



Regular Quantal Response Equilibrium

JACOB K. GOEREE

Division of the Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA

CHARLES A. HOLT

Department of Economics, University of Virginia, Charlottesville, VA 22904-4182, USA

THOMAS R. PALFREY

Department of Economics, Princeton University, 112 Fisher Hall, Princeton, NJ 08540, USA

Received June 27, 2004; Revised June 2, 2005; Accepted July 27, 2005

Abstract

The structural Quantal Response Equilibrium (QRE) generalizes the Nash equilibrium by augmenting payoffs with random elements that are not removed in some limit. This approach has been widely used both as a theoretical framework to study comparative statics of games and as an econometric framework to analyze experimental and field data. The framework of structural QRE is flexible: it can be applied to arbitrary finite games and incorporate very general error structures. Restrictions on the error structure are needed, however, to place testable restrictions on the data (Haile et al., 2004). This paper proposes a reduced-form approach, based on quantal response functions that replace the best-response functions underlying the Nash equilibrium. We define a *regular* QRE as a fixed point of quantal response functions that satisfies four axioms: continuity, interiority, responsiveness, and monotonicity. We show that these conditions are not vacuous and demonstrate with an example that they imply economically sensible restrictions on data consistent with laboratory observations. The reduced-form approach allows for a richer set of regular quantal response functions, which has proven useful for estimation purposes.

Keywords: quantal response equilibrium, discrete choice models, reduced-form approach

JEL Classification: D62, C73

1. Probabilistic choice in individual decision making

Probabilistic choice models were introduced by mathematical psychologists to explain why observed decisions appear to be more random as the strength of a stimulus becomes weaker (e.g. Luce, 1959). For example, a subject in an experiment may be asked which of two lights, 1 or 2, is brighter or which of two sounds is louder. Subjects seldom make mistakes when one signal is a lot stronger than another, but the probability that the subject recognizes which signal is stronger falls to nearly one half as the difference in the two intensities goes to zero. One way to model this is that the perceived intensity of signal $j = 1, 2$ is the actual

intensity π_j plus a random shock ε_j . Thus the subject reports that signal 1 is stronger if its perceived strength is higher:

$$\pi_1 + \mu\varepsilon_1 > \pi_2 + \mu\varepsilon_2,$$

where $\mu > 0$ is an error parameter that measures the importance of the epsilon shocks. In this example, the epsilon shocks represent perception errors that may cause decisions to be different if the same person receives the same stimuli again. In other contexts, epsilon shocks appended to economic payoffs may reflect individual heterogeneity in risk aversion, envy, altruism, etc. Of course, it may be possible to model some of these preference shocks explicitly, e.g. when they are correlated with observable demographics. But such models will always leave out some elements, and some degree of “noise” will therefore remain.¹

Regardless of the interpretation of the epsilon shocks, option 1 is chosen if:

$$\frac{\pi_1 - \pi_2}{\mu} > \varepsilon_2 - \varepsilon_1.$$

The probability that this inequality holds can be expressed as $F((\pi_1 - \pi_2)/\mu)$, where $F(\cdot)$ is the distribution function of the difference in the shocks. For example, if the difference in shocks is uniformly distributed around zero, the probability that option 1 is chosen is a piecewise linear function of the payoff difference, with a slope that depends on the error rate. A desirable feature of a probabilistic choice model is that the stronger stimulus is selected more frequently, which occurs if neither shock is more likely to be larger than the other. In this case, $F(0) = 1/2$ and since $F(\cdot)$ is increasing, option j is more likely to be chosen if and only if $\pi_j > \pi_k$ where $k \neq j$, $j, k = 1, 2$. One obvious way to ensure this monotonicity condition is to assume that the two shocks are independently and identically distributed. The i.i.d. assumption is sufficient but not necessary, e.g. in the above example with two options, monotonicity is guaranteed when $F(0) = 1/2$. Indeed, there are many other natural ways to ensure monotonicity without imposing i.i.d. For example, the epsilon shocks could be joint normal random variables with mean zero and arbitrary variance-covariance matrix. However, allowing for completely arbitrary error distributions can lead to nonsensical formulations of stochastic choice, as we show with a simple example in Section 3.2.

2. Probabilistic choice in games: Background

Stochastic choice has been a central theme in game theory since its initial development by von Neumann (1928) and von Neumann and Morgenstern (1944). Here we briefly discuss some of the historical background for the relationship between stochastic choice and equilibrium in games. Hopefully this will add some perspective about the development of ideas leading up to the quantal response equilibrium discussed next.

First, and most obvious to anyone vaguely familiar with game theory, is the notion of a mixed strategy. Players randomize *intentionally* and the probability mixtures they use are

not arbitrary, but are highly restricted by the structure of the game. In two-person zero-sum games, a player mixes in order to ensure the highest guaranteed expected payoff, the minimax payoff. This concept was introduced by von Neumann and Morgenstern (1944), and is closely related to simultaneous developments of the notion of a statistical decision function in theoretical statistics (Wald, 1945).

The idea of using stochastic choice via a mixed strategy as a solution to an *optimization problem* contrasts sharply with the introduction of decision *errors* in the analysis of refinements, as in Selten (1965, 1975) and Myerson (1978). The highly sophisticated concepts they developed (perfection and properness), are based, ironically, on the notion that players do not always maximize, and so the only reasonable equilibria are ones that are robust to “trembles” by opponents. Also in contrast with mixed strategies, these error-based refinements are defined in terms of asymptotic properties when the errors become negligible.

Third, statisticians realized early on that mixed strategies (or errors, for that matter) could be transformed into pure strategies (*purified*) by introducing private information that may cause one player’s action to be random from another player’s point of view. The first discoveries of this sort emerged in tandem with the theory of statistical decision functions (Dvoretzky et al., 1951), and these ideas resurfaced in a game theoretic context via the work of Harsanyi (1973). These purification results involved extending a game of complete information to one of incomplete information by introducing additive payoff disturbances.² Harsanyi (1973) demonstrated that this purification process yields a refinement of the Nash equilibrium when the payoff disturbances become arbitrarily small. In this landmark paper, three key roles played by stochastic choice in game theory (mixed strategies, equilibrium refinements, and private information) were tied together.

The application of stochastic choice models arrived late to experimental game theory. Articles by Harless and Camerer (1994) and Hey and Orme (1994) use stochastic choice models to analyze data from one-person games of incomplete information, designed to stress-test the axioms of expected utility. The former considers completely random “trembles” while the latter looks at stochastic choice models of the sort developed by econometricians, where decision errors are payoff related. That is, mistakes are inversely related to expected payoffs. Trembles in multi-person decision problems were considered by McKelvey and Palfrey (1992) in their analysis of data from centipede games, but, as in Harless and Camerer (1994), the errors were unrelated to payoffs. In the next section, we discuss a more general approach that allows for payoff sensitive errors and for a wide range of error structures generated by additive, privately-observed, payoff disturbances.

3. Structural approach to quantal response equilibrium

Despite the widespread use of probabilistic choice, or “quantal response,” models in the econometric analysis of individual choice data (e.g. McFadden, 1981), the theoretical analysis of strategic choice has, until recently, relied almost entirely on assumptions of non-stochastic decision making. In particular, the notion of a Nash equilibrium rests on an assumption of perfect rationality, where the choice of one strategy over another is independent

of the magnitude of the payoff difference, and inferior strategies are played with zero probability. About a decade ago, a number of economists began incorporating noise elements into the analysis of games. An early example is provided by Rosenthal (1989) who assumes that players' choice probabilities are *linearly* increasing in expected payoffs, as in the uniform probabilistic choice model mentioned in Section 1. The Quantal Response Equilibrium introduced by McKelvey and Palfrey (1995), provides a general framework to extend the probabilistic choice approach to the case of multiple decision-makers. We summarize that framework below.

