Economics Bulletin

Volume 32, Issue 1

On mu from the logistic quantal-response equilibrium

Robert Ferenc Veszteg
Waseda University

Abstract

This note reflects on the key parameter of the popular logistic quantal-response equilibrium in order to set some common guidelines on its empirical interpretation. It is stressed that the estimated model must be in harmony with the experimental design, because the estimation results on mu prove to be sensitive to changes in the strategy sets of players even if those are unimportant from a game-theoretic point of view. It is also shown that a simple post-estimation correction of mu can help inter-game comparisons, while pre-estimation treatments of the data may introduce unwanted biases.

The author thanks Jack Ochs for providing his database.
Citation: Robert Ferenc Veszteg, (2012) "On mu from the logistic quantal-response equilibrium", *Economics Bulletin*, Vol. 32 No. 1 pp. 102-111.
Contact: Robert Ferenc Veszteg - rveszteg@gmail.com.
Submitted: August 04, 2011. **Published:** January 13, 2012.

1 Introduction

In spite of the interpretational problems of its key parameter μ, the logistic quantal-response equilibrium (QRE) has been enjoying an increasing popularity in the experimental literature. It has appeared in 168 research papers in the leading journals of the field over the past 13 years (table 1), and has been implemented in the game-theoretic analysis software Gambit (McKelvey et al., 2007).

Table 1: Number of articles published in scientific journals in which the concept of "quantal response equilibrium" appears. Sources: ScienceDirect and SpringerLink.

year	ScienceDirect	SprigerLink	sum
1995	1	0	1
1996	1	1	2
1997	3	0	3
1998	8	2	10
1999	5	0	5
2000	6	1	7
2001	3	0	3
2002	18	0	18
2003	11	0	11
2004	12	1	13
2005	16	7	23
2006	14	3	17
2007	15	7	22
2008	27	5	32
1995 - 2008	140	28	168

The goal of this note is to reflect on μ and to set some guidelines for its use with experimental data. First, as any statistical estimate, μ also should always be reported with its standard error and/or its statistical significance should be tested. Second, the QRE estimation result can only be assessed in light of a corresponding Nash equilibrium, as the distance between the observed behavior and the predicted Nash behavior. Third, in order to assess the goodness-of-fit of the QRE model, to interpret its parameter and/or to compare it across different games, μ should be corrected and normalized.

The notion of the quantal response equilibrium has been introduced to the experimental literature by McKelvey and Palfrey (1995) as a possible generalization of the Nash equilibrium. In a quantal-choice model, players make unsystematic mistakes when computing the expected utility attached to their actions and therefore play some noisy best response. Just like the Nash equilibrium, the QRE then is a fixed point, but it is based on these noisy best responses.

In the most-frequently estimated logistic version of QRE, the individual mistakes are modeled with the log-Weibull distribution with a standard deviation proportional to μ.[1] Therefore, the unique μ parameter of the model can be interpreted as a proxy of the players' rationality, i.e. the distance of the observed behavior from the fully-rational behavior assumed by a Nash equilibrium. If μ is above any limit, the considered expected utilities are extremely noisy and

[1] Some authors prefer to characterize QRE by a λ parameter that represents the inverse of μ, i.e. $\lambda = \frac{1}{\mu}$.

players choose their actions practically without looking at them. They randomize among the available actions with equal probability. While as μ approaches 0, players behave more and more in line with the Nash prediction and play best response in the limit, when μ is equal to 0.

Although the two extreme values of μ allow for straightforward interpretations, it is unclear how to interpret the values in between. The literature has fitted the QRE model to numerous data sets on various games and experimental settings. The estimates vary greatly. As μ seems to be game-dependent, the questions of how big is big and how small is small remain open.

