

Risk Averse Behavior in Generalized Matching Pennies Games

*Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey**

April 2002

Abstract. In experimental studies of behavior in 2×2 games with unique mixed strategy equilibria, observed choice frequencies are systematically different from mixed-strategy Nash predictions. This paper examines experimental results for a variety of such games, and shows that a structural econometric model which incorporates risk aversion into a quantal response equilibrium explains the data very well. Moreover, risk aversion estimates are stable across the different games and are close to those obtained from laboratory and field auction data, as well as from individual lottery choice experiments.

Keywords: game theory, matching pennies, Nash equilibrium, quantal response equilibrium, risk aversion, prospect theory, experiments.

JEL Codes: C72, C92

I. Introduction

Many strategic situations have the property that one person benefits from a "match," whereas the other benefits from a "mismatch." A soccer goalie prefers to dive right when the penalty kick goes right, and vice versa, whereas the kicker has opposite preferences. In this case, there is clearly no Nash equilibrium in pure strategies, since the goalie would always want to switch away from a mismatch, and the kicker would always want to switch from a match. Similarly, a business manager may hope that an audit occurs only in cases when preparations have been made, but the auditor may prefer that audits catch sloppy record keepers. The soccer example can be seen as approximating a zero-sum game, but most economic applications, such as auditing, are probably better modeled as non-zero sum games.¹ These are generalized matching pennies games.

* Department of Economics, Rouss Hall, University of Virginia, Charlottesville, VA 22901 (Goeree and Holt); Humanities and Social Sciences, California Institute of Technology, 228-77, Pasadena, CA 99125 (Palfrey). This research was funded in part by the National Science Foundation (SBR-0094800 and SES-0079301) and by the John D. Hacker Social Science Experimental Laboratory at Caltech. We are grateful to two anonymous referees for useful suggestions, and to Sean Gailmard, Serena Guarnaschelli, and Vale Murthy for their assistance in conducting the experiments.

¹ Examples also abound in everyday life and even literature, as documented in Brams (1994), who cites passages from the works of William Faulkner, Edgar Allen Poe, and Sir Arthur Conan Doyle.

An unresolved puzzle for those who study strategic behavior is the failure of standard Nash equilibrium predictions to track laboratory data, even in some of the simplest two-person games. In particular, for generalized matching pennies games that have unique equilibria in mixed strategies, experimental evidence indicates that Nash predictions are unreliable (e.g. Ochs, 1995). Moreover, there is no obvious pattern to departures from Nash across the different games within this class.

There are several possible explanations for these findings. An explanation proposed by McKelvey and Palfrey (1995) is that in general we should expect behavior in games to exhibit more randomness than is often implied by the sharp Nash equilibrium predictions. A natural approach is to treat best response functions as being stochastic: actions with higher expected payoffs are more likely to be chosen, although “best” responses are not chosen with probability one. Just as the intersection of deterministic best responses determines a Nash equilibrium, the intersection of these stochastic (quantal) responses determines a "quantal response equilibrium." This equilibrium can produce what appear to be biases in players' behavior and shift choice probabilities in surprising ways (e.g. Goeree and Holt, 1999). For matching pennies games, McKelvey and Palfrey (1995) have shown that the quantal response equilibrium fits the data much better than the standard Nash equilibrium.

However, follow-up experiments (McKelvey, Palfrey, and Weber, 2000; Goeree and Holt, 2001) have revealed discrepancies, which suggest that other factors may play a role. For instance, when individuals are risk averse or have a “social preference” component to their utility function (Bolton, 1998; Camerer, 1997), equilibrium predictions change. Ideally, one would like to run experiments that control for such factors and reward subjects in “utils” instead of dollars. Unfortunately, the lack of reliable methods for inducing risk neutrality or controlling for risk attitudes and social preferences in experiments complicates direct empirical testing of theories based on richer behavioral assumptions. Our approach is to use a structural econometric model to estimate risk attitudes or other preference characteristics.

The quantal response equilibrium concept captures the feedback effects of noisy behavior that are not considered in standard econometric models. Specifically, standard methods involve calculating the Nash equilibrium for a given set of parameters to be estimated and then adding noise to the equilibrium decisions. In contrast, in a quantal response equilibrium, others' noisy

behavior affect a player's own expected payoffs, which may cause equilibrium behavior to differ systematically from Nash predictions. Using quantal response equilibrium as the basic model of the "error" term, we obtain maximum likelihood estimates of risk aversion and probability weighting parameters.

Section II begins with a discussion of experimental results reported by Ochs (1995). These results reveal "own-payoff effects," which constitute the most common pattern of deviations from Nash predictions in matching pennies games when the equilibrium choice probabilities are not exactly 0.5. In these matching pennies games, a change in a player's own payoffs does not affect that player's equilibrium choice probabilities, since the other player has to be kept indifferent. By incorporating some noise into the decision-making process, quantal response equilibrium probabilities will depend on expected payoff differences, and hence, will be sensitive to one's own payoffs. We show that the inclusion of decision error is a promising explanation of own-payoff effects as suggested by Ochs (1995). This explanation is, however, incomplete since the estimated model overpredicts own-payoff effects.

Risk aversion puts curvature into the utility function and hence diminishes the effects of a very high payoff associated with one decision. To evaluate the role of risk aversion, we consider a matching pennies game in which each player has a risky and a safe strategy. For this game, the Nash equilibrium and *any* quantal response equilibrium have the risky strategy chosen more frequently than the safe strategy. The data reported in section III show the opposite pattern. We prove that the inclusion of risk aversion can explain the observed tendency of subjects to choose the safer strategy more often. In section IV we estimate error and risk aversion parameters using the quantal response structural model applied to data from several matching pennies games. The risk aversion estimates are significant and of approximately the same magnitude as estimates recovered from several auction experiments and auction field data. The estimates are further corroborated by lottery choice experiments discussed in section V.