Let $\Gamma = [I, \{S_i\}_{i=1}^n, \{\pi_i\}_{i=1}^n]$ be a normal-form game, where:

1. $I = \{1, \dots, n\}$ is the set of **players**.
2. $S_i = \{s_{i1}, \dots, s_{iJ_i}\}$ is player i 's set of **strategies** and $S = S_1 \times \dots \times S_n$ is the set of **strategy profiles**.
3. $\pi_i : S \rightarrow \Re$ is player i 's **payoff function**.

Let $\Sigma_i \equiv \Delta^{J_i}$ be the set of probability distributions over S_i . An element $\sigma_i \in \Sigma_i$ is a **mixed strategy**, which is a mapping from S_i to Σ_i , where $\sigma_i(s_i)$ is the probability that player i chooses pure strategy s_i . Let $\Sigma = \prod_{i \in I} \Sigma_i$ be the set of mixed strategy profiles. Given a mixed-strategy profile $\sigma \in \Sigma$, player i 's expected payoff is $\pi_i(\sigma) = \sum_{s \in S} p(s) \pi_i(s)$, where $p(s) = \prod_{i \in I} \sigma_i(s_i)$ is the probability distribution over pure-strategy profiles induced by σ .

For each i and each $j \in \{1, \dots, J_i\}$, and for any $\sigma \in \Sigma$, denote by $\pi_{ij}(\sigma)$ the expected payoff to i of adopting the pure strategy s_{ij} when the other players use σ_{-i} . It is assumed that for each pure strategy s_{ij} , there is an additional privately observed payoff disturbance, ε_{ij} and we denote i 's disturbed payoff by:

$$\hat{\pi}_{ij}(\sigma) = \pi_{ij}(\sigma) + \mu_i \varepsilon_{ij} \quad (3.1)$$

where the error rate μ_i is a strictly positive real number.

For each i , player i 's profile of payoff disturbances, $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iJ_i})$, has a joint distribution with several properties. First, one assumes *absolute continuity*, specifically the distribution of ε_i can be represented by a density function $f_i(\varepsilon_i)$ and the marginal densities exist for each ε_{ij} . Furthermore, the disturbances are *independent across players* (not necessarily across strategies) and *unbiased*, in the sense that $E(\varepsilon_i) = 0$ for all i . Call $f = (f_1, \dots, f_n)$ *admissible* if f_i satisfies the above properties for all i .

The assumed choice behavior is that each player chooses strategy s_{ij} when $\hat{\pi}_{ij}(\sigma) \geq \hat{\pi}_{ik}(\sigma)$ for all $k = 1, \dots, J_i$. Given this choice behavior, $\pi = (\pi_1, \dots, \pi_n)$ and $f = (f_1, \dots, f_n)$ together induce a distribution over the actual choices by each player. To be more specific, for any π , define $B_{ij}(\pi)$ to be the set of realizations of ε_i such that strategy s_{ij} has the highest disturbed expected payoff, $\hat{\pi}$. Then

$$P_{ij}(\pi) = \int_{B_{ij}(\pi)} f(\varepsilon) d\varepsilon \quad (3.2)$$

is the induced choice probability that player i selects strategy j . P_i , which maps π_i into Σ_i is called i 's structural **Quantal Response Function**. Since $P(\pi) \in \Sigma$ and $\pi = \pi(\sigma)$ is defined for any $\sigma \in \Sigma$, $P \circ \pi(\sigma) = P(\pi(\sigma))$ defines a mapping from Σ into itself.

Definition 1. Let $f(\varepsilon)$ be admissible. A structural **Quantal Response Equilibrium** of the normal-form game Γ is a mixed-strategy profile σ^* such that $\sigma^* = P(\pi(\sigma^*))$, with P defined in (3.2).³

Existence of a QRE is proven by McKelvey and Palfrey (1995) and follows from Brouwer's fixed-point theorem since P , and hence $P \circ \pi$, is continuous.

Proposition 1 *There exists a structural Quantal Response Equilibrium of the normal-form game Γ for any admissible $f(\varepsilon)$.*

The main difference with the decision-making context discussed previously is that, in a quantal response *equilibrium*, the exogenously given intensities are replaced by endogenously determined expected payoffs. For example, in a two-by-two game, players' expected payoffs $\pi = (\pi_{11}, \pi_{12}, \pi_{21}, \pi_{22})$ are:

$$\begin{aligned}\pi_{11} &= P_{21}\pi_1(s_{11}, s_{21}) + P_{22}\pi_1(s_{11}, s_{22}) \\ \pi_{12} &= P_{21}\pi_1(s_{12}, s_{21}) + P_{22}\pi_1(s_{12}, s_{22}) \\ \pi_{21} &= P_{11}\pi_2(s_{11}, s_{21}) + P_{12}\pi_2(s_{12}, s_{21}) \\ \pi_{22} &= P_{11}\pi_2(s_{11}, s_{22}) + P_{12}\pi_2(s_{12}, s_{22})\end{aligned}\tag{3.3}$$

where, for each $i, j = 1, 2$:

$$P_{ij} = \text{Prob}\left[\frac{\pi_{ij} - \pi_{ik}}{\mu_i} > \varepsilon_{ik} - \varepsilon_{ij}, \quad k \neq j\right].\tag{3.4}$$

For example, assuming identical and independent extreme-value disturbances implies that P_{ij} is given by the familiar logit formula:

$$P_{ij} = \frac{\exp(\pi_{ij}/\mu_i)}{\exp(\pi_{ij}/\mu_i) + \exp(\pi_{ik}/\mu_i)}, \quad i, j = 1, 2, \quad k \neq j.\tag{3.5}$$

Note that (3.5) is not an explicit solution, since the payoffs on the right-side are themselves functions of the choice probabilities. To compute the (logit) quantal-response equilibrium, the probabilities used to calculate the expected payoffs in (3.3) must match the probabilities that follow from (3.4) as determined by the expected payoffs. That is, the equations given by (3.3) and (3.4) must hold simultaneously.

This fixed-point element complicates the computation of (logit) QRE in all but the simplest games, but also captures the intuitive notion that players react to others' noisiness. Indeed, in some games this type of feedback causes the logit equilibrium to be arbitrarily far away from the Nash equilibrium. For example, for some parameterizations of the traveler's

dilemma (Capra et al., 1999), the Nash equilibrium is at the lower end of the strategy set while the logit equilibrium puts almost all mass at the upper end.

3.1. *Restrictions implied by additivity*

Existing results on the theory of probabilistic choice show that the additive disturbance model in (3.1) puts some restrictions on the choice probabilities. These conditions, listed in Proposition 2 for later reference, are easy to interpret.⁴ Translation Invariance implies that if a constant is added to all payoffs, choice probabilities do not change. Symmetry says that the effect of an increase in strategy k 's payoff on the probability of choosing strategy j is the same as the effect of an increase in strategy j 's payoff on the probability of choosing strategy k . Finally, Strong Substitutability implies (among other things) that if the payoff of strategy k rises, the probability of choosing any of the other strategies $j \neq k$ falls.

Proposition 2 *The structural quantal response functions defined in (3.2) satisfy:*

- (i) **Translation Invariance:** $P_{ij}(\pi_i + ce_{J_i}) = P_{ij}(\pi_i)$ for all $c \in \mathfrak{R}$ and $\pi_i \in \mathfrak{R}^{J_i}$, where $e_{J_i} = (1, \dots, 1)$.
- (ii) **Symmetry:** $\partial P_{ij} / \partial \pi_{ik} = \partial P_{ik} / \partial \pi_{ij}$ for all $i, j = 1, \dots, J_i$ and $\pi_i \in \mathfrak{R}^{J_i}$.
- (iii) **Strong Substitutability:** $(-1)^\ell \partial^\ell P_{ij} / \partial \pi_{ik_1} \dots \partial \pi_{ik_\ell} \geq 0$ for all $1 \leq \ell \leq J_i - 1$, $k_1 \neq \dots \neq k_\ell \neq j$ and $\pi_i \in \mathfrak{R}^{J_i}$.

