

Quantal Response Equilibrium

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey¹

Motivation

Economic theory relies extensively on the assumption of perfect rationality, which makes it possible to construct general models with strong (and sometimes surprising) predictions. The evaluation of these models using field data requires the incorporation of random errors representing unobserved and omitted elements, measurement error, etc. Evaluation of these models using data from laboratory experiments also requires an error structure, since choice behavior in the laboratory is also noisy, showing clear mistakes and inconsistencies over time.

Probabilistic choice models (e.g., logit, probit) have long been used to incorporate stochastic elements into the analysis of individual decisions, and the quantal response equilibrium (QRE) is the analogous way to model games with noisy players. These probabilistic choice models are based on quantal response functions, which have the intuitive feature that deviations from optimal decisions are negatively correlated with the associated costs. That is, individuals are more likely to select better choices than worse choices, but do not necessarily succeed in selecting the very best choice. Formally, a quantal response function maps the vector of expected payoffs from available choices into a vector of choice probabilities that is monotone in the expected payoffs.

In a strategic game environment, a player's expected payoffs from different strategies are determined by beliefs about other players' actions, so beliefs determine expected payoffs, which in turn, generate choice probabilities according to some quantal response function. A quantal response equilibrium (QRE) imposes the requirement that the beliefs match the equilibrium choice probabilities. Thus, QRE requires solving for a fixed point in the choice probabilities, analogous to the Nash equilibrium.

¹ Goeree and Palfrey: Division of the Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA. Holt: Department of Economics, PO Box 400128, University of Virginia, Charlottesville, VA 22904-4182, USA. We acknowledge financial support from the Alfred P. Sloan Foundation, the National Science Foundation (SBR 0094800 and 0551014; SES 0450712 and 0214013), and the Dutch National Science Foundation (VICI 453.03.606).

In fact, QRE is a generalization of Nash equilibrium, which converges to the Nash equilibrium as the quantal response functions become very steep, and approximate best response functions. This approach provides a useful theoretical framework for looking at comparative statics effects of parameter changes that may not alter Nash predictions. The incorporation of random elements also provides a foundation for standard statistical analysis of field and laboratory data in game theoretic applications.

A Motivating Example: Generalized Matching Pennies

Before providing general definitions, it is useful to begin with a simple two-person matching pennies game in which the Row player chooses top (T) or bottom (B) and the Column player chooses left (L) or right (R). Row wins a penny (and column loses a penny) if the outcome is (top, right) or (bottom, left) and Column wins a penny otherwise. Thus Row's expected payoff for Top (U_T) is a function of Column's probability of choosing Right (p_R), which is easily calculated as $U_T(p_R) = p_R - (1-p_R) = 2p_R - 1$, and similarly, $U_B(p_R) = 1 - 2p_R$, so the optimal decision is to choose Top if Column is more likely to choose Right, i.e. if $p_R > 1/2$. Column's expected payoffs are computed analogously, as a function of Row's probability of choosing Top (p_T).

Figure 1 illustrates the best response functions in the unit square of mixed strategies in the game, with the y-axis representing the row player's Top choice probability and the x-axis representing the column player's Right choice. The best response for Row is indicated by the dark step function from that crosses from the bottom to the top at $1/2$. Similarly, the Column player's best response line is the step function, shown in gray, which crosses over from left to right at a height of $1/2$.

Using the same figure, we can represent a quantal response function, which smooths out the discontinuous best response function, reflecting the monotone but stochastic choice as a function of payoffs. Such a quantal response function is illustrated by Row's dark curved line that rises smoothly from the bottom left corner to the top-right corner. The probabilistic choice equals $1/2$ exactly at the point where row player is indifferent between Top and Bottom. A quantal response function is also drawn for Column. The intersection of these two quantal response functions occurs in the center of the figure, and is the quantal response equilibrium, just as the intersection of the sharp

best response function at the same point is the Nash equilibrium in mixed strategies ($1/2$ for each decision).