Motivated by the increasing popularity of the QRE model, Haile et al. (2008) analyze the "empirical content" of the model and argue that the QRE's widely-documented ability to explain experimental data lies in that it can rationalize a large set of behavior patterns in normal form games. The authors call for additional assumptions on the error term, like the ones behind the logistic version of the model studied here, and for presenting the set of possible outcomes for any analyzed game in order to give a clearer interpretation of the experimental evidence. In their answer to the received criticism, Goeree et al. (2005) list such necessary conditions and also propose a regular quantal-response equilibrium (RQRE) whose smoothed best-response functions have enough structure to provide testable predictions. They claim that the imposed restriction are consistent with previous laboratory observations. This note does not pretend to contribute to the outlined discussion on the QRE model. It considers its most-popular restricted logistic version and raises practical questions related to the interpretation of its parameter.

2 Logistic QRE

Let $u_{ij}(\sigma_{-i})$ denote player i's expected utility when she plays action j against the other player strategy σ_{-i}. In the quantal response model, when choosing their strategies players consider

$$\hat{u}_{ij}(\sigma_{-i}) = u_{ij}(\sigma_{-i}) + \varepsilon_{ij} \tag{1}$$

instead that incorporates some unsystematic error represented by ε_{ij}. Similarly to the Nash equilibrium, the QRE is a fixed point of the best response function, but it is based on the noisy expected utilities.

In the logistic QRE model (McKelvey and Palfrey, 1995), the error terms ε_{ij} are assumed to have independent log-Weibull distribution with a standard error proportional to μ. For a given μ, the players' equilibrium strategies (marked with a star) are computed according to

$$\sigma^*_{ij}(u^*_i) = \frac{\exp(u^*_{ij}/\mu)}{\sum_{k=1}^{n} \exp(u^*_{ik}/\mu)}, \tag{2}$$

where σ^*_{ij} is the probability assigned to action j by player i in the equilibrium, u^*_i is the vector of the u^*_{ij}s that denote player i's expected utility when she plays action j in a situation where the other player plays σ^*_{-i}. The above expression is typically used to build the likelihood function for the estimation of μ from experimental data.

3 Interpreting μ

3.1 The games

In this section, the database from Ochs (1995) is used to illustrate the difficulty of interpreting the μ estimates and to give some practical advice. The data was generated by observing subjects in the laboratory who played three different two-player simultaneous-move games whose main characteristics is a unique Nash equilibrium in mixed strategies. Given that for Ochs' game 1 the Nash equilibrium coincides with random play and therefore with the QRE for any μ, here only games 2 and 3 are analyzed. The upper part of figure 1 shows the strategic form definition of the two games. McKelvey and Palfrey (1995) use the same database to present the model and the predictive power of the QRE.

Since μ is payoff-dependent and the estimation result depends on the unit in which (expected) utility is measured, McKelvey and Palfrey (1995) propose to study a payoff-distorted version of the games. They estimate μ using the same database, but they use the payoff matrices displayed in the lower part of figure 1 in constructing the QRE model.

On one hand, they take into account that Ochs paid his subjects by lotteries (in order to balance expected payoffs between the row- and the column-players). On the other hand, they transform the original payoff matrix into 1982 pennies to have a common currency among all analyzed games. The use of 1982 pennies as a common currency for all the analyzed games is only justified by a failed attempt to find consistency in the parameter estimates across several games. While it is plausible to assume that subjects with different conversion rules face different payoffs, it is unclear why subjects would consider the 1982 equivalent of their cash amount in making their decisions. Money is assumed to give the ultimate incentive in laboratory experiments related to economics, however the decision problem is presented and the decisions are made in experimental monetary units. Therefore participants are more likely to think and make considerations (in particular, to compute expected payoffs) in experimental monetary units than in dollars.