Proof: Translation Invariance and Strong Substitutability follow more or less directly from the definition of the structural quantal response functions:

$$P_{ij} = \int_{-\infty}^{\infty} \int_{-\infty}^{\frac{\pi_{ij} - \pi_{i1}}{\mu_i} + x} \dots \int_{-\infty}^{\frac{\pi_{ij} - \pi_{iJ_i}}{\mu_i} + x} f(z_1, \dots, x, \dots, z_{J_i}) dz_{J_i} \dots dz_1 dx,$$

where x corresponds to ε_{ij} in (3.4). Since the expression for P_{ij} involves only payoff differences, adding a constant to all payoffs changes nothing in the limits of integration, or elsewhere. Differentiating with respect to $\partial \pi_{ik_1} \dots \partial \pi_{ik_\ell}$ and multiplying by $(-1)^\ell$ removes ℓ integrals and leaves us with $J_i - \ell$ integrals of a (positive) density function. We next prove Symmetry.

$$\frac{\partial P_{ij}}{\partial \pi_{ik}} = \frac{1}{\mu_i} \int_{-\infty}^{\infty} \int_{-\infty}^{\frac{\pi_{ij} - \pi_{i1}}{\mu_i} + x} \dots \int_{-\infty}^{\frac{\pi_{ij} - \pi_{iJ_i}}{\mu_i} + x} f\left(z_1, \dots, x, \dots, \frac{\pi_{ij} - \pi_{ik}}{\mu_i} + x, \dots, z_{J_i}\right) dz_{J_i} \dots dz_1 dx,$$

where the integral with respect to z_k was removed. Now transform $t = (\pi_{ij} - \pi_{ik}) / \mu_i + x$, and rewrite:

$$\frac{\partial P_{ij}}{\partial \pi_{ik}} = -\frac{1}{\mu_i} \int_{-\infty}^{\infty} \int_{-\infty}^{\frac{\pi_{ik} - \pi_{i1}}{\mu_i} + t} \dots \int_{-\infty}^{\frac{\pi_{ik} - \pi_{iJ_i}}{\mu_i} + t} f\left(z_1, \dots, \frac{\pi_{ik} - \pi_{ij}}{\mu_i} + t, \dots, z_{J_i}\right) dz_{J_i} \dots dz_1 dt,$$

$$\begin{aligned}
& +t, \dots, t, \dots, z_{J_i} \Big) dz_{J_i} \cdots dz_1 dt \\
& = \frac{\partial P_{ik}}{\partial \pi_{ij}},
\end{aligned}$$

where the last step follows from the definition of P_{ik} . \square

In the next section we show that the restrictions of Proposition 2 are not sufficient to rule out unintuitive results.

3.2. Empirical restrictions of structural QRE

The quantal response equilibrium has been primarily used to analyze data patterns generated in laboratory experiments, e.g. auctions, bargaining, market entry, voting, voluntary contributions, rent seeking, coordination, and social dilemma games (McKelvey and Palfrey, 1995; Anderson et al., 1998, 2001) and extensive-form and repeated games (McKelvey and Palfrey, 1995; Camerer et al., 2002; Chong et al., 2005). In past studies, the incorporation of noise elements has resulted in much better explanations of laboratory data than other equilibrium models. For example, changes in payoff parameters that have no effect on Nash equilibrium predictions often have large, dramatic effects on subjects' decisions, and these effects can be explained by a quantal response equilibrium (e.g. Capra et al., 1999; Goeree and Holt, 2001). In all of these applications, the theoretical analysis and the estimation were based on an assumption that the epsilon shocks are i.i.d. Even though the most general versions of probabilistic choice and quantal response models do not restrict the nature of the shock distributions (beyond admissibility), such restrictions are needed to obtain useful empirical predictions.

Indeed, Haile et al. (2004) showed quite generally that any pattern of choice probabilities can be rationalized with an additive random utility model if one relaxes the i.i.d. assumption *and* if one has the freedom to customize the joint distribution of payoff disturbances. In order to assess the methodological implications of the Haile et al. (2004) result, it is useful to consider how *fully general* probabilistic choice models can be indeterminate in single-person decision-making problems like the early perception experiments.

We next present two simple examples to illustrate how, in the absence of restrictions on the joint distribution of disturbances, a standard, single decision-maker, two-alternative, probabilistic choice model can be rigged so that the probability of choosing the low-payoff outcome is arbitrarily close to 1. Suppose payoffs are $\pi_{11} = 1$ and $\pi_{12} = 2$. The probability of choosing option 1 is given by⁵

$$P_{11} = \text{Prob}[1 + \varepsilon_{11} > 2 + \varepsilon_{12}],$$

and the probability of choosing option 2 is $P_{12} = 1 - P_{11}$. The epsilon shocks are not required to be independently and identically distributed, but each will have mean 0.⁶ We next show that for any $p \in (0, 1)$ we can find an admissible $f(\varepsilon_{11}, \varepsilon_{12})$ such that $P_{11} = p$,

if the ε_{1j} are either (i) independently but not identically distributed or (ii) identically but not independently distributed.

- (i) *Independent but not Identical Marginal Distributions*: Since $\pi_{11} - \pi_{12} = 1$, the idea is to let ε_{11} exceed ε_{12} by more than 1 with probability p . To ensure that ε_{11} has mean zero, it has to be sufficiently negative with the complementary probability, $1 - p$. For example, let $\varepsilon_{12} = 0$ with probability 1, and let $\varepsilon_{11} = 2$ with probability p and $\varepsilon_{11} = -2p/(1 - p)$ with probability $1 - p$. Then

$$P_{11} = \text{Prob}[1 + \varepsilon_{11} > 2] = p.$$

- (ii) *Identical but not Independent Marginal Distributions*: Here the idea is to let ε_{11} be perfectly correlated with ε_{12} , such that ε_{11} exceeds ε_{12} by more than 1 for a large set of ε_{12} values. For example, let ξ be a random variable, uniformly distributed on $[0, 2/(1 - p)]$. Furthermore, let $\varepsilon_{12} = \xi$ and $\varepsilon_{11} = (\xi + 2) \bmod (2/(1 - p))$. That is, ε_{11} is just equal to $\varepsilon_{12} + 2$, except when this shift to the right puts ε_{11} above the upper bound of the support of the distribution of ε_{12} , in which case it is inserted in the lower part of the support. Hence $\varepsilon_{11} = \varepsilon_{12} + 2 > \varepsilon_{12}$ when ε_{12} is between 0 and $2/(1 - p) - 2$, and $\varepsilon_{11} = \varepsilon_{12} + 2 - 2/(1 - p) < \varepsilon_{12}$ otherwise. By construction, the marginal distributions of ε_{11} and ε_{12} are both uniform on $[0, 2/(1 - p)]$. So

$$P_{11} = \frac{2/(1 - p) - 2}{2/(1 - p)} = p.$$

To summarize, *without further restrictions on the error distributions, probabilistic choice models can be constructed to predict any observed behavior*. In particular, when the i.i.d. assumption is dropped, the choice probability of the less attractive option can be made arbitrarily close to 1. While the examples above concern a one-person game with only two possible strategies, the logic of the construction can be extended in a straightforward way to allow for an arbitrary number of players and strategies (Haile et al., 2004).

This raises important questions about fitting the Quantal Response Equilibrium to data, because QRE can fit the aggregate frequencies of any data set perfectly if one allows arbitrary error structures. However, this does not imply that probabilistic choice models, or QRE in particular, are without empirical content. It simply means that (economically sensible) assumptions about the epsilon shocks are needed before one can reach meaningful conclusions. Haile et al. (2004, p. 2) emphasize this point as well: “Any restrictions on outcomes obtainable from a QRE comes only from a priori restrictions on the distributions of payoff perturbations.”

Before we can discuss the necessary restrictions on the structure of the payoff perturbations, we first determine what economically meaningful properties the implied quantal

response functions should satisfy, which is the topic of the next section. In Section 5 we then derive conditions on the shock distributions that implement these properties.

