{ik} \geq \rho_{ik} \forall k = 1, \dots, J - 1$.*

Proof. Under the Invariance assumption, one may treat the normalized expected payoffs \tilde{u}_{ij} and \tilde{u}'_{ij} as the true expected payoffs without loss of generality. Further, we can write probabilities over realizations of ϵ_i without conditioning on which game is being played. We then have

$$\begin{aligned} \rho_{ik} &= \Pr \left(\max_{l \in \{1, \dots, k\}} \{ \tilde{u}_{il} + \epsilon_{il} \} \geq \tilde{u}_{ij} + \epsilon_{ij} \quad \forall j > k \right) \\ &= \Pr \left(\max_{l \in \{1, \dots, k\}} \{ d_{i(lj)} + \epsilon_{il} \} \geq \epsilon_{ij} \quad \forall j > k \right). \end{aligned}$$

²²For this trivial case, under the Invariance assumption QRE requires that the distribution of behavior be identical in the two games. Goeree, Holt and Palfrey (2005) refer to this property as “scale invariance.”

Similarly,

$$\rho'_{ik} = \Pr \left(\max_{l \in \{1, \dots, k\}} \{d'_{i(lj)} + \epsilon_{il}\} \geq \epsilon_{ij} \quad \forall j > k \right).$$

Since $d'_{i(lj)} \geq d_{i(lj)}$ for all $j > l$, the result follows. \square

A special case of increasing rank-cumulative probabilities arises when $u'_i(s_{ij}, p'_{-i}) > u_i(s_{ij}, p_{-i})$ for exactly one $j \in J_i$, with $u'_i(s_{ik}, p'_{-i}) = u_i(s_{ik}, p_{-i})$ for all $k \neq j$. Theorem 2 then requires the probability of play of this action to increase, giving Goeree, Holt and Palfrey's (2005) Responsiveness property. While Invariance is sufficient for Responsiveness, Theorem 2 shows that Invariance also implies a richer set of testable restrictions. These restrictions may also be more useful in practice, since they are not limited to cases in which only one of player 1's actions experiences a change in expected payoff. Note that we need not predict how the payoffs $\bar{u}_i(p)$ will change when the baseline payoffs of the game are changed: this is determined by direct observation of the new equilibrium probabilities.

The Invariance assumption may be strong, especially if one believes that the significance (e.g., variance) of the disturbance term is dependent on the "stakes" faced by the agent. It is not obvious whether this should be the case, and indeed one interpretation of the QRE is that it provides a way for "mistakes" to be made more often when they are not very costly (when mean payoff differences are small)—a property that could be undone with perturbations scaled by the associated payoff differences.

One can also obtain a testable restriction with the weaker and arguably more plausible assumption that each f_i is invariant to variations in the *baseline* payoffs u_{-i} of *other* players. For example, consider a two-player game, where one examines behavior under variations in player 2's payoff matrix under the assumption that f_1 does not change in response. After recalculating \bar{u}_1 based on player 2's new behavior, it is immediate that p_{1j} must be monotonic in $(\bar{u}_{1j} - \bar{u}_{11})$ across j .

C. Approach 3: Block-Marschak Polynomials

The final approach we discuss is based on an extension of results obtained by H. D. Block and Jacob Marschak (1960) and Jean-Claude Falmagne (1978) for random utility models of

discrete choice.²³ We first review this result and then describe an extension to strategic contexts.

Begin with a decision problem in which a utility-maximizing agent must choose one alternative from a finite choice set A . The utility from alternative $j \in A$ is given by the sum of a mean utility \bar{u}_j and a zero-mean random component ε_j . The mean utilities may or may not be known to the researcher, and no restriction is made on the joint distribution of the perturbations $\varepsilon_j, j \in A$. Let $p(A) = \{p_1(A), \dots, p_{|A|}(A)\}$ denote the resulting choice probabilities. Now consider a series of related choice experiments in which the agent chooses from restricted choice sets $B \subset A$. In these experiments, the mean utilities $\bar{u}_j, j \in B$, are identical to those in the original choice problem. We refer to this as a *sequence of choice experiments based on a master choice set A* . For each choice set B , the joint distribution of the latent utility perturbations (i.e., $\varepsilon_j, j \in B$) is also assumed to be identical to the marginal (with respect to the relevant set of choices) distribution of these perturbations in the original problem. We will refer to this as an assumption of a *fixed stochastic structure*. Under this assumption, the underlying random utility structure is held fixed across the sequence of experiments, but the choice sets are restricted to subsets of the master set A .