II. Own Payoff Effects in Matching Pennies Games

We focus on the family of 2×2 games which have a unique Nash equilibrium in totally mixed strategies. Consider the matching pennies game in Table I, where the row player chooses between up (U) or down (D), and the column player chooses between left (L) or right (R). The

Table I. Matching Pennies With Own-Payoff Effects (Ochs, 1995)
 Game 1 ($X = 1$), Game 2 ($X = 9$), and Game 3 ($X = 4$)

		Column Player's Decision	
		<i>L</i>	<i>R</i>
Row Player's Decision	<i>U</i>	(X , 0)	(0, 1)
	<i>D</i>	(0, 1)	(1, 0)

value of row's payoff for the up/left outcome, X , determines the degree of "asymmetry." With $X = 1$, laboratory behavior is characterized by probabilities of 0.5 for each decision. This equal-probability outcome corresponds to the unique Nash equilibrium: if one player is choosing with a coin flip, then the other's expected payoff is the average of 0 and 1 for each decision, so a coin flip is as good as any other response. The situation changes if the value of X increases to 9 in Ochs' game 2. The row player's best response is now pulled towards U as shown by the dark line in Figure 1 that has the shape of a "step-function" demand curve. Since column's payoffs are still 0 or 1 for each decision, column's best response crosses over horizontally at a height of 0.5, creating a shape analogous to a single-step supply function. The intersection of these two curves, at $\Pr(R) = 0.9$ and $\Pr(U) = 0.5$, is the unique Nash equilibrium in mixed strategies, labeled N9.

A reduction in X from 9 to 4 reduces the probability of R that makes row indifferent, as shown by the thin vertical line, and the new Nash equilibrium choice probability for column, labeled N4, is $\Pr(R) = 0.8$. Since column's best response function is horizontal, the Nash equilibrium choice probability for row remains unchanged at $\Pr(U) = 0.5$. This is the reason for the absence of own-payoff effects: even though row's large payoff for U is reduced, row's equilibrium probabilities have to remain at 0.5 to keep column indifferent.

Ochs used the three values of X (1, 4, and 9) to construct payoffs in an experiment designed to evaluate these to own-payoff effects. Payoffs were actually in lottery tickets (probabilities) with the prizes chosen to equate expected earnings for row and column players. The sixteen subjects who participated in each game were randomly matched in a series of 56-64

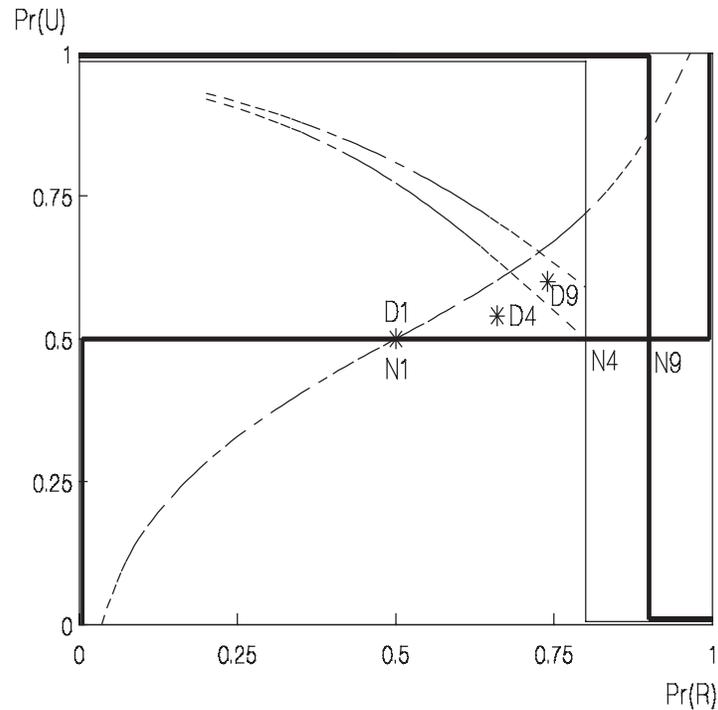


Figure 1. Ochs' (1995) Matching Pennies Games
 Laboratory Data (D1, D4, and D9) and Nash Predictions (N1, N4, and N9)

rounds. Subjects kept the same role, row or column, in all rounds. The aggregate choice percentages for Games 1 to 3 are shown as asterisk points D1, D4, and D9 in Figure 1. As noted by Ochs (1995), the Nash equilibrium prediction of own-payoff invariance is inconsistent with the positive correlation between row's high payoff and the observed probability of choosing up.

McKelvey and Palfrey (1995, 1996) have shown that own-payoff effects are explainable as equilibrium phenomena, by introducing some randomness into the decision-making process. Without such noise, the probability of choosing a decision jumps discontinuously from 0 to 1 as soon as its expected payoff is the highest available. Following Luce (1959), suppose instead that the probability of choosing each decision is a smoothly increasing function of the expected payoff for that decision. Luce provided an axiomatic derivation of the popular "logit" rule, which is

based on exponential functions.² For the column player with expected payoffs $\pi^e(L)$ and $\pi^e(R)$ for decisions L and R respectively, the logit probability of choosing strategy R , say, is given by:

$$p_R = \frac{\exp(\pi^e(R)/\mu)}{\exp(\pi^e(L)/\mu) + \exp(\pi^e(R)/\mu)}, \quad (1)$$

and the probability of choosing L is simply $1 - p_R$. Similarly, row's logit equilibrium probability is determined by:

$$p_U = \frac{\exp(\pi^e(U)/\mu)}{\exp(\pi^e(D)/\mu) + \exp(\pi^e(U)/\mu)}. \quad (2)$$

The denominators in (1) and (2) ensure that probabilities lie between 0 and 1, and μ is a "noise" or "error" parameter. As μ goes to zero, the decision with the highest expected payoff is selected with probability one. Thus higher values of μ correspond to "flatter," i.e. less responsive logit best-response curves. For example, when $\mu = 0$ the best response lines are the solid straight lines in Figure 1. The smooth dashed lines represent the relevant parts of the logit best responses for a higher value of μ .

The intersections of the dashed lines in Figure 1 are *logit equilibria* (McKelvey and Palfrey, 1995) for the Ochs' games for an error parameter estimated from the Ochs data ($\mu = 0.66$). They are equilibria in the sense that each of these probabilities is a logit stochastic best response to the other one. The logit equilibrium is a generalization of the Nash equilibrium, and it converges to a Nash outcome as μ goes to zero (perfect rationality). Since column's stochastic best response function has a positive slope, shifts in row's response function due the payoff changes from 1 to 4 to 9 result in a positive association between row's high payoff and the predicted probability of choosing up. But it appears that the logit equilibrium overpredicts the own-payoff effects, a result that is consistent with the data reported in McKelvey, Palfrey, and Weber (2000) discussed below. Essentially what is happening is that the high payoff seems to

² The necessary axioms are that choice probabilities be unaffected by adding a constant to all payoffs, and that ratios of probabilities for two decisions be independent of the payoffs associated with any other decision. An alternative derivation of the logit rule is based on an assumption that the payoffs for each decision are augmented by an unobserved preference shock, with a double exponential distribution. This random-preference model was used by Harsanyi (1973) in an equilibrium analysis, which is closely related to the quantal response equilibrium discussed below.

be undervalued by subjects, which moves row's stochastic best response in the direction of the data point.