4. Regular quantal response equilibrium

When the QRE concept is applied to analyze experimental data, one typically imposes restrictions that require choice probabilities to be monotone in expected payoffs. Indeed, in most papers, attention is further restricted to an i.i.d. logit probabilistic choice model. Sometimes a different error parameter, μ , is estimated for different treatments, but this one-dimensional degree of freedom is generally not sufficient to generate *any* outcome as will be shown below. Also theoretical work on quantal-response comparative statics has often assumed an i.i.d. error structure, although sometimes more general results are possible (e.g. with the $F(0) = 1/2$ restriction in the two-option case). The intuitive assumption of monotonicity has always been imposed, since otherwise economically uninteresting outcomes arise, as the examples above demonstrate.

Definition 2. $P_i : \mathfrak{R}^{J_i} \rightarrow \Delta^{J_i}$ is a **regular quantal response function** if it satisfies the following four axioms.

- (A1) **Interiority:** $P_{ij}(\pi_i) > 0$ for all $j = 1, \dots, J_i$ and for all $\pi_i \in \mathfrak{R}^{J_i}$.
- (A2) **Continuity:** $P_{ij}(\pi_i)$ is a continuous and differentiable function for all $\pi_i \in \mathfrak{R}^{J_i}$.
- (A3) **Responsiveness:** $\frac{\partial P_{ij}(\pi_i)}{\partial \pi_{ij}} > 0$ for all $j = 1, \dots, J_i$ and for all $\pi_i \in \mathfrak{R}^{J_i}$.
- (A4) **Monotonicity:** $\pi_{ij} > \pi_{ik} \Rightarrow P_{ij}(\pi_i) > P_{ik}(\pi_i)$ for all $j, k = 1, \dots, J_i$.

Intuitively, regular quantal response functions are generalizations (smoothings) of best response functions, and as such, replace strict rational choice with a weaker version. Non-regular quantal response functions that violate Monotonicity would seem to indicate a weak form of *anti*-rationality, which is not the intent of the approach.

Axioms (A1)–(A4) are economically and intuitively compelling. Interiority ensures that the model has full domain, that is, it is logically consistent with all possible data sets. This is important for empirical application of the model. Continuity is a technical restriction, which ensures that P_i is non-empty and single-valued. Furthermore, it seems a natural assumption since arbitrarily small changes in expected payoffs should not lead to jumps in choice probabilities. Responsiveness requires that if the expected payoff of an action increases, ceteris paribus, the choice probability must also increase. Monotonicity is a weak form of rational choice that involves binary comparisons of actions: an action with higher expected payoff is chosen more frequently than an action with a lower expected payoff. One implication, which follows directly from Continuity and Monotonicity, is that strategies with the same expected payoff are chosen with equal probability: $\pi_{ij} = \pi_{ik} \Rightarrow P_{ij}(\pi_i) = P_{ik}(\pi_i)$ for all $j, k = 1, \dots, J_i$.

At first sight, Responsiveness and Monotonicity may seem to have similar consequences. To see that their implications are different consider the two-option case and let player 1's

probability of choosing option 1 be given by $P_{11} = G(\pi_{11} - \pi_{12})$. Then P_{11} is responsive if $G'(x) > 0$ for all x but it is not monotonic if $G(0) \neq \frac{1}{2}$. Likewise, P_{11} is monotonic if $G(x) < \frac{1}{2}$ for $x < 0$ and $G(x) > \frac{1}{2}$ for $x > 0$ but it is not responsive if $G(x)$ is decreasing for some x .

Define $P(\pi) = (P_1(\pi_1), \dots, P_n(\pi_n))$ to be regular if each P_i satisfies the regularity axioms (A1)–(A4). Since $P(\pi) \in \Sigma$ and $\pi = \pi(\sigma)$ is defined for any $\sigma \in \Sigma$, $P \circ \pi$, defines a mapping from Σ into itself.

Definition 3. Let P be regular. A **Regular Quantal Response Equilibrium** (R-QRE) of the normal-form game Γ is a mixed-strategy profile σ^* such that $\sigma^* = P(\pi(\sigma^*))$.

Since regularity of P includes continuity, $P \circ \pi$ is a continuous mapping. Existence of an R-QRE therefore follows directly from Brouwer’s fixed-point theorem.

Proposition 3 *There exists a Regular Quantal Response Equilibrium (R-QRE) of Γ for any regular P .*

4.1. *Empirical restrictions of regular QRE: An example*

Axioms (A1)–(A4) collectively have strong empirical implications, even without any parametric assumptions on P . To illustrate the nature of these restrictions, consider the asymmetric matching-pennies game in figure 1, where Row’s payoff is $X > -1$ when the outcome is (“Up”, “Left”). (See Loomes (2005, this issue) for some complementary discussion about applying the QRE approach to matching-pennies games.) Let p denote the probability with which the Column player chooses “Left” and q denote the probability with which the Row player chooses “Up”. First, consider the case $X \geq 1$. It is readily verified that Row’s expected payoff of choosing “Up” is lower (higher) than of choosing “Down” when $p < 2/(X + 3)$ ($p > 2/(X + 3)$). Monotonicity therefore implies that, if (p^*, q^*) defines an R-QRE, it must satisfy the inequalities:

$$\begin{cases} q^* \leq 1/2 & \text{if } p^* \leq 2/(X + 3), \\ q^* \geq 1/2 & \text{if } p^* \geq 2/(X + 3). \end{cases}$$

		Column	
		Left	Right
Row	Up	X, -1	-1, 1
	Down	-1, 1	1, -1

Figure 1. An asymmetric matching pennies game.

Likewise, Column’s expected payoff of choosing “Left” is higher (lower) than of choosing “Right” when $q < 1/2$ ($q > 1/2$). Thus, (p^*, q^*) must satisfy:

$$\begin{cases} p^* \geq 1/2 & \text{if } q^* \leq 1/2, \\ p^* \leq 1/2 & \text{if } q^* \geq 1/2. \end{cases}$$

The region defined by these inequalities defines the set of possible R-QRE. For the specific case of $X = 9$, this area is shown by the dark gray shaded area in figure 2. The measure of the set of possible R-QRE is $1/2 * 1/3 = 1/6$ in this case, i.e. only 16.7% of all possible outcomes is consistent with some R-QRE.⁷

The case $-1 < X < 1$ can be analyzed in a similar way. The set of regular-QRE for $X = 0$, for instance, is given by the light shaded area in figure 2. Note that the Row player is predicted to choose “Up” more often than “Down” in an R-QRE when $X > 1$, while the reverse is true for $X < 1$. In fact, we next show that if Row’s payoff of the (“Up”, “Left”) outcome rises, Row’s probability of choosing “Up” increases. This intuitive “own payoff” effect contrasts with the Nash equilibrium prediction of no change in Row’s choice probabilities (since they are determined by the requirement that Column is indifferent).

Proposition 4 *In a regular QRE of the asymmetric matching pennies game in figure 1, Row’s probability of choosing “Up” is strictly increasing in X and Column’s probability of choosing “Left” is strictly decreasing in X .*

Proof: The equilibrium probabilities satisfy:

$$\begin{cases} q^* = P_{11}((X + 1)p^* - 1, 1 - 2p^*) \\ p^* = P_{21}(1 - 2q^*, 2q^* - 1) \end{cases} \tag{4.1}$$

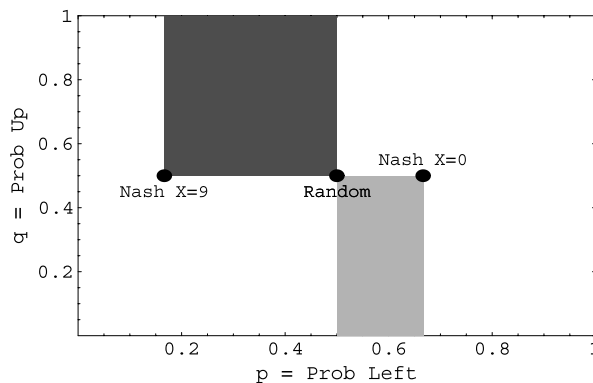


Figure 2. On the horizontal axis is Column’s probability of choosing “Left” and on the vertical axis is Row’s probability of choosing “Up”. The dark gray shaded area contains the set of possible regular QRE for $X = 9$, and the light gray shaded area contains the set of possible regular QRE for $X = 0$.