Figure 1: Games 2 and 3 defined by Ochs (1995), and their payoff-distorted versions by McKelvey and Palfrey (1995).

	game 2			game 3	
Ochs		*L*	*R*	*L*	*R*
	T	9, 0	0, 1	4, 0	0, 1
	B	0, 1	1, 0	0, 1	1, 0
McK-P		*L*	*R*	*L*	*R*
	T	1.1144, 0.0000	0.0000, 1.1141	1.1144, 0.0000	0.0000, 1.1141
	B	0.0000, 1.1141	0.1238, 0.0000	0.0000, 1.1144	0.2785, 0.0000

The main purpose of looking at these games through experimental data was to analyze mixed-strategy Nash equilibria. The unique Nash equilibrium is $(\sigma_r^*, \sigma_c^*) = \left(\{\frac{1}{2}, \frac{1}{2}\}, \{\frac{1}{10}, \frac{9}{10}\}\right)$ in game 2, and $(\sigma_r^*, \sigma_c^*) = \left(\{\frac{1}{2}, \frac{1}{2}\}, \{\frac{1}{5}, \frac{4}{5}\}\right)$ in game 3, where σ_r^* denotes the row player's equilibrium strategy and σ_c^* the column player's.

According to the original experimental design (Ochs, 1995), in each round participants had to decide how many times to play the action T (row players) or L (column players) in 10 consecutive hypothetical repetitions of the game. Then the computer prepared a randomized list of actions, matched the players and computed the payoffs accordingly.

The complication arises, because this design introduces a super-game in which players' action sets are different from the ones defined in figure 1. In each round, both the row- and the column-player subject had to choose from the set $\{0,1,2,3,4,\ 5,6,7,8,9,10\}$ instead of $\{T,B\}$ or $\{L,R\}$, respectively. From a game-theoretic point of view the two games are equivalent, as it is easy to find an isomophism that matches the (equilibrium) strategies of the original and the ones of the super-game.[2] Nevertheless, the QRE is not immune to the above change in the experimental design. To show this in a simple way, let us consider the so-called n-full version of the game whose payoff matrix has $(n+1)$ rows and $(n+1)$ columns. In the n-full version of the game, players have to choose from a different action sets. The cell in the ith row and the jth column represents a situation in which the row players chooses T randomly $(n+1-i)$ times and the column player chooses L randomly $(n+1-j)$ times in the n consecutive repetitions of the original game. For example in game 2, the corresponding payoffs therefore are $\frac{1}{n}[4(n+1-i)(n+1-j)+(i-1)(j-1)]$ for the row player and $\frac{1}{n}[(i-1)(n+1-j)+(n+1-i)(j-1)]$ for the column player.

The set of Nash equilibria of the n-full game is much larger than the one of the reduced game, as the former typically contains a continuum of mixed-strategy equilibria. For example, in the 2-full game any profile $(\sigma_r^*,\sigma_c^*) = (\{p, 1-2p, p\},\ \{q_1, q_2, 4q_1+1.5q_2\})$ with $p \in [0,\frac{1}{2}]$ and $q_1, q_2 \in [0,1]$ such that $4q_1+1.5q_2 \leq 1$ represents a Nash equilibrium. The good news is that all of them correspond to the unique mixed-strategy equilibrium of the original game. The QRE represents an equilibrium refinement (McKelvey and Palfrey, 1995), given that it selects a unique strategy profile for any value of μ. In particular, for $\mu = 0$, the QRE of the 2-full game is $\left(\{\frac{1}{3},\frac{1}{3},\frac{1}{3}\},\{\frac{1}{5},0,\frac{4}{5}\}\right)$.[3]

3.2 Payoff and strategy-set dependence

As noted by McKelvey and Palfrey (1995), μ is payoff dependent. If the payoff matrix (including both players' payoffs) is multiplied by a constant, the corresponding estimate of μ changes proportionally, it is multiplied by the same constant. McKelvey and Palfrey (1995) also fit the QRE model to data previously studied by Lieberman (1960) using a distorted payoff matrix. In their attempt to use 1982 pennies as a common currency for all analyzed games, they multiply all the elements in the payoff matrix by 3.373 in this case. This modification renders μ estimates that are 3.373 times larger than the those that can be obtained by considering the original game.