Let $p(B) = \{p_1(B), \dots, p_{|A|}(B)\}$ denote the corresponding choice probabilities for each choice set B , where of course $p_j(B) = 0$ whenever $j \notin B$. Let $F(B, m)$ denote the set of all subsets of B containing exactly m elements. Then for all $B \subset A$ with $B \neq A$ and any $j \in A - B$,²⁴ define the *Block-Marschak polynomial*

$$K_{j,B} = \sum_{k=0}^{|B|} (-1)^k \sum_{C \in F(B, |B|-k)} p_j(A - C).$$

LEMMA 1 (Falmagne (1978)): *Consider a sequence of choice experiments based on a master choice set A . Under the assumption of a fixed stochastic structure, choice probabilities $p(B), B \subset A$, are consistent with utility maximization in a random utility model if and only if all Block-Marschak polynomials are nonnegative.*

²³See also Salvador Barbera and Prasanta Pattanaik (1986) and McFadden (2005). We thank a referee for making us aware of Falmagne (1978).

²⁴The notation $A - B$ denotes the difference between sets A and B .




Figure 3. Each set E_B represents the region of realizations of perturbations making alternative j the utility-maximizing choice when B is the choice set, so that $p_j(B) = \Pr(E_B)$. Note that both $E_{A-\{k\}} \subset E_{A-\{k,\ell\}}$ and $E_{A-\{\ell\}} \subset E_{A-\{k,\ell\}}$ must hold. $K_{j,\{k,\ell\}}$ gives the probability of the unshaded area.

While the compact notation for the Block-Marschak polynomials is opaque, the underlying idea is simple: the probability of choosing a given alternative j is higher when the set of other available choices is smaller. For $B = \emptyset$, $K_{j,B} = p_j(A)$, which obviously cannot be negative. If $|B| = 1$,

$$K_{j,B} = p_j(A - B) - p_j(A)$$

and the requirement is that dropping B from the choice set weakly raises the choice probability on each remaining alternative. When B contains 2 elements, k and ℓ ,

$$K_{j,B} = p_j(A - \{k, \ell\}) - [p_j(A - \{k\}) + p_j(A - \{\ell\})] + p_j(A)$$

which again must be positive (see Figure 3). For larger $|B|$, the polynomials have a higher order but a similar interpretation: the probability of a given choice goes up whenever other alternatives are removed from the choice set.

To extend Falmagne's result to strategic settings, we will consider variation in the set of pure strategies available to a player in a game that is otherwise held fixed. This extension would be immediate were it not for the fact that changing player i 's strategy space may lead to changes in other players' equilibrium behavior. Since $\bar{u}_{ij}(p)$ depends on p_{-i} , this will typically lead to a violation of the condition that the expected payoff from each given action (i.e., that prior to the realization of ϵ_i) be the same across all games in which that action is available.

This can be overcome in some types of games, and the laboratory provides an ideal environment for considering them. A trivial approach is to examine the behavior of a player whose payoffs from each action are independent of the actions of his opponents. In a QRE, such a player acts as if he faces a random utility discrete choice problem, and the extension of Lemma 1 is immediate. This is not very satisfying as a test of QRE, since the essence of QRE is strategic behavior. However, games of this sort can provide more meaningful restrictions if we consider an *opponent* of players who face non-strategic decisions.

For example, consider a game Γ_0 that is *non-strategic for player 1*, i.e., letting S_{i0} denote i 's strategy space

$$u_1(s_{1j}, s_{-i}) = u_1(s_{1j}) \text{ for all } s_{1j} \in S_{10}, s_{-i} \in S_{-i0}.$$