Responsiveness implies that P_{ij} is strictly increasing in π_{ij} , and since $P_{i2}(\pi_{i1}, \pi_{i2}) = 1 - P_{i1}(\pi_{i1}, \pi_{i2})$, P_{i1} is strictly increasing (decreasing) in its first (second) argument. It follows from the top line that q^* is a strictly increasing function of p^* , while the bottom line implies that p^* is a strictly decreasing function of q^* . Hence, the solution to the fixed-point conditions (4.1) is unique. Moreover, the bottom equality of (4.1) implies that either q^* rises with X and p^* falls, q^* falls with X and p^* rises, or both q^* and p^* remain constant as X rises. The latter two cases are impossible, however, since the top equality of (4.1) implies that q^* rises with X if p^* remains constant or rises with X . Hence, in a regular QRE, an increase in X results in a strict increase of q^* and a strict decrease of p^* . \square

Note that the non-parametric comparative statics of figure 2 are mainly driven by Monotonicity, while Proposition 4 rests heavily on Responsiveness. The own-payoff effect predicted by regular QRE accords with data from laboratory experiments that employ an asymmetric matching-pennies structure, e.g. Ochs (1995), McKelvey et al. (2000), Goeree and Holt (2001), and Goeree et al. (2003).

Goeree et al. (2005) discuss the empirical restrictions of regular QRE in more general n -player games. They prove that the measure of equilibria is bounded above by $1/K!$ where $K = \max_i K_i$ and K_i is the number of pure strategies for player $i = 1, \dots, n$. In other words, the measure of the set of regular equilibria shrinks fast as the number of strategies and players increase.

5. Regularity within the structural approach

In the definition of structural QRE in Section 3, axioms (A1)–(A3) are satisfied by any choice probabilities generated by admissible payoff disturbances satisfying a full support condition.⁸ But in general, it is the Monotonicity axiom (A4) that can be violated for some disturbance distributions (such as the ones in the examples of Section 3.2). However, as an axiom about economic choice behavior, Monotonicity seems quite reasonable. So, a natural question to ask is: “What admissible distributions of payoff disturbances generate regular quantal response functions?”

McKelvey and Palfrey (1995, 1996) point out that a sufficient condition for Monotonicity is i.i.d., but this restriction can obviously be relaxed. For example, a joint normal distribution with identical means, variances, and covariances will generate monotone responses. In McKelvey and Palfrey (1994), a condition called “label independence” is identified to ensure regularity. Label independence means that choice probabilities depend on expected payoffs only and not on the labels of the strategies. For example, suppose there are three options with payoffs $\pi_{11} = 1$, $\pi_{12} = 2$, and $\pi_{13} = 3$ and corresponding choice probabilities $P_{11} = \frac{1}{6}$, $P_{12} = \frac{1}{3}$ and $P_{13} = \frac{1}{2}$. Label independence implies that if payoffs change to $\pi_{11} = 2$, $\pi_{12} = 3$, and $\pi_{13} = 1$, the choice probabilities become $P_{11} = \frac{1}{3}$, $P_{12} = \frac{1}{2}$, and $P_{13} = \frac{1}{6}$. In other words, when the payoffs are permuted so are the choice probabilities.

The proof of Proposition 5 below shows that label independence of the choice probabilities can be ensured by requiring the additive payoff disturbances to be interchangeable random variables, which includes i.i.d. as a special case (Karlin, 1966).⁹ Let Ψ_J denote the set of all

possible permutations of J_j objects. We say that the density of payoff disturbances satisfies **interchangeability** if, for all $\psi \in \Psi_{J_i}$, $f(\varepsilon_1, \dots, \varepsilon_{J_i}) = f(\varepsilon_{\psi(1)}, \dots, \varepsilon_{\psi(J_i)})$.

Proposition 5 *The quantal response function defined in (3.2) is regular if $f(\varepsilon)$ satisfies admissibility and interchangeability.*

Proof: Admissibility of $f(\varepsilon)$ guarantees axioms (A1)–(A3). To prove Monotonicity, we first demonstrate that interchangeability implies label independence:

$$\begin{aligned} P_{ij}(\pi_{i\psi(1)}, \dots, \pi_{i\psi(J_i)}) &= \text{Prob}[\pi_{i\psi(j)} + \varepsilon_{ij} \geq \pi_{i\pi(k)} + \varepsilon_{ik}, k = 1, \dots, J_i] \\ &= \text{Prob}[\pi_{i\psi(j)} + \varepsilon_{i\psi(j)} \geq \pi_{i\psi(k)} + \varepsilon_{i\psi(k)}, k = 1, \dots, J_i] \\ &= P_{i\psi(j)}(\pi_{i1}, \dots, \pi_{iJ_i}), \end{aligned} \quad (5.1)$$

where we used interchangeability of the random variables ε_i in going from the first to the second line. Label independence implies that $P_{ij}(\pi) = P_{ik}(\pi)$ if $\pi_{ij} = \pi_{ik}$, as can be seen by restricting ψ to be a pairwise permutation of j and k in (5.1). Recall from Proposition 2 that P_{ij} rises with π_{ij} and falls with all the π_{ik} for $k \neq j$. Likewise, P_{ik} falls with π_{ij} when $j \neq k$. Together with $P_{ij} = P_{ik}$ when $\pi_{ij} = \pi_{ik}$ this implies that $P_{ij} > P_{ik}$ when $\pi_{ij} > \pi_{ik}$, because as π_{ij} rises from π_{ik} , P_{ij} strictly increases and P_{ik} strictly decreases. \square

Interchangeability is not a necessary condition in the sense that the quantal response function defined in (3.2) may be regular without it. Consider, for instance, the case of two options and let ε_{11} and ε_{12} be independent normal random variables, both with mean zero but with different variances σ_{11}^2 and σ_{12}^2 respectively, so that the difference $\varepsilon_{12} - \varepsilon_{11}$ is normally distributed with mean zero and variance $\sigma^2 = \sigma_{11}^2 + \sigma_{12}^2$. The quantal response functions are given by $P_{11} = \Phi((\pi_{11} - \pi_{12})/\sigma)$ and $P_{12} = \Phi((\pi_{12} - \pi_{11})/\sigma)$, where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal variable. In this example, Monotonicity of the quantal response functions holds even though ε_{11} and ε_{12} are not interchangeable random variables.

We next show with a simple three-option example that without interchangeability, regularity may fail to hold. The construction used in Section 3.2 suggests that non-monotone response functions may only arise if one makes unrealistic assumptions about the disturbances (see Haile et al., 2004). However, more reasonable assumptions about disturbances can lead to non-monotone choice probabilities with three or more alternatives. Consider the following individual choice problem with disturbed payoffs (where we normalized $\mu_1 = 1$):

$$\begin{aligned} \hat{\pi}_{11} &= 1 + \varepsilon_{11}, \\ \hat{\pi}_{12} &= 2 + \varepsilon_{12}, \\ \hat{\pi}_{13} &= 3 + \varepsilon_{13}. \end{aligned}$$

Further suppose the disturbances are independent and normally distributed, each with mean 0, but with different variances. Let the variance of ε_{11} be 10 and the variances of ε_{12} and

ε_{13} be 0.01. Then alternatives 1 and 3 will each be chosen with probability approximately 1/2, and alternative 2 will be chosen with probability close to 0. However alternative 2 has a higher expected payoff than 1. The intuition is that alternatives with highly variable payoffs can be chosen with fairly high probability even when their expected payoff is low.

Modifying this example shows that the highest ranked alternative may not be the one that is most likely to be chosen. Simply change the variance of ε_{12} from 0.01 to 10. Now alternatives 1 and 2 are each chosen with probability approximately 3/8, while alternative 3 is chosen with probability approximately 1/4. Indeed, if we change alternative 1's disturbance variance from 10 to 100, then the worst (in expected payoff) alternative is chosen with probability 1/2 and the other two are chosen with probability 1/4. By continuity, it is easy to see that one could find a variance for ε_{11} somewhere between 10 and 100 so that the choice probabilities of the alternatives are *in reverse order* of their expected payoffs.