One possible explanation for this undervaluation is risk aversion, since concavity of utility causes a relative undervaluation of higher payoffs. Ochs (1995) attempted to control for the effects of risk aversion by paying subjects in "points" that correspond to lottery tickets for a cash prize. In theory this procedure makes expected utility a linear function of these payoffs, which should induce risk neutrality. However, the experimental evidence on the validity of this procedure is at best mixed (e.g. Davis and Holt, 1993, Ch. 8). In addition, our econometric analysis in section IV indicates that the binary lottery procedure did not remove the effects of risk aversion in the Ochs experiment. First, however, we consider an experiment designed to separate the implications of risk aversion and noisy behavior as modeled by the quantal response equilibrium.

III. An Experiment with Safe and Risky Choices

The payoffs for Game 4 in Table II are chosen to exaggerate the effects of possible risk aversion, by giving each player a choice between a safe decision with low payoff variability (200 or 160) and a risky decision with variable but potentially high payoffs (370 or 10). This is also a matching pennies game. Under risk neutrality, the best response functions would be horizontal (for column) and vertical (for row) lines that intersect at the unique Nash equilibrium: $\Pr(U) = \Pr(R) = 0.53$.

Table II. Game 4: A Matching Pennies Game with "Safe" (200/160) and "Risky" (370/10) Decisions

		Column Player's Decision	
		<i>L</i>	<i>R</i>
Row Player's Decision	<i>U</i>	(200, 160)	(160, 10)
	<i>D</i>	(370, 200)	(10, 370)

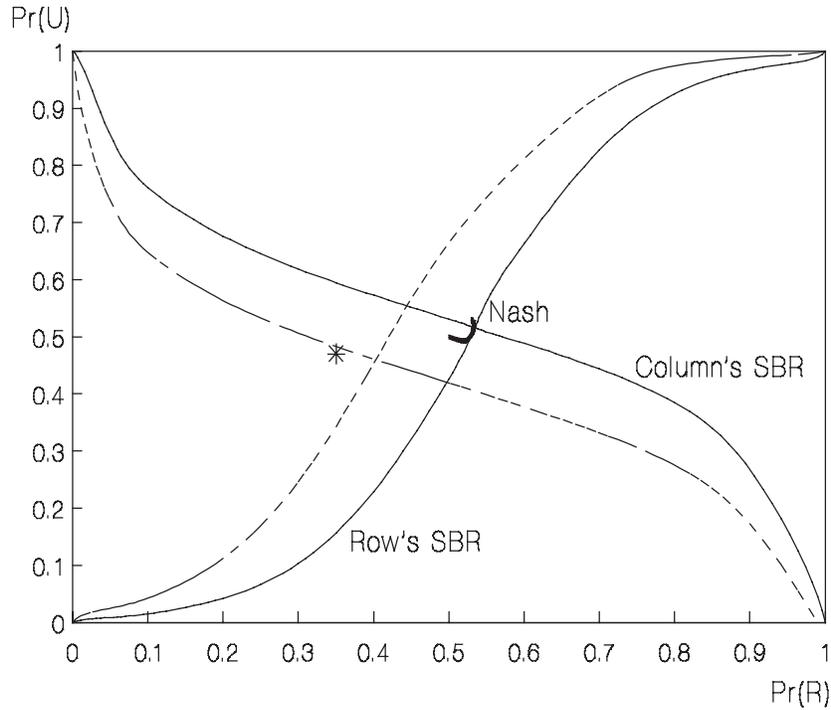


Figure 2. Game 4: Stochastic Best Response Lines for Risk Neutrality (Solid Lines) and for Risk Aversion (Dashed Lines) using Error and Risk Aversion Parameters from Lottery Choice Decisions

With some noise, the best response functions are smooth, as shown by the thin line curves in Figure 2 (please ignore the dashed lines for now). These thin line functions intersect at a point that is essentially on top of the Nash equilibrium. As the error parameter is increased, the locus of logit equilibria traces out a curve slightly down and to the left from this intersection, ending up at the midpoint (0.5, 0.5) for very high values of μ . This locus is shown in the figure as a small semi-circular dark line in the center. As one can see from this locus, the effects of noise appear to be minimal in this game, since smoothing the stochastic best reply functions has little effect on the central intersection point. In fact, we can show that any quantal response equilibrium (logit or otherwise) will always stay on the right side of the figure, with $\text{Pr}(R) \geq 0.5$, when players are risk neutral with respect to money payments in the game. (See Appendix A for proofs of the propositions.)

Proposition 1. *With risk neutrality, the probability of R will be greater than or equal to 0.5 in Game 4 for any quantal response equilibrium.*

If players are sufficiently risk averse, the predictions are quite different. Intuition suggests that the effect of risk aversion would be to increase reliance on column's safe strategy, L , which could put choice frequencies in the left part of the figure. The decision probabilities for row and column are determined jointly, however, so that the first-order intuitive effects are confounded by equilibrium effects that may go in the opposite direction. First consider what happens if we replace the expected payoff expressions in equation (1) with expected utilities, using $U(x) = x^{1-r}$, where x is the money payoff and r is the non-negative coefficient of constant relative risk aversion.³ The resulting concavity changes the scale of payoffs, effectively rescaling the error parameter μ . To correct for this scale effect, we normalize the utilities such that the utility of the worst possible outcome is zero and that of the best possible outcome is 1, i.e. $U(10) = 0$ and $U(370) = 1$, which implies:

$$U_r(x) = \frac{x^{1-r} - 10^{1-r}}{370^{1-r} - 10^{1-r}}, \quad (3)$$

which is defined for all r . (For $r = 1$, the natural logarithm is used.)