Since the the original QRE model operates with a unique parameter, i.e. the same μ for both players, it is not immune to payoff distortions that affect player asymmetrically. Also, as it is shown below, the QRE estimates are not robust for changes in the strategy space as the one described in the previous section.

Table 2 contains the QRE estimation results for the four versions (two sets of payoffs and two strategy sets) of the two games (games 2 and 3) taking into account different periods in time. Not only the μs change in a non-monotonic way across these versions, but more importantly, they lead to qualitatively different conclusions on the best-performing model among random choice, QRE and Nash equilibrium. Table 3 displays the name of the best-fitting model in each of the studied cases. The model choice is based on the Akaike (AIC) and the Bayesian (BIC) model-selection criteria and the likelihood-ratio test for embedded models.[4] With two

[2]Action i in the super-game correspond to a mixed strategy in the original game according to which the player mixes between her two available action with probabilities $\frac{i}{10}$ and $\frac{10-i}{10}$, respectively. Mixed strategies of the super-game can also be transformed into mixed-strategies of the original game.

[3]It is true that the unique mixed-strategy Nash equilibrium of the original game coincides with the simplified version of the strategy profile selected by the QRE for $\mu = 0$ in the n-full game. Note that, in the QRE of the n-full game for $\mu = 0$, action T of the reduced game is played $\sum_{k=0}^{n} k\frac{1}{n} = \frac{1}{n}\sum_{k=0}^{n} k = \frac{n}{2}$ times over the n repetitions of the game. Normalized to one game it is precisely $\frac{1}{2}$.

[4]All the comparisons are based on the log-likelihood ($\ln\mathscr{L}$) of the estimation results. AIC is computed as $2k-2\ln\mathscr{L}$, where k is the number of estimated parameters, while BIC is defined as $k\ln n - 2\ln\mathscr{L}$, where n is the

exception they pick the same competing model.

Recall that the full games with the original payoff table (Ochs) describe best the situation that subjects faced in the experimental laboratory. It seems that the strategies observed in the laboratory during the early periods are best described as random. Then periods characterized by noisy best-response (QRE) follow. Observed behavior in game 3 converges to the Nash equilibrium in the latest rounds, while in game 2 it does not. It is important to note that the distortions of the payoff matrix and the strategy sets in McKelvey and Palfrey (1995) tend to favor the QRE model.

A simple Monte Carlo experiment shows that switching between the reduced and the *n*-full game does result in different estimates for μ, and there does not exist a regular pattern or systematic bias. In any case, these results suggest caution and that the strategy set implied by the experimental instructions should be used in the QRE estimation process.

Monte Carlo experiment 1 *In this experiment, players are asked in each of a total of 10 rounds to imagine 2 hypothetical repetitions of game 3 and to decide how often she would like to play the action T. For technical reasons, in order to be able to study all possible strategy profiles, the reduced game is studied along with the 2-full game. Figure 2 plots the estimation results from the reduced game as a function of the 2-full-game estimates. It includes 4,356 points that show some positive correlation, but no clear pattern or bias other than some white areas due to lack of precision in the estimation process. Also, given that the optimal μ has been estimated by selecting the value that belongs to the highest log-likelihood on a grid with 360 evaluation points, the points in the figure tend to stay in regular rows and columns.*

3.3 Normalization

The model behind the QRE assumes that players are rational in the sense that they try their best to play some sort of best response to the opponent's behavior, but they may make some unsystematic errors when computing the expected payoffs that their decisions are based on (equation 1). If we were to observe these subjectively-computed expected payoffs, the goodness-of-fit of the QRE model could be assessed by comparing their variance to μ^2. This would shed light on what proportion of the variation in the subjective expected payoffs is due to errors, i.e. the variance of the error term.[5] Unfortunately this normalization is almost never feasible in practice, therefore we must rely on proxies.