These examples suggest that it may be possible to find some plausible non-interchangeable error structures for games (or individual choice problems) that could lead to non-monotone choice probabilities. To our knowledge, however, there is no systematic empirical evidence of such behavior.

6. Reduced-form approach to regular QRE

In the previous section we have shown that regularity can be salvaged within the structural approach by imposing interchangeability of the additive payoff disturbances. The resulting label independence of the choice probabilities is an appealing property, which implies that choice probabilities depend on expected payoffs but not on the precise labeling of the strategies. However, as Proposition 2 makes clear, the additive structural approach also imposes other conditions on the choice probabilities: Translation Invariance, Symmetry, and Strong Substitutability. These latter conditions do *not* translate into sensible empirical restrictions on possible QRE outcomes (see Section 3.2), since they are not related to Monotonicity. Moreover, they result from a modeling assumption (of *additive* payoff disturbances) but are *not* derived from economic principles. As a consequence, they may lead to implausible or empirically false restrictions in certain contexts.

For example, Translation Invariance is not plausible in settings where the magnitudes of perception errors or preference shocks depend on the magnitudes of expected payoffs. For instance, a 25 cent error is unlikely for a decision involving pennies, but such an error would be common in decisions involving hundreds of dollars. One way to model scale-dependent shocks, when payoffs are positive, is to have them be multiplicative so that the person selects the decision with the highest product: $(\pi_{ij})^{1/\mu_i} \varepsilon_{ij}$, where the shock has mean 1. This product can be transformed into a sum by taking a natural log, which shows that the option for which $\log(\pi_{ij})/\mu_i + \log(\varepsilon_{ij})$ is the highest is chosen. If the logs of the shocks have identical extreme-value distributions, then the π_{ij}/μ_i term in (3.5) is replaced $\log(\pi_{ij})/\mu_i$ and the resulting probabilistic choice function obeys a power law:

$$P_{ij} = \frac{(\pi_{ij})^{1/\mu_i}}{\sum_{k=1}^{J_i} (\pi_{ik})^{1/\mu_i}}, \quad i = 1, \dots, n, \quad j = 1, \dots, J_i. \quad (6.1)$$

This formulation was first derived by Luce (1959) from simple choice axioms. Note that the choice probabilities in (6.1) are invariant when multiplying the payoffs by a positive constant, but not when adding a constant. Power-law choice rules have been shown to provide better fits to laboratory choice data when the experiment involves large changes in payoff scale across treatments, e.g. from several dollars to several hundred dollars (Laury and Holt, 2002). In particular, they found that the variation in choice data did not diminish with increases in the payoff scale, so multiplicative nature of the shock specification provided a better fit. The power-law choice model was also used by Goeree et al. (2002) to analyze data from different first-price auction treatments.

The Luce power-law rule in (6.1) can be generalized easily to allow for negative payoffs. For instance, let $g : \Re \rightarrow \Re_+$ be a strictly positive and strictly increasing function, and define:

$$P_{ij} = \frac{g(\pi_{ij}/\mu_i)}{\sum_{k=1}^{J_i} g(\pi_{ik}/\mu_i)}, \quad i = 1, \dots, n \quad j = 1, \dots, J_i. \quad (6.2)$$

For general $g(\cdot)$ functions, these choice probabilities do not satisfy Translation Invariance nor Symmetry, and, hence, cannot be derived via the structural approach. Note that the choice probabilities in (6.2) do satisfy Strong Substitutability. However, also this condition may not be desirable since there is no *a priori* reason why the probability of choosing strategy 1, say, should fall when the payoff of strategy 3 rises.^{10,11} Responsiveness implies that the probability of choosing a certain strategy should go up if its payoff rises, but there is no reason that the probabilities with which the other strategies are chosen should all go down.

To summarize, while the structural approach can be amended to ensure monotonicity (Proposition 5), this approach restricts the quantal response functions in undesirable ways likely to be contradicted by laboratory and field data. We therefore introduce a more general approach to modeling regular quantal response equilibrium, which we call the **reduced-form** approach. This approach takes axioms (A1)–(A4) as primitives and allows for any quantal response functions that satisfy these axioms without requiring that they are derived from some underlying model of stochastic utility maximization.¹²

The main point of this section is that there are many quantal response functions that satisfy the regularity conditions (A1)–(A4), but that cannot be derived via the structural approach. One class of examples is generated by the generalized Luce model in (6.2), which violates Translation Invariance and Symmetry but satisfies Strong Substitutability because of an underlying independence-of-irrelevant-alternatives (IIA) property. Also Substitutability is not implied by Regularity, however, as the next example illustrates.

Proposition 6 *The following quantal response functions*

$$P_{11}(\pi_{11}, \pi_{12}, \pi_{13}) = \frac{\frac{e^{\pi_{11}}}{1+\pi_{12}^2+\pi_{13}^2}}{\frac{e^{\pi_{11}}}{1+\pi_{12}^2+\pi_{13}^2} + \frac{e^{\pi_{12}}}{1+\pi_{11}^2+\pi_{13}^2} + \frac{e^{\pi_{13}}}{1+\pi_{11}^2+\pi_{12}^2}},$$

$$P_{12}(\pi_{11}, \pi_{12}, \pi_{13}) = \frac{\frac{e^{\pi_{12}}}{1 + \pi_{11}^2 + \pi_{13}^2}}{\frac{e^{\pi_{11}}}{1 + \pi_{12}^2 + \pi_{13}^2} + \frac{e^{\pi_{12}}}{1 + \pi_{11}^2 + \pi_{13}^2} + \frac{e^{\pi_{13}}}{1 + \pi_{11}^2 + \pi_{12}^2}},$$

$$P_{13}(\pi_{11}, \pi_{12}, \pi_{13}) = \frac{\frac{e^{\pi_{13}}}{1 + \pi_{11}^2 + \pi_{12}^2}}{\frac{e^{\pi_{11}}}{1 + \pi_{12}^2 + \pi_{13}^2} + \frac{e^{\pi_{12}}}{1 + \pi_{11}^2 + \pi_{13}^2} + \frac{e^{\pi_{13}}}{1 + \pi_{11}^2 + \pi_{12}^2}},$$

are Regular, but violate Translation Invariance, Symmetry, and Strong Substitutability, and, hence, cannot be generated by the structural approach.

Proof: Interiority and Continuity are obvious. To show Responsiveness, consider P_{11} :

$$P_{11}(\pi_{11}, \pi_{12}, \pi_{13}) = \left[1 + \frac{e^{\pi_{12}}}{e^{\pi_{11}}} \left(\frac{1 + \pi_{12}^2 + \pi_{13}^2}{1 + \pi_{11}^2 + \pi_{13}^2} \right) + \frac{e^{\pi_{13}}}{e^{\pi_{11}}} \left(\frac{1 + \pi_{12}^2 + \pi_{13}^2}{1 + \pi_{11}^2 + \pi_{12}^2} \right) \right]^{-1}.$$

Define $f(\pi) = e^\pi (1 + \pi^2 + \tilde{\pi}^2)$ and note that its derivative is $f'(\pi) = e^\pi ((1 + \pi)^2 + \tilde{\pi}^2)$, so $f(\pi)$ is a strictly increasing function of π . Hence, P_{11} is strictly increasing in π_{11} . The proofs of Responsiveness of P_{12} and P_{13} are similar. To show Monotonicity consider

$$\frac{P_{11}(\pi_{11}, \pi_{12}, \pi_{13})}{P_{12}(\pi_{11}, \pi_{12}, \pi_{13})} = \frac{e^{\pi_{11}}}{e^{\pi_{12}}} \left(\frac{1 + \pi_{11}^2 + \pi_{13}^2}{1 + \pi_{12}^2 + \pi_{13}^2} \right).$$

Again, since $f(\pi)$ is strictly increasing in π , $P_{11} > P_{12}$ ($P_{11} < P_{12}$) if and only if $\pi_{11} > \pi_{12}$ ($\pi_{11} < \pi_{12}$). The proof that Monotonicity holds generally can be done in an analogous manner by considering P_{11}/P_{13} and P_{12}/P_{13} . Translation Invariance is obviously not satisfied, and the fact that Symmetry does not generally hold can be verified by direct computation. The novel feature of this example is that substitutability does not generally hold. For instance, $\partial P_{11}/\partial \pi_{13} = 2e^3/(3 + 2e + 3e^2)^2 > 0$ when $\pi_{11} = 1$, $\pi_{12} = 0$, and $\pi_{13} = -1$, which illustrates the possibility that the probability of choosing option 1 may be increasing in the payoff of option 3. \square

6.1. Comparing structural and reduced-form regular QRE

Proposition 6 shows that the reduced-form approach allows for more general regular quantal response functions. However, the set of possible regular quantal response equilibria generated by regular reduced-form response functions is not necessarily larger than that generated by regular structural response functions.