Suppose for simplicity that Γ_0 is a 2-player game. We focus on player 2, whose payoffs *do* depend on the actions of player 1. Consider a sequence of games $\Gamma_1, \dots, \Gamma_G$ which differ from Γ_0 only in that player 2's strategy spaces S_{21}, \dots, S_{2G} are subsets of the original S_{20} . For each player i , the payoff function $u_i(s_i, s_{-i})$ is the same in each game.²⁵ No restriction is placed on the joint distribution of perturbations in the master game. However, analogous to the discrete choice case, we assume a *fixed stochastic structure*, i.e., in each game Γ_g the payoff perturbations $\epsilon_{ij}, j \in S_{ig}$ have joint distribution equal to the marginal distribution of these same perturbations in Γ_0 . Because all of these games remain nonstrategic for player 1, player 1's QRE behavior is the same in every game. Thus, player 2's mean payoff \bar{u}_{2j} from any available pure strategy j is the same in each game. In a QRE player 2 anticipates the equilibrium distribution of behavior by his opponent, and his behavior is then as if he were maximizing his utility in a sequence of choice experiments based on a fixed stochastic structure, with the expected payoffs $\bar{u}_{2j}(p)$ replacing the mean utilities \bar{u}_j . Lemma 1 then can be applied, giving the following result.

THEOREM 3: *Consider a game Γ_0 that is nonstrategic for all players except i and sequence of games $\Gamma_1, \dots, \Gamma_G$ that differ from Γ_0 only in that player i 's strategy spaces S_{i1}, \dots, S_{iG} are subsets of the original S_{i0} . Let $p_{ij}(S_{ig})$ denote the probability that i plays pure strategy j in game g . Under the assumption of a fixed stochastic structure, behavior is consistent with QRE if and only if, for all $B \subset S_{i0}$ with $B \neq S_{i0}$ and all $j \in S_{i0} - B$, $K_{j,B} = \sum_{k=0}^{|B|} (-1)^k \sum_{C \in F(B, |B|-k)} p_{ij}(S_{i0} - C) \geq 0$.*

As an example, consider an n -player game Γ_0 that is nonstrategic for all players except player 1, who has pure strategies $S_{10} = \{1, 2, 3\}$. Suppose we observe probabilities $\{p_{11}(\Gamma_0), p_{12}(\Gamma_0), p_{13}(\Gamma_0)\}$ in the game Γ_0 , and $\{p_{11}(\Gamma_1), p_{12}(\Gamma_1)\}$ in game Γ_1 where player 1's strategy set is reduced to $S_{11} = \{2, 3\}$. Then Theorem 3 gives the testable restrictions

$$(5) \quad p_{1j}(\Gamma_1) - p_{1j}(\Gamma_0) \geq 0 \quad j \in \{2, 3\}$$

²⁵The domain of $u_2(\cdot)$ does change across games, but its value at any given profile of feasible actions does not.

i.e., player 1 must be more likely to choose a given action when the set of alternative strategies available is smaller. Suppose that, in addition, we observe play probabilities $\{p_{11}(\Gamma_2), p_{13}(\Gamma_2)\}$ from Γ_2 , where player 1's strategy set is reduced to $\{1, 3\}$. Along with the two inequalities analogous to (5), this gives us the additional testable restriction

$$1 + p_{13}(\Gamma_0) \geq p_{13}(\Gamma_1) + p_{13}(\Gamma_2).$$

Note that observing player 1's behavior in three games gives 5 testable restrictions. Observing behavior in the game Γ_3 with $S_{13} = \{1, 2\}$ would give 4 more

$$\begin{aligned} p_{1j}(\Gamma_3) - p_{1j}(\Gamma_0) &\geq 0 & j \in \{1, 2\} \\ 1 + p_{11}(\Gamma_0) &\geq p_{11}(\Gamma_2) + p_{11}(\Gamma_3) \\ 1 + p_{12}(\Gamma_0) &\geq p_{12}(\Gamma_1) + p_{12}(\Gamma_3) \end{aligned}$$

yielding 9 restrictions from 4 treatments. Considering games with larger numbers of strategies, and running experiments with larger numbers of subsets of the original strategy set allow use of higher-order Block-Marschak polynomials, leading to a large number of testable restrictions.