Let u_i^{min*} denote player i's smallest possible payoff when she plays a best-response strategy. Similarly, let u_i^{max*} be her maximum best-response payoff. Formally, these are $u_i^{min*} = \min_{\sigma_{-i}} \max_{\sigma_j} u_i(\sigma_i, \sigma_{-i})$ and $u_i^{max*} = \max_{\sigma_{-i}} \max_{\sigma_j} u_i(\sigma_i, \sigma_{-i})$, where $u_i(\sigma_i, \sigma_{-i})$ is player i's expected payoff when using strategy σ_i against the opponent's σ_{-i} strategy. Now we can use the difference $d_i = u_i^{max*} - u_i^{min*}$ as a proxy for the total variation in the subjectively-computed expected payoffs by player i. Note that it is always non-negative and only can be equal to zero in uninteresting degenerate games. Therefore, the average of these differences across players can then be used to normalize μ: $\mu_n = \frac{\mu}{\frac{1}{2}(d_i + d_{-i})}$. In case the payoff matrix (for both players)

number of observations. The model with the smaller AIC and/or BIC value is preferred. The test statistics of the likelihood-ratio test is $-2\frac{\mathscr{L}_0}{\mathscr{L}_1}$, where $\mathscr{L}_0$ is the likelihood of the restricted model and $\mathscr{L}_1$ is the one of the general model. For the nested model, the test statistics follows the χ^2 distribution. The degree of freedom is equal to 1, i.e. the difference in the number of parameters between the two models.

[5]The same philosophy is used to compute relative variances, and it lies also behind the well-known R^2 statistics for various regression models.

Table 2: Likelihood estimates of μ and log-likelihood values corresponding to the QRE model using experimental data gathered by Ochs (1995). μ_n: normalized μ; Random: QRE with $\mu \to \infty$; Nash: QRE with $\mu = 0$. Columns 2 to 6: payoffs as given by Ochs (1995). Columns 7 to 11: payoffs as given by McKelvey and Palfrey (1995). Reduced and full game: as defined in the main text. ***Estimate significantly different from zero at 1%.

periods	μ	μ_n		$\ln \mathscr{L}_{Ochs}$		μ	μ_n		$\ln \mathscr{L}_{McK-P}$	
			QRE	Random	Nash			QRE	Random	Nash
				REDUCED GAME 2						
1−16	19.83	4.41	-1769	-1774	-1938	0.51***	0.49	-1721	-1774	-1938
17−32	0.19***	0.04	-1587	-1774	-1664	0.27***	0.25	-1517	-1774	-1664
33−48	0.10***	0.02	-1703	-1774	-1725	0.29***	0.27	-1605	-1774	-1725
49−56	0.15***	0.03	-767	-887	-792	0.22***	0.21	-743	-887	-792
all	7.70***	1.71	-6103	-6211	-6120	0.31***	0.29	-5612	-6211	-6120
				REDUCED GAME 3						
1−16	8.73	4.36	-1772	-1774	-1821	0.54***	0.55	-1747	-1774	-1822
17−32	3.50***	1.75	-1755	-1774	-1870	0.64***	0.65	-1735	-1774	-1870
33−48	0.26***	0.13	-1661	-1774	-1708	0.31***	0.31	-1640	-1774	-1708
49−64	0.00	0.00	-1678	-1774	-1678	0.00***	0.00	-1679	-1774	-1679
all	4.93***	2.46	-7066	-7098	-7079	0.38***	0.39	-6864	-7098	-7079
				FULL GAME 2						
1−16	79.22	1.84	-613	-614	-684	2.04***	0.26	-601	-614	-913
17−32	0.70***	0.02	-575	-614	-598	1.02***	0.13	-548	-614	-839
33−48	0.40***	0.01	-610	-614	-617	1.12***	0.14	-571	-614	-840
49−56	0.53***	0.01	-279	-307	-286	0.82***	0.10	-270	-307	-413
all	0.51***	0.01	-2147	-2149	-2185	1.20***	0.15	-1997	-2149	-3005
				FULL GAME 3						
1−16	34.62	1.89	-613	-614	-630	2.15***	0.30	-607	-614	-646
17−32	13.93	0.75	-609	-614	-643	2.53***	0.35	-604	-614	-669
33−48	1.04***	0.06	-586	-614	-599	1.18***	0.16	-580	-614	-617
49−64	0.10***	0.01	-592	-614	-592	0.01***	0.00	-574	-614	-630
all	0.48***	0.03	-2451	-2455	-2462	1.51***	0.21	-2397	-2455	-2563