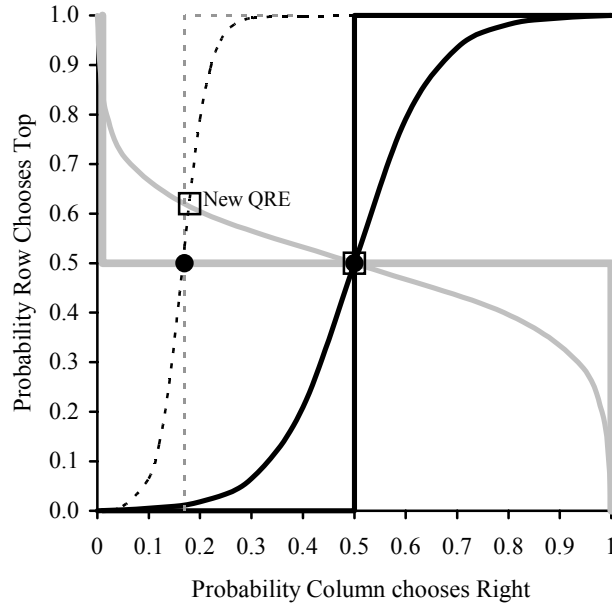


Figure 1. Players' Best Responses and Quantal Responses For a Generalized Matching Pennies Game

Now suppose that all payoffs stay the same except for the Top-Right outcome, which gives Row a higher payoff of 9 and Column -1 as before. The increase in Row's Top-Right payoff shifts Row's Best response line upward, as indicated by the dotted line step function in the figure, and it also shifts Row's quantal response (smooth dotted line). The new Nash equilibrium (dot at the intersection of the step functions) is at $p_T = 0.5$, whereas the new QRE is at a higher level of the probability of Top. 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). 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, Palfrey, and Weber (2000), Goeree and Holt (2001), and Goeree, Holt, and Palfrey (2003).

Definitions

Let $G = (N, \{S_1, \dots, S_n\}, \pi_1, \dots, \pi_n)$ be a normal-form game, where $N = \{1, \dots, n\}$ is the set of players, $S_i = \{s_{i1}, \dots, s_{iJ(i)}\}$ is player i 's set of strategies and $S = S_1 \times \dots \times S_n$ is the set of strategy profiles, and $\pi_i : S_i \rightarrow \mathbb{R}$ is player i 's payoff function. Furthermore, let $\Sigma_i = \Delta^{J(i)}$ 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 $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ 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 $p(s) = \prod_{i \in N} \sigma_i(s_i)$ is the probability distribution over pure-strategy profiles induced by σ .

Let P_{ij} denote the probability that player i selects strategy j . Recall that the main idea behind QRE is that strategies with higher expected payoffs are more likely to be chosen, although the best strategy is not necessarily chosen with probability 1. In other words, QRE replaces players' strict rational choice best-responses by *smoothed best responses* or *quantal responses*.

Definition 1: $P_i : R^{J(i)} \rightarrow \Delta^{J(i)}$ is a **regular quantal-response function** if it satisfies the following four axioms.

- **Interiority:** $P_{ij}(\pi_i) > 0$ for all $j = 1, \dots, J(i)$ and for all $\pi_i \in R^{J(i)}$.
- **Continuity:** $P_{ij}(\pi_i)$ is a continuously differentiable function for all $\pi_i \in R^{J(i)}$.
- **Responsiveness:** $\partial P_{ij}(\pi_i) / \partial \pi_{ij} > 0$ for all $j = 1, \dots, J(i)$ and for all $\pi_i \in R^{J(i)}$.
- **Monotonicity:** $\pi_{ij} > \pi_{ik}$ implies that $P_{ij}(\pi_i) > P_{ik}(\pi_i)$ for all $j, k = 1, \dots, J(i)$.

These axioms are economically and intuitively compelling. Interiority ensures the model has full domain, i.e. is logically consistent with all possible data sets. This is important for empirical applications 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.