To see this, consider again the asymmetric matching pennies game in figure 1 with $X = 9$, for which the set of possible reduced form R-QRE are given by the dark gray shaded area in figure 2.¹³ We next demonstrate that any point in this area can be obtained as a structural QRE with additive i.i.d. payoff disturbances. Let the ε_{ij} for $i, j = 1, 2$ be i.i.d. with admissible density \tilde{f} . Furthermore, let f denote the density of the differences $\varepsilon_{i1} - \varepsilon_{i2}$

for $i = 1, 2$, with cumulative distribution F . Note that f is symmetric so $F(0) = 1/2$. The equations defining the structural QRE are

$$\begin{cases} q^* = F\left(\frac{12p^*-2}{\mu_1}\right) \\ p^* = F\left(\frac{2-4q^*}{\mu_2}\right) \end{cases}$$

where, as before, p^* is the equilibrium probability with which Column chooses “Left” and q^* is the equilibrium probability with which Row chooses “Up”. The fixed-point equations can be inverted to yield

$$\begin{cases} \mu_1 = (12p^* - 2)/F^{(-1)}(q^*) \\ \mu_2 = (2 - 4q^*)/F^{(-1)}(p^*) \end{cases}$$

where $F^{(-1)}$ denotes the inverse of the distribution function, with $F^{(-1)}(p) < 0$ when $p < \frac{1}{2}$ and $F^{(-1)}(p) > 0$ when $p > \frac{1}{2}$. Note that μ_1 and μ_2 so defined are positive for all $q \in (\frac{1}{2}, 1)$ and $p \in (\frac{1}{6}, \frac{1}{2})$.

To summarize, for the asymmetric matching pennies game of figure 1, any of the reduced form R-QRE can be obtained from *any* admissible i.i.d. structural QRE by choosing the error rates appropriately. In particular, the same set of R-QRE can be obtained by assuming a Logit or Probit model with different error rates for the two players.

Whether this equivalence is special to the 2×2 case is not immediately clear and it is an interesting open question whether the set of R-QRE can be obtained from a structural model with additive disturbances more generally. If so, then in terms of the *equilibria* that can be generated, the reduced form approach would be essentially equivalent to the structural approach with regularity imposed.

7. Conclusion

Models of stochastic choice are widely employed by mathematical psychologists, applied econometricians, and experimental game theorists. They were introduced to the psychology literature by Luce (1959) and others to explain why a subject’s behavior becomes “noisier” as the stimuli become weaker. In this context the stochastic elements represent intra-personal variations in utility levels, or perception errors, which may cause a subject to choose differently when faced with the same stimuli. In the econometrics literature (e.g. McFadden, 1981), stochastic choice models are usually applied to cross sectional data, containing decisions of many individuals. Here the stochastic elements are interpreted as inter-personal variation, or heterogeneity, in preferences. Experimental game theorists are mostly agnostic about the interpretation of the stochastic elements. In the laboratory, noise may be due to either distractions, perception biases, or miscalculations, or due to heterogeneous preference shocks such as feelings of envy, spite, or altruism.

Regardless of the interpretation of the noise, the effect can be particularly important in an interactive context where players’ payoffs are sensitive to others’ decisions. In some versions of the traveler’s dilemma (Capra et al., 1999), for example, small amounts of noise

have a large “snowball” effect when endogenous interactions are considered. McKelvey and Palfrey (1995) introduced the notion of a Quantal Response Equilibrium (QRE) to capture the feedback effects of noisy behavior. The basic idea behind QRE is that players are “better responders” rather than best responders, and they are aware that others are better responders. In other words, QRE imposes a consistency condition on players’ beliefs about others’ noisy behavior.

In past work, the structural QRE approach has been successfully applied to analyze data from a wide variety of laboratory experiments. Almost all of these applications involve Logit probabilistic choice models that result when extreme-value i.i.d. errors are assumed. One advantage of the i.i.d. assumption is that choice probabilities are monotonic, i.e. options with higher expected payoffs are chosen more frequently. This economically sensible criterion may be invalidated when more general error structures are used. Without monotonicity, players are not necessarily “better responders,” suggesting some form of anti-rationality, and *any outcome* can be rationalized as Haile et al. (2004) have pointed out. As we argue in Section 3.2, their critique applies to all (econometric) applications of stochastic choice models, not just QRE, and underscores the importance of imposing restrictions on the error structure.

Even with a parametric specification for the error distribution, Haile et al. (2004) argue that the fit to the data fails to be perfect only to the extent that the specification does not allow for a sufficiently rich family of distributions. As a fix, they advocate using data from alternative environments (treatments), which is done in most of the experimental economics papers they cite where the same error distribution is used for all treatments. They note, however, that this approach would still be “uninformative” if the estimation allows a new distribution for each treatment. The force of this critique depends on the game being considered and the class of distributions one is limited to. For the generalized matching-pennies game in Section 4 with $X = 9$, the set of all possible regular QRE is given by the dark shaded area in the upper left part of figure 2. A second treatment with $X = 0$ produces an analogous rectangle given by the light shaded area in figure 2, with the only point of overlap being the midpoint $(1/2, 1/2)$. This sharp prediction would not depend on an assumption that the error distributions are the same for the two treatments.¹⁴

In this paper, we propose two constructive solutions to the problem signaled by Haile et al. (2004). First, we derive restrictions on the error distribution such that monotonicity is guaranteed. In particular, we show that interchangeability of the additive random payoff disturbances is a sufficient condition. It is important to note that interchangeability is weaker than the commonly employed i.i.d. assumption. Interchangeability of the shocks implies that choice probabilities are label independent, i.e. they are determined by expected payoffs only and not by the labels of decisions. Such label independence seems to be a sensible restriction in the context of abstract (matrix) games.

Second, and more importantly, we introduce a reduced-form approach to QRE. Specifically, we define a regular quantal response equilibrium (R-QRE) as the fixed point of regular quantal response functions, which are smoothed best response functions that satisfy interiority, continuity, responsiveness, and monotonicity. In general R-QRE does impose strong, testable empirical restrictions on data, and these restrictions are consistent with laboratory observations. We believe the reduced-form approach is a simpler, more intuitive

approach to modeling stochastic choice in games. Furthermore, if one imposes regularity, the reduced-form approach is more general than the structural approach and allows for a richer set of parametric models for estimation purposes (e.g. the Luce probabilistic choice model).

Acknowledgment

Financial support from the National Science Foundation NSF (SBR-0098400 and SES-0079301) and the Alfred P. Sloan Foundation is gratefully acknowledged. The paper has benefitted from discussions with Michelle Goeree, Philip Haile, Ali Hortacsu, David Levine, Angela Moore, Brian Rogers and Larry Samuelson. The Haile et al. (2004) paper inspired us to think more deeply about the restrictions implied by QRE models. We also acknowledge helpful comments from seminar participants at the 2003 ESA North American Meeting and the Nottingham workshop on “The Role of Experimental Methods in Economics,” an anonymous referee and the associate editors for this special issue.