The limitation to games that are strategic for only one player can be overcome if one considers some simple sequential games.²⁶ For example, consider the class of finite “Stackelberg games,” i.e., 2-period 2-player sequential move games. Let $\Gamma_0, \Gamma_1, \dots$ be a sequence of Stackelberg games with the first mover labeled player 1. Let the strategy space for player 2 be identical in all games, while the strategy space for player 1 in each game is a subset of that in the master game Γ_0 . Let the payoff functions $u_i(s_i, s_{-i})$ be the same across all games. As above, we refer to this as a *sequence of Stackelberg games based on a master game* Γ_0 . Under the assumption of a fixed stochastic structure, in each game Γ_g the payoff

²⁶McKelvey and Palfrey (1998) discuss the extension of QRE to general extensive form games, using the agent-normal form representation of the game. For the Stackelberg games we consider here, the distinction between the agent-normal form and standard normal form specification is not important, since each player “moves” only once. Hence for these games the agent-QRE notion developed by McKelvey and Palfrey is an obvious extension of the QRE for normal form games discussed above.

perturbations ϵ_{ij} , $j \in S_{ig}$ have joint distribution equal to the marginal distribution of these same perturbations in the original game Γ_0 .

It is easy to see that the sequential nature of these games leaves the play of the Stackelberg follower constant across all of these games *conditional on the action taken by player 1*. This implies that player 1’s expected payoff from taking an action s_{ij} is the same in any game g for which this action is available. Lemma 1 is then again immediately applicable, giving the following result.

THEOREM 4: *Consider a sequence of Stackelberg games based on a master game Γ_0 . Let S_{10} denote player 1’s strategy space in the original game Γ_0 and let $p_{1j}(S_{1g})$ denote the probability that 1 plays pure strategy j in the game Γ_g in which $S_{1g} \subset S_{10}$ is his strategy space. Under the assumption of a fixed stochastic structure, behavior is consistent with QRE if and only if, for all $B \subset S_{10}$ with $B \neq S_{10}$ and all $j \in S_{10} - B$, $K_{j,B} = \sum_{k=0}^{|B|} (-1)^k \sum_{C \in F(B, |B|-k)} p_{1j}(S_{10} - C) \geq 0$.*

The maintained hypothesis for these results (the “fixed stochastic structure” assumption) is that the joint distribution of perturbations $\{\epsilon_{ij}, j \in B\}$ is the same in any game in which i ’s strategy space includes B . This assumption is implied by Invariance, but is clearly weaker.²⁷ Because the Block-Marschak polynomials provide a rich set of restrictions with quite weak maintained hypotheses, we view this as a particularly promising approach for testing.²⁸ The obvious limitation is in the set of games that can be considered. This is a less serious limitation for experimental work, however, and further work may reveal additional types of games to which this approach can be applied.

²⁷Setting $\bar{u}'_{ij}(p) = -\infty$ is equivalent to removing s_{ij} from i ’s choice set. The fixed stochastic structure assumption places no restriction on perturbations across games that differ only in payoffs, and implies neither Responsiveness nor the “scale invariance” property discussed by Goeree et al. (2005). Unlike Invariance, the fixed stochastic structure assumption is consistent with mixed logit/probit models of multinomial choice.

²⁸For the discrete choice setting, Joe (2000) has described additional intuitive implications that can be tested if one assumes that perturbations are mutually independent.

V. Conclusion

QRE is an appealing equilibrium notion with several compelling interpretations. It is natural to ask how well it does in explaining/predicting behavior in practice. However, we have pointed out that evaluating the fit of a QRE in a single game is uninformative without significant *a priori* restrictions on the distributions of payoff perturbations. This should not be mistaken for a critique of the QRE notion itself.²⁹ Rather, our aim has been to clarify some limitations of examining behavior one game at a time and to develop approaches for more informative evaluation of QRE.

In general this will require maintaining assumptions beyond those that define a QRE, and the most appropriate set of maintained hypotheses may depend on the application. The standard logit specification imposes such restrictions and can starkly limit the set of possible QRE outcomes in many games. Presenting this set of possible outcomes for the games examined in practice will enable clearer evaluation of the evidence.

In some cases, the logit (or i.i.d.) assumption may not sufficiently restrict the set of outcomes to provide a powerful test, or may be viewed as restricting outcomes in undesirable ways. The latter is especially likely when examining out-of-sample predictions or behavior across games. In such cases alternative testing approaches would be useful. We have pointed readers to the axiomatic approach proposed by Goeree, Holt and Palfrey (2005) and suggested two new approaches. Each approach examines comparative statics predictions, relying explicitly or implicitly on maintained assumptions about how the distributions of payoff perturbations are related across different games. Each approach can produce a large number of testable restrictions from a relatively small number of different experimental treatments.³⁰ Currently, evidence regarding such comparative statics predictions of QRE is limited, and we hope that attention to this topic here and in Goeree, Holt and Palfrey (2005) leads researchers to explore richer sets of testable restrictions in order to better evaluate QRE

²⁹See John O. Ledyard (1986) for discussion in a similar spirit of the empirical restrictions imposed by Bayesian Nash equilibrium.