Define $P(\pi) = (P_1(\pi_1), \dots, P_n(\pi_n))$ to be regular if each P_i satisfies the above regularity axioms. 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 2: *Let P be regular. A **Regular Quantal Response Equilibrium (R-QRE)** of the normal-form game G is a mixed-strategy profile σ^* such that $\sigma^* = P(\sigma^*)$.*

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

Theorem: *There exists a Regular Quantal Response Equilibrium (R-QRE) of G for any regular P .*

Empirical Implications of Regular QRE

The axioms underlying regular QRE collectively have strong empirical implications, even without any parametric assumptions on P . To illustrate the nature of these restrictions, consider again the generalized matching-pennies game, where Row's payoff is X when the outcome is (top, right). If $X > 1$, it is readily verified that Row's expected payoff of choosing "top" is higher than of choosing "bottom" when $p_R < 2/(X+3)$ ($p_R > 2/(X+3)$). Monotonicity therefore implies that, if (p_R^*, p_T^*) defines a regular QRE, it must satisfy the inequalities: $p_T^* \geq 1/2$ if $p_R \geq 2/(X+3)$ and vice versa. Likewise, Column's expected payoff of choosing "right" is higher (lower) than of choosing "left" when $p_T < 1/2$ ($p_T > 1/2$). Thus (p_R^*, p_T^*) must satisfy $p_R^* \leq 1/2$ if $p_T \geq 1/2$, and vice versa. The region defined by these inequalities defines the set of possible regular QRE. For the specific case of $X = 9$, this area is shown by the dark gray shaded area in Figure 2. The three black dots show the Nash equilibria for $X = 9$ (left), $X = 1$ (center) and $X = 0$ (right).




Figure 2. QRE Sets for Generalized Matching Pennies with $X = 9$ (dark) and $X = 0$ (light)

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 "top" more often than "bottom" in an R-QRE when $X > 1$, while the reverse is true for $X < 1$. In fact, the responsiveness axiom can be used to show that if Row's payoff of the (top, right) outcome rises, Row's probability of choosing "top" increases.

Proposition (Goeree, Holt, and Palfrey, 2005). *In any regular QRE of the asymmetric matching pennies game, Row's probability of choosing "top" is strictly increasing in X and Column's probability of choosing "right" is strictly decreasing in X .*

Quantal Response Equilibrium: A Structural Definition

The original definition of QRE (McKelvey and Palfrey 1995) adopts an approach in the spirit of Harsanyi (1973) and McFadden (1976) whereby the choice probabilities are rationalized by privately observed, mean zero random disturbances to the expected payoffs. These disturbances are assumed to be private information to the players, thereby converting the original game into special kind of game of incomplete information. Any

Bayesian equilibrium of this disturbed game is a QRE of the underlying game. The quantal response function generating the QRE is determined by the probability distribution of the random payoff disturbances.

Thus a smoothed response line can be interpreted to be the (inverse) distribution function of the differences between the disturbances, which has a value of $\frac{1}{2}$ when the expected payoffs are equal. For example, if the disturbances are iid and normally distributed, then the quantal response functions will take the shape of a “probit” curve, while if the iid disturbances are distributed according to an extreme value distribution the quantal response functions will have a logistic form. For example, the logit QRE for the generalized matching pennies game is a pair of probabilities that solve:

$$p_T = \frac{\exp(\lambda[(X+1)p_R - 1])}{\exp(\lambda[(X+1)p_R - 1]) + \exp(\lambda[1 - 2p_R])}, \quad p_R = \frac{\exp(\lambda[1 - 2p_T])}{\exp(\lambda[1 - 2p_T]) + \exp(\lambda[2p_T - 1])}$$

where the numerators are exponential functions of the expected payoffs for the corresponding decision (T or R), and the denominators are normalizing factors that force probabilities to sum to 1. As the logit precision parameter λ increases, the response functions become more responsive to payoff differences, and the logit response functions converge to the sharp step functions shown in Figure 1.

The disturbances in the structural approach to QRE can be interpreted in several ways. One possibility is to interpret them literally. That is, one views the underlying game as simply a model of the average game being played, with each actual game player being a mean zero perturbation of the basic game. With this view, one can think of the payoff disturbances as reflecting the effects of unobservable components such as a player’s mood, perceptual variations, etc. A second possibility is to think of the players as statisticians, whose objective is to estimate the payoff of each strategy using some unknown set of instruments to perform the estimation. For general abstract games, a reasonable first cut is to suppose that their estimation errors are unbiased. The players then choose the strategy with the highest estimated expected payoffs, implicitly taking into account that the other players are also estimating payoffs with some error, with the resulting equilibrium corresponding to QRE.