The dashed-line stochastic response curves in Figure 2 are drawn by holding the error parameter constant and increasing the risk aversion from 0 to 0.31, an estimate (discussed below) based on individual lottery-choice decisions made separately by these same subjects. Notice that row's dashed-line stochastic response function tends to shift upward, with higher probabilities on the safe decision U for each value of $\Pr(R)$. Similarly, column's stochastic response function tends to shift left, with higher probabilities for the safe decision L at each given level of $\Pr(U)$. The net effect of these two shifts is to reduce the logit equilibrium probability of R , but the

³ We choose to use the constant relative risk aversion for several reasons. This formulation is convenient for auction models and has been used by a number of others in the analysis of auction data both from the field (Campo, Perrigne, and Vuong, 2000), Cox and Oaxaca (1996), and others discussed below (see footnote 7). Obviously, there will be some heterogeneity in risk attitudes and these variations will have effects similar to the noise in a quantal response model. In fact, we report some estimates below for a model that combines a heterogenous random effects model into a quantal response equilibrium (see footnote 8).

probability of row's safe decision U actually goes *down* in this case. This is analogous to the observation that an upward shift in supply and a leftward shift in demand will reduce quantity, but the effect on price can go either way. The risk aversion effects shown in Figure 2 are easily generalized to the case of an arbitrary concave utility function, $U(x)$:

Proposition 2. *In the logit equilibrium, an increase in risk aversion leads to a decrease in p_R in Game 4.*

To summarize, the probability of R is greater than 0.5 in *any* quantal response equilibrium assuming risk neutrality, and the effect of risk aversion is to reduce p_R . To evaluate these predictions, we used Game 4 in an experiment with 6 cohorts of 10-12 subjects each. Subjects were randomly matched for 10 periods with a fixed role (row or column). Half of the sessions were run at the University of Virginia and half were done at Caltech using Caltech students (one session) and Pasadena City College students (2 sessions).⁴

The aggregate choice frequencies were 0.35 for R and 0.47 for U , as shown by the asterisk in Figure 2. Consistent with the risk averse quantal response equilibrium model, the probability of L was greater than 0.5 in all six sessions, which is highly significant using a binomial test with each session as an independent observation. Also notice that the data average point in Figure 2 is very close to the intersection of the stochastic response lines drawn for parameter estimates that were obtained using lottery choice data (see section V).

IV. Risk Aversion Estimates for Generalized Matching Pennies Games

In this section we use the logit equilibrium to obtain structural estimates of risk aversion parameters using data from several different generalized matching pennies experiments. We obtain separate estimates for each game and find they are virtually the same across games. This demonstrates that risk aversion parameters estimated from one game can be used to predict behavior in the other games. Maximum likelihood estimates are obtained using the utility

⁴ The choice percentages for each of the sessions are provided in the top four rows of the summary data table in the Appendix. Individual data and instructions are available at: <http://www.people.virginia.edu/~cah2k/datapage.html>.

function in (3); an alternative approach based on prospect theory is discussed below. For each value of the error (μ) and risk aversion parameter (r), the logit equilibrium probabilities are computed from the fixed-point equations in (5), where F is specified by the logistic distribution, $F(x) = \exp(x)/(1+\exp(x))$.⁵ Let the logit equilibrium choice probabilities for row to choose U and for column to choose R be denoted by $p_U(\mu, r)$ and $p_R(\mu, r)$, respectively. These equilibrium choice probabilities imply a likelihood function for the observed choice frequencies. To be more precise, for any given data set, X , one observes the aggregate number of times all row and column players choose U , D , L , or R , denoted N_U , N_D , N_L , and N_R , respectively. The log of the likelihood function for a particular game, G , is then simply:

$$\begin{aligned} \log(L^G(\mu, r)) = & N_U \log(p_U(\mu, r)) + N_D \log(1 - p_U(\mu, r)) \\ & + N_R \log(p_R(\mu, r)) + N_L \log(1 - p_R(\mu, r)). \end{aligned} \quad (4)$$

For each game we use a standard algorithm for finding the values of μ and r that maximize $\log(L^G)$, given the actual choice frequencies of the subjects who played that game.⁶

Table III. Maximum Likelihood Estimates
Asymptotic Standard Errors in Parentheses

Experiment	Risk Aversion r	Logit Error μ	LogLikelihood (number observations)
Matching Pennies Games 1-3 (Ochs, 1995)	0.42 (0.06)	0.010 (0.003)	-2071.9 (1560)
Matching Pennies Game 4 (this paper)	0.44 (0.07)	0.150 (0.060)	-455.8 (340)
Matching Pennies Games A-D (McKelvey, Palfrey, and Weber, 2000)	0.43 (0.02)	0.042 (0.004)	-5929.0 (4800)

The estimates from the three matching pennies experiments that Ochs conducted are

⁵ The logit equilibrium for this game is unique for all values of μ and r .

⁶ The likelihood function is different for each of the three games, since the logit equilibrium choice probabilities, $P^*(\mu, r)$ and $Q^*(\mu, r)$, depend on the payoff matrix of the game.

Table IV. Matching Pennies Games A ($X = 90$) and D ($X = 40$)

		Column Player's Decision	
		L	R
Row Player's Decision	U	$(X, 0)$	$(0, 10)$
	D	$(0, 10)$	$(10, 0)$

reported in the first row of Table III and the second row shows analogous estimates for our Game 4. Asymptotic standard errors are in parenthesis, and indicate that r is significantly different from one (risk neutrality) and μ is significantly different from zero (no error). Thus *both* error and risk aversion are significant components for explaining deviations from the minimax strategies of his game.⁷

Note that the risk aversion estimate for the Ochs (1995) data is virtually identical to the other risk aversion estimates in the table. This suggests that the binary lottery procedure failed to induce risk neutrality. While the risk aversion estimates are stable, the error parameter estimates are not, which we conjecture to be caused by to differences in payoff scales, subject pools, and procedures. For example, some sessions were done by hand and others were computerized. Furthermore, subjects in Ochs' experiments played the mixed extension of the game. Another relevant factor is the scale and normalization of payoffs. For instance, the estimated error parameter would be 100 times as high when payoffs are measured in pennies instead of dollars. We attempted to construct comparable payoff measures, but the appropriate normalizations become less obvious when payoffs are in lottery tickets (as in Ochs, 1995) or when only one out of ten decisions is selected *ex post* to determine earnings, as in our lottery choice experiment discussed below.

⁷ Observed autocorrelation in some people's decisions might be due to heterogeneity in individual risk attitudes. One can accommodate such heterogeneity by specifying a random effects model, which requires that there be a separate choice probability for each possible draw from the distribution of risk aversion parameters. Therefore, a system of simultaneous equations has to be solved to determine the likelihood function for each iteration over the parameter space, which includes an error parameter and the mean and variance of the distribution of risk aversion parameters. We estimated such a model for Game 4 and the estimated mean and variance of the risk aversion distribution are 0.55 and 0.40 respectively. The mean is close to the point estimate reported in the second line of Table III.