³⁰We have focused on the theoretical question of falsifiability, leaving open interesting questions of formal testing procedures.

as a tool for understanding and predicting strategic behavior.

*Haile: Department of Economics and Cowles Foundation, Yale University, and NBER, 37 Hillhouse Avenue, Box 208264, New Haven, CT 06520, philip.haile@yale.edu; Hortaçsu: Department of Economics, University of Chicago, and NBER, 1126 E. 59th St., Chicago, IL 60637, hortacsu@uchicago.edu; Kosenok: New Economic School, Nakhimovsky Prospekt 47, Suite 1721 (11), 117418 Moscow, Russia, gkosenok@nes.ru. This paper subsumes material in two earlier drafts: one, first circulated widely in August 2003, with the same title, and a second with the title “On the Empirical Content of Quantal Choice Models.” We received helpful comments from a number of colleagues and seminar audiences. We are also grateful to the referees for very helpful comments and for encouraging us to develop our ideas for testing more fully. Kun Huang and Dmitry Shapiro provided capable research assistance. Financial support from the National Science Foundation (grants SES-0112047 and SES-0449625) and the Alfred P. Sloan Foundation is gratefully acknowledged.

Appendix

Proof of Theorem 1 with $J_i > 2$ pure strategies:

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

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

for some $\alpha_j > 0$ to be determined below, where $q_{J_i} \in (0, 1)$ is arbitrary and

$$(6) \quad q_j = \frac{p_{ij}}{1 - \sum_{k < j} p_{ik}} \quad \forall j < J_i.$$

Note that $q_j \in (0, 1) \forall j$ because $p_{ij} \in (0, 1) \forall j$ and $\sum_{j=1}^{J_i} p_{ij} = 1$. By construction, each ϵ_{ij} has expectation zero. Let $A_{jk} = \Pr \left\{ \epsilon_{ij} \geq \epsilon_{ik} + H_i^{jk}(p_{-i}) \right\}$. Now begin by fixing $\alpha_{J_i} > 0$ at an arbitrary value. For sufficiently large $\alpha_{J_{i-1}}$ we have

$$\begin{aligned} \alpha_{J_{i-1}} &> \alpha_{J_i} + H_i^{(J_i-1)J_i}(p_{-i}) \\ \frac{-q_{J_{i-1}}}{1 - q_{J_{i-1}}}\alpha_{J_{i-1}} &< \frac{-q_{J_i}}{1 - q_{J_i}}\alpha_{J_i} + H_i^{(J_i-1)J_i}(p_{-i}) \end{aligned}$$

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, $\alpha_{J_{i-1}}^*$. Of course 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 sufficiently large $\alpha_{J_{i-2}}$

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

Proceeding in this fashion, given any $q_j \forall j$, we can choose each α_j so that

$$(7) \quad A_{jk} = \begin{cases} q_j & \text{if } j < k \\ 1 - q_k & \text{if } j > k. \end{cases}$$

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 $j = J_i$. By (6), $q_j \prod_{k < j} (1 - q_k) = p_j$ for $j < J_i$, while $\prod_{k < J_i} (1 - q_k) = 1 - \sum_{k < J_i} p_{ij} = p_{iJ_i}$. 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.³¹ 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 (see Haile, Hortaçsu and Kosenok, 2003).³²

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

$$(8) \quad \varepsilon_{ij} = \begin{cases} \xi + \delta_j & \xi + \delta_j < \kappa \\ \xi + \delta_j - 2\kappa & \xi + \delta_j > \kappa \end{cases}$$

where each δ_j will be 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 $\delta_1 > \delta_2 > \dots > \delta_{J_i}$. Suppose for the moment that $H_i^{jk} = 0$ for all j and k . One can then confirm that for each j

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

³¹There 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).

³²It is intuitive that mixtures of normals could approximate the two-point distributions above arbitrarily well. In Haile, Hortaçsu and Kosenok (2003), however, we show that one can match the probabilities p_{ij} exactly.