Figure 2: Monte Carlos experiment results. Estimates of μ for the reduced game (horizontal axis) and the n-full game (vertical axis), for $n = 2$.

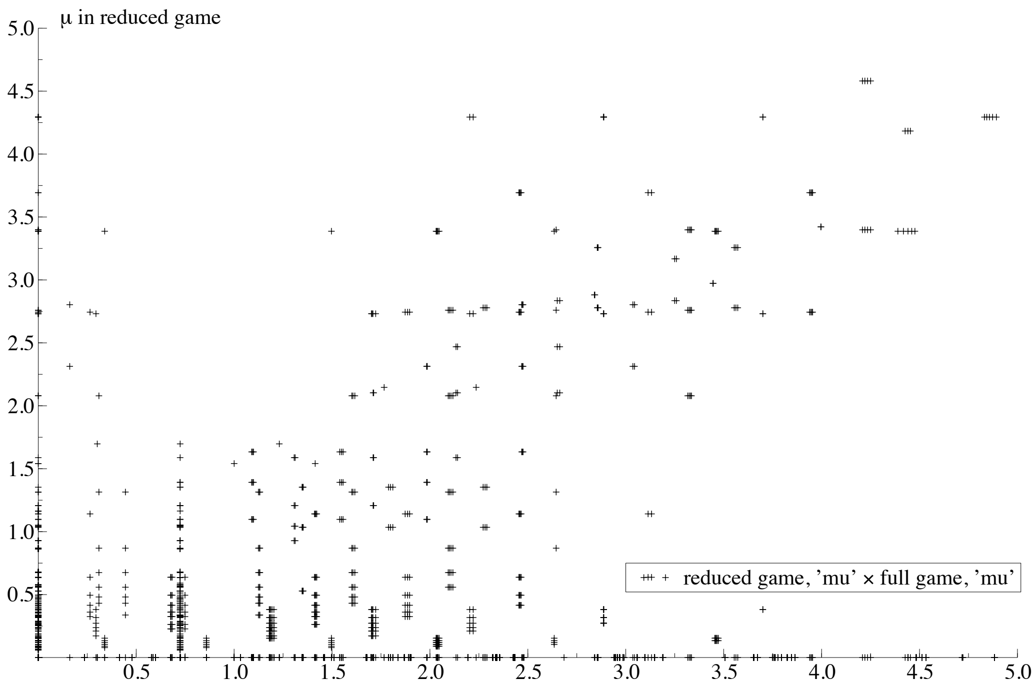

is multiplied by the same constant, like the reconsideration of Lieberman (1960) by McKelvey and Palfrey (1995), both μ and $(d_i + d_{-i})$ change proportionally. Therefore μ_n stays constant and can be compared across different games. However, as the normalized values in table 2 show, μ_n is only robust to small changes in payoffs like the ones between games 2 and 3 (every other characteristic constant). Larger asymmetric changes in payoffs and modifications in the strategy sets can cause large changes even in μ_n.

4 Conclusions

We have seen that μ is game-dependent: changes in the players' payoffs and/or strategy sets may render completely different estimates for μ, even if such changes do not alter incentives and the main game-thoretic characteristics of the conflict.