One can show that the quantal response function generated by iid disturbances will always have the continuity and monotonicity properties of regular quantal response functions, and therefore will lead to regular QRE. In particular, the comparative static result of the previous section holds for the logit QRE (McKelvey and Palfrey, 1996). If disturbances are not iid, however, non-monotonicities are possible. Haile et al. (2004) use this observation to show that without restrictions on the disturbances, structural QRE can explain *any* data. One way to avoid this problem is to make the iid assumption or impose the weaker notion of *interchangeability* (Goeree, Holt, and Palfrey, 2005). A second way is to generate restrictions is to constrain the same structural assumptions to hold across different data sets, thereby generating comparative static predictions. Another solution is to simply work directly with the regularity axioms of Definition 1 – the resulting quantal response functions do impose empirical restrictions on the data and do not inherit the unintuitive features of the structural approach such as symmetry and strong substitutability.²

Applications: Quantal Response Equilibrium in Normal-Form Games

In an individual choice problem, the addition of “noise” spreads out the distribution of decisions around the expected-payoff-maximizing decision. In contrast, expected payoffs in a game depend on other players’ choice probabilities, and this interactive element can magnify the effects of noise via feedback effects. One notable example is a coordination game where each person’s payoff is the minimum of all player’s efforts, minus a cost parameter, c , times a player’s own effort. If $c < 1$, any common effort level is a Nash equilibrium, since a unilateral decrease below a common effort reduces the minimum and saves the cost, for a net loss of $1-c$ for each unit reduction in effort. Conversely, a unilateral effort increase does not alter the minimum, so the loss is $-c$. Therefore c affects the downward slopes of the expected payoff function in each direction, and if there is any uncertainty about others’ decisions, low values of c make effort increases less risky. It is not surprising that reductions in the

² *Symmetry* requires 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 . *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.

effort cost c tend to increase average efforts, both in laboratory experiments with human subjects and in a quantal response equilibrium with noisy behavior (Anderson, Goeree, and Holt, 2001).³

The traveler's dilemma is another game where small amounts of noise can have large effects. As in the coordination game, the payoffs depend on the minimum of all decisions ("claims"). This is a two-person "lost luggage" problem, where the airline representative interprets unequal damage claims as evidence that the high claimant is inflated unjustly. Each player earns the minimum of the two claims, with a penalty of R subtracted from the payoff for the high claimant and added to that of the low claimant. As with the coordination game, claims must be in a specified interval, but in the Traveler's Dilemma the unique Nash equilibrium is the lowest possible claim in this interval, irrespective of the magnitude of R , as long as the benefit R from a reduction below a common claim is greater than the smallest permitted claim reduction. In contrast, intuition suggests that claims will be high when the penalty from having the higher claim is low. In the Capra et al. (1999) experiment, reductions in the penalty parameter, R , induce dramatic increases in claims, moving the average from near Nash levels for high values of R to the opposite side of the range of feasible claims for low values of R . This strong treatment effect is tracked well by the quantal response equilibrium with the same precision parameter used to track the coordination game data.

In addition to these applications, the QRE has been used to explain "anomalous" behavior in a wide variety of games, including signaling games, centipede games, two-stage bargaining, and overbidding in auctions (Goeree, Holt, and Palfrey, 2002). In addition, the quantal response equilibrium has proven to be quite useful in the analysis of data from political science experiments: jury voting (Guarnaschelli, McKelvey, and Palfrey, 2000), voter turnout (Levine and Palfrey, 2005), and behavior in participation games (Goeree and Holt, 2005; Cason and Van Lam, 2005).

³ For the two-person coordination game experiments reported in Goeree and Holt (2005), a reduction in the effort cost parameter from 0.75 to 0.25 raised average effort levels from 126 to 159 in the final 5 rounds. With payoffs in pennies and a precision parameter of 0.1, the logit QRE predictions for this game are 126 and 154 for the high and low effort costs. Thus the QRE tracks the strong behavioral response to the treatment change, whereas the range of Nash equilibria (from 110 to 170) is unaffected by this change.