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

$$(9) \quad \delta_j = \left(1 - \sum_{\ell=1}^j p_{i\ell}\right) 2\kappa \quad j = 1, \dots, J_i - 1.$$

We now drop the assumption that each $H_i^{jk} = 0$. We do this by ensuring the equivalence

$$(10) \quad \{\epsilon_{ij} > \epsilon_{ik} + H_i^{jk}\} \iff \{\epsilon_{ij} > \epsilon_{ik}\}$$

which holds if ϵ_{ij} and ϵ_{ik} are always sufficiently different when $j \neq k$. When (9) holds we know $|\epsilon_{ij} - \epsilon_{ik}| \geq 2\kappa (\min_{j=1, \dots, J_i} p_{ij})$ for all $j \neq k$.³³ So for any $\kappa > \frac{\max_{k,j} |H_i^{jk}|}{2 \min_{j=1, \dots, J_i} p_{ij}}$, (10) holds and we have

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

Repeating this construction for every player completes the proof. □

³³When $\xi + \delta_j$ and $\xi + \delta_k$ both exceed κ or are both smaller than κ , this is immediate from (9). When $\xi + \delta_j > \kappa > \xi + \delta_k$, $|\epsilon_j - \epsilon_k| = |\delta_k - \delta_j + 2\kappa|$, and the claim then follows from (9).

References

- Akerberg, Daniel, and Rysman, Marc. 2005. "Unobservable Product Differentiation in Discrete Choice Models: Estimating Price Elasticities and Welfare Effects." *RAND Journal of Economics*, 36: 771–788.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt. 1998. "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction." *Journal of Political Economy*, 106(4): 828-853.
- Anderson, Simon P., Jakob K. Goeree, and Charles A. Holt. 2001. "Minimum Effort Coordination Games: Stochastic Potential and the logit Equilibrium." *Games and Economic Behavior*, 34(2): 177-199.
- Anderson, Simon P., Jakob K. Goeree, and Charles A. Holt. 2002. "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies." *Southern Economic Journal*, 69(1): 21-47.
- Augereau, Angelique, Shane Greenstein and Marc Rysman. Forthcoming. "Coordination vs. Differentiation in a Standards War: 56K Modems." *RAND Journal of Economics*.
- Bajari, Patrick. 1998. "Econometrics of Sealed-bid Auctions." *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*.
- Bajari, Patrick, and Ali Hortacısu. 2005. "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data." *Journal of Political Economy*. 113: 703-741.
- Barbera, Salvador, and Prasanta K. Pattanaik. 1986. "Falmagne and the Rationalizability of Stochastic Choices in Terms of Random Orderings." *Econometrica*, 54 (3): 707-715.
- Baye, Michael R. and John Morgan. 2004. "Price Dispersion in the Lab and on the Internet." *RAND Journal of Economics*, 35(3): 448–466.

- Berry, Steven. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, 25(2): 242-262.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841-890.
- Block, H.D., and Jacob Marschak. 1960. "Random Orderings and Stochastic Theories of Responses." In *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, ed. Olkin et al. Stanford University Press.
- Camerer, Colin F., Teck-Hua Ho, and Juin Kuan Chong. 2004. "Behavioral Game Theory: Thinking, Learning, and Teaching." In *Advances in Understanding Strategic Behavior, Essays in Honor of Werner Guth*, ed. Steffan Huck. London: Palgrave MacMillan.
- Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt. 1999. "Anomalous Behavior in a Travelers Dilemma?" *American Economic Review*, 89(3): 678-690.
- Cason, Timothy N., and Stanley S. Reynolds. 2005. "Bounded Rationality in Laboratory Bargaining with Asymmetric Information." *Economic Theory*, 25: 553-574.
- Chen, Hsuan-Chi, Jed Friedman, and Jacques-Francois Thisse, J. 1997. "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach." *Games and Economic Behavior*, 18(1): 32-54.
- Dagsvik, John K. 1995. "How Large is the Class of Generalized Extreme Value Random Utility Models." *Journal of Mathematical Psychology*, 39(1): 90-98.
- Debreu, Gerard. 1960. Review of R.D. Luce, "Individual Choice Behavior: A Theoretical Analysis." *American Economic Review*, 50(1): 186-188.
- Dufwenberg, Martin, Uri Gneezy, Jacob K. Goeree, and Rosemarie Nagel. 2002. "Price Floors and Competition." Working Paper.
- Falmagne, Jean-Claude. 1978. "A Representation Theorem for Finite Random Scale Systems." *Journal of Mathematical Psychology*, (18): 52-72.