Therefore, it is important to use that model in the QRE estimation that lies the closest to the experimental design. The normalization technique—of converting experimental monetary units to real money, or finding a common monetary unit for games played in different point in time and/or space—proposed by McKelvey and Palfrey (1995) make questionable adjustments that unnecessarily affects all players in the same way and introduces a bias into the estimation results.

In other words, μ is situation- and game-specific, therefore can not be generalized or interpreted across games even if those are very similar to each other. This also implies that the underlying QRE concept has constrained predictive power, although may successfully be used to identify learning effect in the laboratory.

The normalized μ, i.e. μ_n, proposed in this note allows for further comparisons among

Table 3: The best fitting model to the data gathered by Ochs (1995) based on the likelihood-ratio (LR) test, the AIC and the BIC criterion. Random: QRE with $\mu \to \infty$. Nash: QRE with $\mu = 0$. Columns 2 to 5: payoffs as given by Ochs (1995). Columns 6 to 9: payoffs as given by McKelvey and Palfrey (1995). Reduced and full game: as defined in the main text.

Periods		$\ln \mathscr{L}_{Ochs}$			$\ln \mathscr{L}_{McK-P}$	
	LR	AIC	BIC	LR	AIC	BIC
			REDUCED GAME 2			
1–16	QRE	QRE	QRE	QRE	QRE	QRE
17–32	QRE	QRE	QRE	QRE	QRE	QRE
33–48	QRE	QRE	QRE	QRE	QRE	QRE
49–56	QRE	QRE	QRE	QRE	QRE	QRE
all	QRE	QRE	QRE	QRE	QRE	QRE
			REDUCED GAME 3			
1–16	Random	QRE	Random	QRE	QRE	QRE
17–32	QRE	QRE	QRE	QRE	QRE	QRE
33–48	QRE	QRE	QRE	QRE	QRE	QRE
49–64	Nash	Nash	Nash	Nash	Nash	Nash
all	QRE	QRE	QRE	QRE	QRE	QRE
			FULL GAME 2			
1–16	Random	Random	Random	QRE	QRE	QRE
17–32	QRE	QRE	QRE	QRE	QRE	QRE
33–48	QRE	QRE	QRE	QRE	QRE	QRE
49–56	QRE	QRE	QRE	QRE	QRE	QRE
all	Random	QRE	QRE	QRE	QRE	QRE
			FULL GAME 3			
1–16	Random	Random	Random	QRE	QRE	QRE
17–32	QRE	QRE	QRE	QRE	QRE	QRE
33–48	QRE	QRE	QRE	QRE	QRE	QRE
49–64	Nash	Nash	Nash	QRE	QRE	QRE
all	QRE	QRE	QRE	QRE	QRE	QRE

"similar" games. Nevertheless, the ultimate point of reference for the QRE as an equilibrium concept based on noisy best responses remains the corresponding Nash equilibrium.

References

[1] Goeree, J.K., Holt, C.A. and T.R. Palfrey (2005) "Regular quantal response equilibrium" *Experimental Economics* **8**, 347-367.

[2] Haile, P.A., Hortaçsu, A. and G. Kosenok (2008) "On the empirical content of quantal response equilibrium" *American Economic Review* **98**, 180–200.

[3] Lieberman, B. (1960) "Human behavior in a strictly determined 3 x 3 matrix game" *Behavioral Science* **5**, 317-322.

[4] McKelvey, R.D., McLennan, A.M. and T.L. Turocy (2007) *Gambit: Software tools for game theory*, version 0.2007.01.30, http://www.gambit-project.org.

[5] McKelvey, R.D. and T.R. Palfrey (1995) "Quantal response equilibria for normal form games" *Games and Economic Behavior* **10**, 6-38.

[6] Ochs, J. (1995) "Games with unique mixed strategy equilibria: An experimental study" *Games and Economic Behavior* **10**, 174-189.