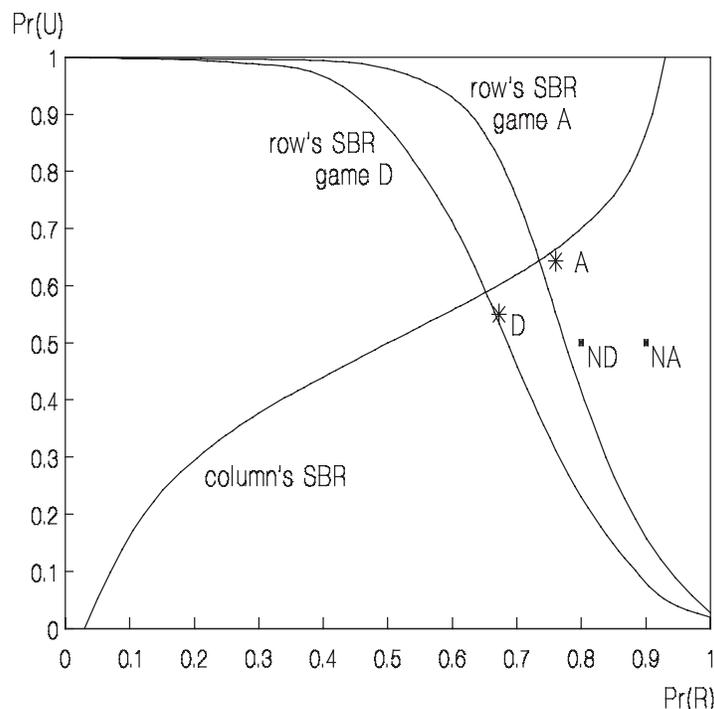


Figure 3. Stochastic Best Response Functions for Games A and D

We also estimate the quantal response model for the four matching pennies games in McKelvey, Palfrey, and Weber (2000); see the Appendix for game payoffs and aggregate results. The third row of Table III shows that the resulting risk aversion estimates are approximately the same. The estimates are also comparable to those from some other studies.⁸

In order to evaluate the fit, consider the McKelvey, Palfrey, and Weber (2000) games A and D, which are shown in Table IV. The reduction in row's upper left payoff from 90 in game A to 40 in game D leaves the mixed Nash prediction unchanged at 0.5 for the row player, but the prediction for column falls from 0.9 to 0.8, as shown by the shift from the point NA to ND in Figure 3. With the pooled estimates ($\mu = 0.042$ and $r = 0.43$), the stochastic best response function for column stays the same for games A and D, and the reduction in row's payoff, X ,

⁸ Goeree, Holt, and Palfrey (2002) report a risk aversion estimate of $r = 0.52$ (0.10) based on a maximum likelihood analysis of data from private-values auction experiments. Campo, Perrigne, and Vuong (2000) analyze field data from timber auctions and find an almost identical result: $r = 0.56$. These estimates are similar to estimates reported in Chen and Plott (1998) but somewhat lower than the individual-specific estimates reported by Cox and Oaxaca (1996). Finally, the risk aversion estimate of $r = 0.44$ reported in Table III matches exactly the coefficient reported in an analysis of 37 two-person matrix games played only once (Goeree and Holt, 2002).

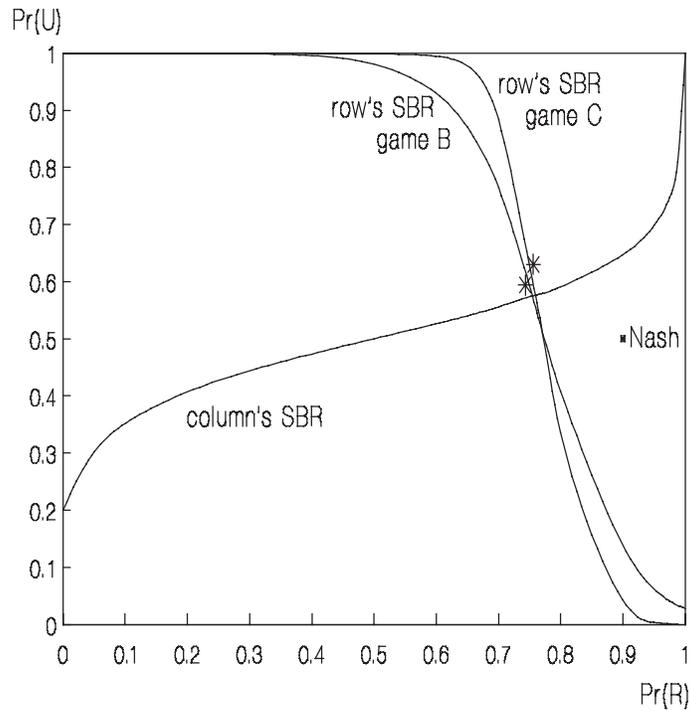


Figure 4. Stochastic Best Response Functions for Games B and C

shifts row’s stochastic best response downward in Figure 3. This shift produces quantal response intersections that are remarkably close to the asterisks labeled A and D that represent the decision averages for games A and D respectively. Obviously, these data points fall on a line with a positive slope, and the quantal response equilibrium predictions (with risk aversion) track these own-payoff effects very well.

The McKelvey, Palfrey, and Weber (2000) experiments included two other variations of the matching pennies game in order to determine the effects of increasing payoff magnitudes for one player, without altering the Nash equilibrium. Game B was obtained from A by multiplying column’s payoffs by 4, i.e. raising the payoffs of 10 to 40. Game C was then obtained from B by multiplying row’s payoffs in B by a factor of 4. With risk neutrality, payoff magnitude changes of this kind affect the logit equilibrium because the increase in row’s payoffs will make row’s best-response function in Figure 3 decline more sharply. McKelvey, Palfrey, and Weber (2000), however, reported that “the hypothesized payoff magnitude effects are not borne out in our data” (p. 534). They attempt to explain this puzzling finding in terms of subject

heterogeneity. Here we argue that, even in the absence of heterogeneity, an alternative explanation is risk aversion. Although payoff magnitudes do matter in a logit equilibrium, they also affect risk aversion, and these effects can work in opposite directions. Indeed, row's payoff increase from game B to C sharpens row's SBR function, but the net effect on the logit equilibrium (with $\mu = 0.042$ and $r = 0.43$) is essentially zero, as shown by the two intersection points in Figure 4. The asterisks representing the data averages for these two treatments lie almost on top of one another, and both data averages in the figures are well away from the Nash equilibrium – and close to the logit equilibrium – for both games. Thus, the absence of a clear treatment effect when row's payoff parameters are scaled up by a factor of 4 is not perplexing when risk aversion is incorporated into the stochastic response functions.