Notes

1. A third source of noise is due to differences in individuals' experiences that causes their beliefs to differ. The latter does not come up in individual choice experiments, but may play a role in interactive contexts (see Section 2). Hey (2005, this issue) and Loomes (2005, this issue) explore in more depth these questions about possible sources of noise and how to account for it in the data analysis.
2. Of course, it was known before that one could purify mixing by adding payoff-irrelevant information, such as the private observation of a randomizing device. The innovation here was that the private information actually could change the payoff of the game, and players were using *strictly* optimal strategies.
3. In other words it is a Bayesian-Nash equilibrium of the Bayesian game defined by Γ and f .
4. See, e.g., Anderson, et al. (1992) for a derivation and discussion of these properties.
5. Without loss of generality we will take $\mu_1 = 1$ in the examples to follow.
6. Note that without the mean 0 restriction required by admissibility, it would be trivial to find a distribution of disturbances such that $\text{Prob}[1 + \varepsilon_{11} > 2 + \varepsilon_{12}] = p$ for any $p \in (0, 1)$. Simply let ε_{11} be uniform over $[0, 1]$ and let ε_{12} be degenerate at $\varepsilon_{12} = -p$.
7. The measure for general $X \leq 1$ is $(X - 1)/(4X + 12)$. The upper bound is $1/4$ as X grows large, and the lower bound is 0 when $X = 1$. In other words, for the symmetric matching pennies game any regular QRE predicts that both players randomize uniformly.
8. Without full support, the inequalities in conditions (A1) and (A3) hold only weakly. For example, interiority and *strict* responsiveness are violated with uniformly distributed disturbances.
9. We are grateful to David Levine for pointing out that the conventional terminology for label independence is interchangeability.
10. Indeed, in the structural empirical Industrial Organization literature, Symmetry and Strong Substitutability are often seen as weaknesses of the Logit model because of the implied substitution patterns. One solution in this literature is to define demand over product attributes (rather than products) and to allow for interaction effects between consumer demographics and product attributes. This solution is not applicable in the context of abstract (matrix) games studied here.
11. The choice probabilities in (6.2) also satisfy the IIA property, which may be undesirable in some contexts.
12. Our terminology of reduced-form versus structural quantal response functions follows the econometrics literature where the reduced-form approach typically involves the specification of a demand function, while in the structural approach this demand function is derived from optimal consumer choice behavior based on underlying (parametric) utility functions.
13. Recall that the inequalities used to construct the shaded area were derived from Monotonicity arguments.

14. See Chapter 11 in Holt (2004) for data from two treatments of a matching pennies game with reflected rectangular prediction areas.

References

- Anderson, S.P., Jacob, K.G., and Charles, A.H. (1998). "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction." *Journal of Political Economy*. 106(4), 828–853.
- Anderson, S.P., Jacob K.G., and Charles, A.H. (2001). "Minimum-Effort Coordination Games: Stochastic Potential and Logit Equilibrium." *Games and Economic Behavior*. 34(2), 177–199.
- Anderson, S.P., Andre de Palma, and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.
- Camerer, C.F., H. T.-H., and Chong, J.-K. (2002). "Sophisticated Learning and Strategic Teaching." *Journal of Economic Theory*. 104, 137–188.
- Capra, C.M., Jacob, K.G., Rosario, G., and Charles, A.H. (1999). "Anomalous Behavior in a Traveler's Dilemma." *American Economic Review*. 89(3): 678–690.
- Chong, J.-K., Colin, F.C., and H., T.-H. (2005). "A Learning-based Model of Repeated Games with Incomplete Information." *Games and Economic Behavior*. forthcoming.
- Dvoretzky, A., Abraham, W., and Jacob, W. (1951). "Elimination of Randomization in Certain Statistical Decision Problems and Zero-Sum Two-person Games." *Annals of Mathematical Statistics*. 22, 1–21.
- Goeree, J.K. and Charles, A.H. (2001). "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." *American Economic Review*. 91(5), 1402–1422.
- Goeree, J.K. and Charles, A.H. (forthcoming) "An Explanation of Anomalous Behavior in Models of Political Participation." *American Political Science Review*.
- Goeree, J.K., Charles, A.H., and Thomas, R.P. (2002). "Quantal Response Equilibrium and Overbidding in First-Price Auctions." *Journal of Economic Theory*. 104, 247–272.
- Goeree, J.K., Charles, A.H., and Thomas, R.P. (2003). "Risk Averse Behavior in Generalized Matching Pennies Games." *Games and Economic Behavior*. 45, 97–113.
- Goeree, J.K., Thomas, R.P., and Brian, W.R. (2005). "Rank Dependent Choice in Games." working paper, California Institute of Technology.
- Haile, P.A., Ali, H., and Grigory, K. (2004). "On the Empirical Content of Quantal Response Models." working paper, University of Wisconsin.
- Harless, D.W. and Colin, F.C. (1994). "The Predictive Utility of Generalized Expected Utility Theories." *Econometrica*. 62:6 (November), 1251–1289.
- Harsanyi, J.A. (1973). "Games with Randomly Distributed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of Game Theory*. 2, 1–23.
- Hey, J.D. and Chris, O. (1994). "Investigating Generalization of Expected Utility Theory Using Experimental Data." *Econometrica*. 62:6 (November), 1291–1326.
- Hey, J.D. (2005). "Why We Should Not Be Silent About Noise." working paper, Universities of York and Bari.
- Holt, C.A. (2004). "Markets, Games, and Strategic Behavior." unpublished manuscript: <http://www.people.virginia.edu/~cah2k/expbooknsf.pdf>
- Karlin, S. (1966). *A First Course in Stochastic Processes*. Academic Press: New York.
- Laury, S.K. and Charles, A.H. (2002). "Risk Aversion and Incentive Effects." *American Economic Review*. December, 93(5).
- Loomes, G. (2005). "Modelling the Stochastic Component of Behaviour: Some Implications and Problems for Testing Models." working paper, University of East Anglia.
- Luce, D. (1959). *Individual Choice Behavior*. New York: Wiley.
- McFadden, D. (1981). "Econometric Models of Probabilistic Choice." in *Structural Analysis of Discrete Data with Econometric Applications*. Manski C. and McFadden D., eds., MIT Press, Cambridge, 198–272.
- McKelvey, R.D. and Thomas, R.P. (1992). "An Experimental Study of the Centipede Game." *Econometrica*. 60(4), 803–836.
- McKelvey, R.D. and Thomas, R.P. (1994). "Quantal Response Equilibrium for Normal Form Games." *Social Science Working Paper # 883*. California Institute of Technology, March.

- McKelvey, R.D. and Thomas, R.P. (1995). "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*. 10(1), 6–38.
- McKelvey, R.D. and Thomas, R.P. (1996). "A Statistical Theory of Equilibrium in Games." *Japanese Economic Review*. 47(2): pp. 186–209.
- McKelvey, R.D. and Thomas, R.P. (1998). "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*. 1(1), 9–41.
- McKelvey, R.D., Thomas, R.P., and Roberto, W. (2000). "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2×2 Games with a Unique Mixed-Strategy Equilibrium." *Journal of Economic Behavior and Organization*. 42, 523–548.
- Myerson, R.B. (1978). "Refinements of the Nash Equilibrium Concept." *International Journal of Game Theory*. 7, 73–80.
- Ochs, J. (1995). "Games with a Unique Mixed-Strategy Equilibrium: An Experimental Study." *Games and Economic Behavior*. 10, 202–217.
- Rosenthal, R.W. (1989). "A Bounded Rationality Approach to the Study of Non-Cooperative Games." *International Journal of Game Theory*. 2, 65–67.
- Selten, R. (1965). "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrage-tragheit." *Zeitschrift für die Gesamte Staatswissenschaft*. 121, 301–324, 667–689.
- Selten, R. (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory*. 4, 25–55.
- von Neumann, J. (1928). "Zur Theorie der Gesellschaftsspiele." *Mathematische Annalen*. 100, 295–320.
- von Neumann, J. and Oskar, M. (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Wald, A. (1945). "Statistical Decision Functions Which Minimize the Maximum Risk." *Annals of Mathematics*. 46, 265–80.