- Fey, Mark, Richard D. McKelvey, and Thomas R. Palfrey. 1996. "An Experimental Study of Constant-Sum Centipede Games." *International Journal of Game Theory*, 25(3): 269–287.
- Goeree, Jacob K., and Charles A. Holt. 1999. "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory." *Proceedings of the National Academy of Sciences*, 96 (September) 10564-10567.
- Goeree, Jacob K., and Charles A. Holt. 2000. "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games." *European Economic Review*, 44: 1079-1089.
- Goeree, Jacob K., and Charles A. Holt. 2001. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." *American Economic Review*, 91(5): 1402-1422.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. 2002. "Quantal Response Equilibrium and Overbidding in Private-Value Auction." *Journal of Economic Theory*, 104(1): 247-272.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. 2003. "Risk Averse Behavior in Generalized Matching Pennies Games." *Games and Economic Behavior*, (45): 97–113.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. 2005. "Regular Quantal Response Equilibrium." *Experimental Economics*, 8(December) 347–367.
- Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey. 2000. "An Experimental Study of Jury Decision Rules." *American Political Science Review*, 94(2): 407-423.
- Haile, Philip, Ali Hortaçsu, and Grigory Kosenok. 2003. "On the Empirical Content of Quantal Response Equilibrium." Cowles Foundation Discussion Paper No. 1432.
- Hausman, Jerry A., and David A. Wise. 1978. "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences." *Econometrica*, 46(2): 403–426.

- Ho, Teck-Hua, Colin F. Camerer and Juin-Kuan Chong. 2007. "Self-Tuning Experience Weighted Attraction Learning in Games." *Journal of Economic Theory*, 133: 177–198.
- Joe, Harry. 2000. "Inequalities for Random Utility Models, with Applications to Ranking and Subset Choice Data." *Methodology and Computing in Applied Probability*, 2: 359–372.
- Joe, Harry. 2001. "Multivariate Extreme Value Distributions and Coverage of Ranking Probabilities." *Journal of Mathematical Psychology*, 45: 180–188.
- Ledyard, John O. 1986. "The Scope of the Hypothesis of Bayesian Equilibrium." *Journal of Economic Theory*, 39(1): 59–82.
- Luce, Robert D. 1958. *Individual Choice Behavior*, New York: Wiley.
- Manski, Charles F. 1988. "Identification of Binary Response Models." *Journal of the American Statistical Association*, 83 (403): 729–738.
- Matzkin, Rosa L. 1992. "Nonparametric and Distribution-Free Estimation of the Binary Threshold Crossing and the Binary Choice Models." *Econometrica*, 60(2): 239–270.
- McFadden, Daniel L. 1978. "Modelling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models*, eds. Anders Karlqvist et al., Amsterdam: North Holland.
- McFadden, Daniel L. 1981. "Econometric Models of Probabilistic Choice." In *Econometric Models of Probabilistic Choice*, eds. Daniel L. McFadden and Charles F. Manski, Cambridge: MIT Press.
- McFadden, Daniel L. and Kenneth Train. 2000. "Mixed MNL Models for Discrete Response." *Journal of Applied Econometrics*, 15(5): 447–470.
- McFadden, Daniel L. 2005. "Revealed Stochastic Preference: A Synthesis." *Economic Theory*, 26(2): 245–264.
- McKelvey, Richard D. and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10(1), 6–38.

- McKelvey, Richard D. and Thomas R. Palfrey. 1998. "Quantal Response Equilibrium in Extensive Form Games." *Experimental Economics*, 1(1): 9–41.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber. 2000. "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(4): 523–548.
- Rosenthal, Robert W. 1989. "A Bounded-Rationality Approach to the Study of Noncooperative Games." *International Journal of Game Theory*, 18(3): 273–292.
- Seim, Katja. 2002. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." Stanford University Working Paper.
- Signorino, Curtis S. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review*, 93(2): 279–297.
- Sweeting, Andrew. 2004. "Coordination Games, Multiple Equilibria, and the Timing of Radio Commercials." Northwestern University Working Paper.