Applications: Quantal Response Equilibrium in Extensive Form Games

The QRE approach has also been developed for extensive form games (McKelvey and Palfrey 1998), where the analysis is done using behavioral strategies. In the extensive form QRE, players follow Bayes' rule and calculate expected continuation payoffs based on the QRE strategies of the other players. Interiority implies that over the nodes at every information set are uniquely defined for any QRE strategy profile. Therefore issues related to belief-based refinements do not arise and a quantal response version of sequential rationality follows immediately. When quantal response functions approach best response functions, then the limiting QRE of the extensive form game will select a subset of the *sequential equilibria* of the underlying game.

QRE in extensive form games will typically have different implied choice probabilities than would obtain if the same quantal response function were applied to the same game in its reduced normal form. This occurs for two reasons. First, QRE is not immune to "reduction" of equivalent strategies, since duplicate strategies will generally change the quantal response choice probabilities, for much the same reason as the "red bus – blue bus" example in discrete choice econometrics. Second, expected payoff differences are different when one collapses an extensive form game into its normal form: with behavior strategies, expected payoffs are computed at the interim stage, conditioning on previous actions in the game; in contrast, normal form mixed strategies are calculated *ex ante*.

Summary

The quantal response approach to the analysis of games has proven to be a useful generalization of the Nash equilibrium, especially when dealing with "noisy decisions" made by boundedly-rational players and by subjects in experiments. It can be extended to allow for learning and cognitive belief formation in one-shot games where learning is not possible. This approach provides a coherent framework for analyzing an otherwise bewildering array of "biases" and anomalies in economics.

References

- Anderson, S.K., J.K. Goeree, and C.A. Holt (2001) "Minimum Effort Coordination Games: Stochastic Potential and the Logit Equilibrium," *Games and Economic Behavior*, 34, 177-199.
- Capra, C.M, J.K. Goeree, R. Gomez, and C.A. Holt (1999) "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89, 678-690.
- Cason, T.N. and V.L. Mui (2005) "Uncertainty and Resistance to Reform in Laboratory Participation Games," *European Journal of Political Economy*, 21(3), 708-737.
- Goeree, J.K. and C.A. Holt (2005) "An Experimental Study of Costly Coordination," *Games and Economic Behavior*, 46(2), 281-294.
- Goeree, J.K. and C.A. Holt (2005b) "An Explanation of Anomalous Behavior in Models of Political Participation," *American Political Science Review*, 99(2), 201-213.
- Goeree, J.K., C.A. Holt and T.R. Palfrey (2002) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), 247-272.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey (2003) "Risk Averse Behavior in Asymmetric Matching Pennies Games," *Games and Economic Behavior*, 45, 97-113.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey (2005) "Regular Quantal Response Equilibrium," *Experimental Economics*, 8(4), 347-67.
- Guarnaschelli, S., R.D. McKelvey, and T. R. Palfrey (2000) "An Experimental Study of Jury Decision Rules," *American Political Science Review*, 94, 407-423.
- Haile, P., A. Hortacsu, and G. Kosenok (2004) "On the Empirical Content of Quantal Response Equilibrium," working paper, Yale University.
- Harsanyi, J. (1973) "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium," *International Journal of Game Theory*, 2, 1-23.
- Levine, D. and T.R. Palfrey (2005) "A Laboratory Test of the Rational Choice Theory of Voter Turnout," forthcoming *American Political Science Review*.
- McFadden, D. (1974) "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka (ed.) *Frontiers in Econometrics*, New York: Academic Press.
- McKelvey, R. D. and T. R. Palfrey R. (1995) "Quantal Response Equilibrium for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- McKelvey, R. D. and T. R. Palfrey (1996) "A Statistical Theory of Equilibrium in Games," *Japanese Economic Review*, 47(2), 186-209.
- McKelvey, R. D. and T. R. Palfrey (1998) "Quantal Response Equilibrium for Extensive Form Games," *Experimental Economics*, 1(1), 9-41.
- Ochs, J. (1994) "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.