An alternative explanation of seemingly risk averse behavior is based the notion of probability weighting, i.e. that low probabilities tend to be overweighted and high probabilities tend to be underweighted (Kahneman and Tversky, 1979). Prelec (1998) argues that the weighting function, $w(p)$, has an inverted "S" shape: concave for low probabilities, convex for high probabilities, which crosses the diagonal at a probability of about 1/3. A one-parameter formula that allows inverted S-shapes is:

$$w(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}}, \quad (5)$$

where $\delta \geq 0$. Notice that (5) satisfies the endpoint conditions: $w(0) = 0$ and $w(1) = 1$, and has the inverted S-shape when $\delta < 1$. In this case, probabilities close to 0 will be overweighted and those close to 1 will be underweighted.⁹

This weighting function is applied to the probability with which the other player chooses the "safe" strategy (U for row and L for column). The weight for the "risky" strategy is then $1 - w(p)$. The estimated model fits the data well and yields roughly the same likelihood. For example, for the Ochs data we estimate $\delta = 1.32$ (0.09) and $\mu = 0.47$ (0.10) with a loglikelihood of -2075.1. Note, however, that the estimated weighting function does not have the inverted "S"

⁹ This can explain why individuals are willing to bet on low probability gains, and why they shy away from prospects when sure gains are diluted by the small chance of a loss.

shape that has been used to explain anomalies in other contexts. The estimated δ is greater than one and thus yields the wrong curvature: small probabilities are underweighted not overweighted.

V. Lottery Choice Data

As a further calibration of our risk aversion estimates, all subjects who participated in the experiments based on Game 4 also made ten choices between "safe" (S) and "risky" (R) lotteries, which were constructed from the payoffs of game 4. Lottery S paid 200 cents with probability p and 160 cents with probability $1 - p$, where p varied from 0.1, 0.2, ..., 1.0 in the 10 different choice problems. Similarly, lottery R yielded 370 cents with probability p and 10 cents with probability $1 - p$. A subject's pattern of S and R choices gives an independent measure of that subject's aversion to risk. To see this, note that if utilities are described by (3), the constant of relative risk aversion, r , determines a unique "switch point," p^* : for $p < p^*$, lottery S is preferred while R is better for $p > p^*$. From (3), the critical level is easily determined as

$$p^*(r) = \frac{160^{1-r} - 10^{1-r}}{370^{1-r} - 10^{1-r} + 160^{1-r} - 200^{1-r}}. \quad (6)$$

If the probability p would have been varied continuously in the experiment, the point at which a subject switches from S to R uniquely ties down the risk aversion parameter by (6). However, in the experiment we considered only ten possible values of p , and the number of S choices therefore determines a range of possible risk aversion parameters. The ranges that correspond to a given number of "safe" choices are shown in the second column of Table V. The observed choice frequencies sorted by subject pool are given in columns 4-6, and the final column gives the aggregate data.

After a subject made all of the ten paired lottery choice decisions, one was selected at random, *ex post*, to be played out, and the outcome of the selected lottery was added to the earnings of that subject. This lottery choice part was prior to Game 4 in the Virginia sessions, and it was done afterwards in the sessions conducted with Caltech and Pasadena City College students. The percentages of safe choices were remarkably stable across sessions, i.e. in the interval from 52-56 percent in each session. The average across all subjects was 54 percent, which corresponds to a risk aversion parameter of about 0.3. There is some variability across

Table V. Risk Aversion Classifications Based on Lottery Choices

Number of Safe Choices	CRRA Range for $U(x) = x^{1+r}$	Classification	Numbers of Choices			
			UVA	Caltech	PCC	Total
2	$r < -0.56$	very risk loving	1	0	0	1
3	$-0.56 < r < -0.2$	risk loving	4	1	5	10
4 ^a	$-0.20 < r < 0.09$	risk neutral	7	2	4	13
5	$0.09 < r < 0.36$	slightly risk averse	11	1	3	15
6	$0.36 < r < 0.64$	risk averse	15	8	3	26
7	$0.64 < r < 0.93$	very risk averse	5	0	3	8
8 or more ^b	$0.93 < r$	highly risk averse	1	0	4	5

^a Does not include subject 7-11 whose behavior (RRRSRSSRRS) seems random.

^b Only one subject selected the safe option 9 times.

subjects, as can be seen from the risk classifications by individual in Table V. A few (14%) seem to be risk seeking or close to risk neutral (16%), with the rest in the risk averse categories. The modal choice pattern (32%) involves six S choices, which is consistent with the interval of relative risk aversion coefficients centered around $r = 0.5$, or square root utility.¹⁰

The average risk aversion parameter of about 0.3 is corroborated by a maximum likelihood analysis of the 800 individual binary choices. With $U_r(10) = 0$ and $U_r(370) = 1$, the expected utility difference between the safe and risky choices is: $p - p U_r(200) - (1 - p) U_r(160)$. Hence, the logit model predicts that the probability of choosing lottery S is

$$p_s = \frac{1}{1 + \exp((p - p U_r(200) - (1 - p) U_r(160))/\mu)}, \quad (7)$$

where the treatment variable p runs from 0.1, 0.2, ..., 1.0, and the utility function U_r is given in

¹⁰ Holt and Laury (2002) also use a lottery choice menu to obtain individual risk aversion estimates. The mean of these estimates is similar to the estimate reported here, although risk aversion increased as payoffs were scaled up by factors of 20, 50, and 90. These menu-based methods of eliciting risk aversion are alternatives to the procedure of asking a buying or selling price for a lottery. It is well known that such willingness-to-pay/willingness-to-accept methods may yield inconsistent conclusions about risk aversion (Kachelmeier and Shehata, 1992). Furthermore, risk aversion estimates obtained from willingness-to-accept methods may be inversely related to estimates inferred from bids in an auction (Isaac and James, 1999).

(3). Maximizing the likelihood of observed choices yields a risk aversion estimate $r = 0.31$ (.03).

In fact, the lottery choice error and risk aversion estimates were the parameters used to draw the dashed stochastic best response lines for Game 4 in Figure 2. Since these intersect a little to the right of the asterisk that plots the data average, it is not surprising that a slightly higher risk aversion estimate (0.44) provides a better fit by shifting row's quantal response curve up and column's to the left. This higher level of risk aversion is in the "risk averse" interval that contains the modal choice in Table V.

VI. Conclusion

This paper considers causes for the disparity between the Nash equilibrium and experimental data for matching pennies games. A simple two-parameter model that combines quantal response equilibrium and risk aversion explains the observed choice patterns. Neither risk aversion nor quantal response equilibrium *alone* can account for the data across the range of games we study. In a model with only risk aversion, the mixed equilibrium choice probability for one player is independent of that player's own payoffs, which is inconsistent with observed data patterns. The Ochs (1995) data exhibit own-payoff effects, i.e. row's tendency to choose strategy "Up" when the payoffs for this strategy are increased. Since column's best-response line for these games is horizontal for any risk attitudes, all Nash equilibria lie on a flat line at a height of 0.5 in Figure 1. Hence, risk aversion alone cannot provide a full explanation of observed behavior.

The quantal response equilibrium introduces noise into the decision-making process thereby "smoothing" the best-responses. The resulting S-shaped response functions can produce own-payoff effects in a way that is analogous to the price increase that results when demand shifts and supply is not perfectly elastic. Game 4 was designed to determine whether noise-based theories are able to explain observed deviations from Nash, without resorting to risk aversion. Proposition 1 shows that, with risk neutrality, all quantal response equilibria are on the right side of Figure 2, with $\Pr(R) > 0.5$. However, the observed data lie on the other side of the figure, so the incorporation of noise into the equilibrium analysis will not provide good predictions for this game, as long as risk neutrality is assumed. Proposition 2 shows that increases in risk aversion shift the response functions toward the safe strategies, so that the quantal response intersection

moves to the left in Figure 2, towards the observed data averages. Maximum likelihood estimates of the risk aversion parameter are relatively stable across the different games and are only slightly higher than an estimate based on individual lottery-choice tasks performed by the same subjects. The resulting risk-averse/quantal response equilibrium model tracks the data closely and explains some perplexing patterns, including the absence of payoff magnitude effects in the McKelvey, Palfrey, and Weber (2000) experiments. The risk aversion estimates reported here, which approximately yield square-root utility, should not be extrapolated to lotteries with losses or large gains. In particular, it may well be the case that people are risk neutral or risk loving in the loss domain.

The focus of this paper is on the aggregate choice frequencies (across time and subjects), rather than on individual differences and the dynamics of learning in games. Adaptive learning models have been used to study the time series of choices in matching pennies games (Mookherjee and Sopher, 1997; Erev and Roth, 1998; and Camerer and Ho, 2000). Moreover, there is evidence of heterogeneity across subjects (McKelvey, Palfrey, and Weber, 2000). Finally, Goeree and Holt (2002) develop a model of iterated noisy introspection that is applied to matching pennies and other games when these are played only once, so that equilibrium and dynamic learning theories are not applicable. These alternative modeling approaches (learning, heterogeneity, and introspection) are useful complements to the equilibrium analysis presented here. For example, introspection models may aid in predicting behavior in the first period of a repeated-play experiment, learning may explain patterns of adjustment, and equilibrium theories may explain data patterns that have stabilized.

Finally, the magnitude of our risk aversion estimates are about the same as has been found in econometric studies of both field and laboratory auction data. The importance of risk aversion is also suggested by the similar estimates we obtained from a risk aversion model applied to individual lottery choice data. The results of this paper suggest that the quantal response equilibrium may be useful in a wide class of applied econometric studies. With respect to the empirical study of auctions, initial attempts to develop methodologies for doing this in the contexts of both laboratory data (Goeree, Holt, and Palfrey, 2002) and field data (Bajari and Hortacsu, 2001) indicate that this approach is both feasible and promising.

Data Appendix: Experimental Design and Summary Data

Treatment	Payoffs in Pennies	Session Matching	Location (No. Subjects)	Aggregate Data
Game 4 Safe/Risky	L R	1 Random	UVA (10)	0.36 U, 0.68 L
	U (200, 160) (160, 10)	2 Random	UVA (12)	0.52 U, 0.72 L
	D (370, 200) (10, 370)	3 Random	UVA (10)	0.56 U, 0.56 L
		4 Random	CIT (12)	0.53 U, 0.72 L
		5 Random	PCC (12)	0.42 U, 0.52 L
			6 Random	PCC (12)
MPW Game A*	L R	1a Random	CIT (12)	0.68 U, 0.32 L
	(90, 0) (0, 10)	2b Random	CIT (12)	0.62 U, 0.24 L
	(0, 10) (10, 0)	5a Random	CIT (12)	0.62 U, 0.11 L
		6b Random	CIT (12)	0.61 U, 0.22 L
		7a Random	CIT (12)	0.64 U, 0.23 L
		8a Random	CIT (12)	0.68 U, 0.24 L
MPW Game B*	L R	1b Random	CIT (12)	0.57 U, 0.20 L
	(90, 0) (0, 40)	2a Random	CIT (12)	0.65 U, 0.24 L
	(0, 40) (10, 0)	3a Random	CIT (12)	0.61 U, 0.36 L
		4b Random	CIT (12)	0.69 U, 0.18 L
MPW Game C*	L R	3b Random	CIT (12)	0.57 U, 0.39 L
	(360, 0) (0, 40)	4a Random	CIT (12)	0.62 U, 0.16 L
	(0, 40) (40, 0)	5b Random	CIT (12)	0.59 U, 0.20 L
		6a Random	CIT (12)	0.59 U, 0.27 L
MPW Game D*	L R	7b Random	CIT (12)	0.46 U, 0.31 L
	(40, 0) (0, 10) (0, 10) (10, 0)	8a Random	CIT (12)	0.64 U, 0.34 L
Lottery Choice	Safe Lottery: (200, p; 160, 1-p)	1 (game 3)	UVA (12)	0.54 Safe
		2 (game 3)	UVA (10)	0.55 Safe
		3 (game 3)	UVA (12)	0.53 Safe
	Risky Lottery: (370, p; 10, 1-p)	4 (game 3)	UVA (10)	0.52 Safe
		5 (game 3)	CIT (12)	0.55 Safe
		6 (game 3)	PCC (12)	0.56 Safe
	with probability p = 0.1, 0.2, ... 1.0	7 (game 3)	PCC (12)	0.52 Safe

* Data for games A-D from McKelvey, Palfrey, and Weber (2000).

References

- Bajari, Patrick and Ali Hortacsu (2001) "Auction Models When Bidders Make Small Mistakes: Consequences for Theory and Estimation," Discussion Paper, Department of Economics, Stanford University.
- Bolton, Gary E. (1998) "Bargaining and Dilemma Games: From Laboratory Data Towards Theoretical Synthesis," *Experimental Economics*, 1:257-81.
- Brams, Steven J. (1994) "Game Theory and Literature," *Games and Economic Behavior*, 6, 32-54.
- Camerer, Colin F. (1997) "Progress in Behavioral Game Theory," *Journal of Economic Perspectives*, 11:167-88.
- Camerer, Colin and Teck-Hua Ho (1999) "Experience Weighted Attraction Learning in Normal-Form Games," *Econometrica*, 67, 827-874.
- Campo, Sandra, Isabelle Perrigne, and Quang Vuong (2000) "Semi-Parametric Estimation of First-Price Auctions with Risk-Averse Bidders," Discussion Paper, University of Southern California.
- Chen, Kay-Yut and Charles R. Plott (1998) "Nonlinear Behavior in Sealed Bid First-Price Auctions," *Games and Economic Behavior*, 25, 34-78.
- Cox, James C., and Ronald L. Oaxaca (1996) "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions," in R. M. Isaac, ed., *Research in Experimental Economics*, Vol. 6, Greenwich, Conn.: JAI Press, 131-148.
- Davis, Douglas D. and Charles A. Holt (1993) *Experimental Economics*, Princeton: Princeton University Press.
- Erev, Ido and Alvin E. Roth (1998) "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88(4), 848-881.
- Fehr, Ernst and Klaus Schmidt (1999) "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, August, 114(3), 817-868.
- Goeree, Jacob K. and Charles A. Holt (1999) "Stochastic Game Theory: For Playing Games, Not Just For Doing Theory," *Proceedings of the National Academy of Sciences*, 96, September, 10564-10567.
- Goeree, Jacob K. and Charles A. Holt (2000) "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," *European Economic Review*, 44, 1079-1089.

- Goeree, Jacob K. and Charles A. Holt (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402-1422.
- Goeree, Jacob, and Charles Holt (2002) "A Model of Noisy Introspection," Discussion Paper, University of Virginia.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," forthcoming *Journal of Economic Theory*.
- Harsanyi, J. (1973) "Games with Randomly Disturbed Payoffs," *International Journal of Game Theory*, 2(1), 1-23.
- Holt, Charles A. and Susan K. Laury (2002) "Risk Aversion and Incentive Effects," *American Economic Review*, forthcoming.
- Isaac, R. Mark and Duncan James (1999) "Just Who Are You Calling Risk Averse?" Discussion Paper, University of Arizona (forthcoming in the *Journal of Risk and Uncertainty*).
- Kachelmeier, Steven J. and Mohamed Shehata (1992) "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, December, 82, 1120-1141.
- Kahneman, Daniel and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 53, 75-100.
- Luce, Duncan (1999) *Individual Choice Behavior*, New York: Wiley.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- McKelvey, Richard D. and Thomas R. Palfrey (1996) "A Statistical Theory of Equilibrium in Games," *Japanese Economic Review*, 47:2 (June), 186-209.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber (2000) "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2×2 Games with Unique Mixed Strategy Equilibria," *Journal of Economic Behavior and Organization*, 42:4 (August), 523-48.
- Mookerjee, Dilip, and Barry Sopher (1997) "Learning and Decision Costs in Experimental Constant Sum Games," *Games and Economic Behavior*, 19(1), 197-132.
- Ochs, Jack (1995) "Games with Unique Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.
- Prelec, Drazen (1998) "The Probability Weighting Function," *Econometrica*, 66(3), 497-527.

Appendix. Proofs of Propositions 1 and 2

Proposition 1. *With risk neutrality, the probability of R will be greater than or equal to 0.5 in Game 4 for any quantal response equilibrium.*

Proof. Let p_R and p_U denote the probability of choosing "Right" and "Up" respectively. Any (risk-neutral) quantal response equilibrium is defined by the two equations:

$$\begin{aligned} p_R &= F((\pi^e(\mathbf{R}) - \pi^e(\mathbf{L}))/\mu) = F((170 - 320p_U)/\mu), \\ p_U &= F((\pi^e(\mathbf{U}) - \pi^e(\mathbf{D}))/\mu) = F((320p_R - 170)/\mu), \end{aligned} \tag{A1}$$

where F is a continuous cumulative distribution function (with full support) that satisfies $F(x) = 1 - F(-x)$, so $F(0) = 1/2$. (Note that (A1) reduces to equations (1) and (2) for the particular logistic distribution $F(x) = \exp(x)/(1+\exp(x))$ that is frequently used in estimation.) Suppose, in contradiction, that $p_R < 1/2$. From the second line of (A1) it then follows that $p_U < 1/2$, since $320p_R - 170 < 0$ and $F((320p_R - 170)/\mu) < F(0) = 1/2$. In turn, $p_U < 1/2$ implies that $170 - 320p_U > 0$, which together with the top line of (3) yields $p_R > 1/2$, a contradiction. Q.E.D.

Proposition 2. *In the logit equilibrium, an increase in risk aversion leads to a decrease in p_R in Game 4.*

Proof. As before we normalize the utility of the worst outcome to be zero and the utility of the best outcome to be one: $U(10) = 0$ and $U(370) = 1$. Equation (A1) that defined the risk-neutral quantal response equilibrium is replaced by:

$$\begin{aligned} p_R &= F(((1 - p_U) - U(160)p_U - U(200)(1 - p_U))/\mu), \\ p_U &= F((U(160)p_R + U(200)(1 - p_R) - (1 - p_R))/\mu), \end{aligned} \tag{A2}$$

where, as before, F is a cumulative distribution function that satisfies $F(x) = 1 - F(-x)$, so $F(0) = 1/2$. It is sufficient to show that an increase in the Arrow-Pratt measured of risk aversion, $-U''(x)/U'(x)$, raises p_U for constant p_R , and lowers p_R for constant p_U , as illustrated by the shift

from solid to dashed lines in Figure 2. In particular, let the more risk averse utility function be denoted by $U_1(x)$, so $-U_1''(x)/U_1'(x) \geq -U_2''(x)/U_2'(x)$ for all $10 < x < 370$. From the top line in (A2), an increase in risk aversion lowers p_R for constant p_U if it raises both $U(160)$ and $U(200)$. In fact, we prove that $U_1(x) > U_2(x)$ for all $10 < x < 370$, i.e. the "more concave" utility function is everywhere higher. Suppose, in contradiction, that $U_1(x) < U_2(x)$ for some x . Let x^* be the point at which $U_2(x) - U_1(x)$ is maximized. At x^* we must have: $U_2'(x^*) = U_1'(x^*)$ and $U_2''(x^*) \leq U_1''(x^*)$, which are the necessary first and second-order conditions for $U_2(x) - U_1(x)$ to be maximized. These conditions imply $-U_2''(x^*)/U_2'(x^*) \geq -U_1''(x^*)/U_1'(x^*)$, which contradicts the original assumption that U_1 is more risk averse. To prove that an increase in risk aversion raises p_U for constant p_R follows immediately from the bottom line of equation (A2). The resulting upward shift in row's stochastic best response and the leftward shift in column's stochastic best response (see Figure 2) results in a lower equilibrium for p_R Q.E.D.