Testing the Quantal Response Hypothesis*

Emerson Melo[†] Kirill Pogorelskiy [‡] Matthew Shum[‡]

March 27, 2015

Abstract

This paper develops a formal test for consistency of players' behavior in a series of games with the quantal response equilibrium (QRE). The test exploits a characterization of the equilibrium choice probabilities in a QRE as the gradient of a convex function, which thus satisfies the *cyclic monotonicity* inequalities. Our testing procedure utilizes recent econometric results for moment inequality models. We assess our test using lab experimental data from a series of generalized matching pennies games. We reject the QRE hypothesis in the pooled data, but it cannot be rejected in the individual data for over half of the subjects.

JEL codes: C12, C14, C57, C72, C92

Keywords: quantal response equilibrium, cyclic monotonicity, moment inequalities, lab experiment

1 Introduction

A vast literature in experimental economics has demonstrated that, across a wide variety of games, behavior deviates systematically from Nash equilibrium-predicted behavior. In order to relax Nash equilibrium in a natural fashion, while preserving the idea of equilibrium, McKelvey & Palfrey (1995) introduced the notion of *Quantal Response Equilibrium (QRE)*. QRE has become a popular tool in experimental economics because typically it provides an improved fit to the experimental data.

In this paper we develop a procedure to test, using experimental data, whether subjects are indeed behaving according to QRE. Our test is based on the notion of *cyclic monotonicity*, a concept from convex analysis which is useful as a characterizing feature of convex potentials. Cyclic monotonicity imposes joint inequality restrictions between the underlying choice frequencies and payoffs from the underlying games; hence, we are able to apply tools and methodologies from the recent econometric literature on moment inequality models to derive the formal statistical properties of our test. Importantly, our test for QRE is *nonparametric* in that one need not specify

*Acknowledgments: We thank Marina Agranov, Larry Blume, Federico Echenique, Ben Gillen, Phil Haile, Jörg Stoye, Leeat Yariv and especially Tom Palfrey for support and insightful comments. We are grateful to Michael McBride and Alyssa Acre at UC Irvine ESSL laboratory for help with running experiments.

[†]Department of Computer Science, Cornell University. Email: ecm226@cornell.edu

[‡]Division of the Humanities and Social Sciences, California Institute of Technology. Emails: kbp@caltech.edu, mshum@caltech.edu

a particular probability distribution for the random shocks; thus, the results are robust to a wide variety of distributions.

Subsequently, we apply our test to data from a lab experiment on generalized matching pennies games. We find that QRE is rejected soundly when data is pooled across all subjects and all plays of each game. But when we consider subjects individually, we find that the QRE hypothesis cannot be rejected for upwards of half the subjects. This suggests that there is substantial heterogeneity in behavior across subjects. Moreover, the congruence of subjects' play with QRE varies substantially depending on whether subjects are playing in the role of the Row vs. Column player.

Our work here extends Haile, Hortacsu, & Kosenok (2008) (hereafter HHK) who showed that, that without imposing strong assumptions on the shock distributions, QRE can rationalize any outcome in a given game. HHK describe several approaches for testing for QRE, but without discussing the econometric implementation of such tests. We build upon one of these approaches, based on the relation between changes in QRE probabilities across the series of games that only differ in the payoffs and the changes in the respective payoffs, and develop an econometric test for consistency of the data with a QRE in this more general case.

Our use of the notion of cyclic monotonicity to test the QRE hypothesis appears new to the experimental game theory literature. Elsewhere, cyclic monotonicity has been studied in the context of multidimensional mechanism design. In particular, the papers by Rochet (1987), Saks & Yu (2005), Lavi & Swamy (2009), Ashlagi et al. (2010), and Archer & Kleinberg (2014) (summarized in Vohra (2011, Chapter 4)), relate the incentive compatibility (truthful implementation) of a mechanism to its cyclic monotonicity properties. Similarly, the papers by Fosgerau & de Palma (2015) and McFadden & Fosgerau (2012) introduce cyclic monotonicity to study revealed preference in discrete choice models.

The rest of the paper is organized as follows. Section 2 presents the QRE approach. Section 3 introduces the test for the QRE hypothesis, and Section 4 discusses the moment inequalities for testing. Section 5 discusses the statistical properties of the test. Section 6 describes our experiment, with subsections 6.1 and 6.2 presenting the experimental design and results respectively. Section 7 concludes. Appendix A provides additional details about the interpretation of the cyclic monotonicity inequalities. Appendices B and C contain omitted proofs and additional theoretic results. Appendix D contains omitted computational details of the test. Appendix E contains experimental instructions.

2 QRE background

In this section we briefly review the main ideas behind the QRE approach. We use the notation from McKelvey & Palfrey (1995).

Consider a finite n-person game $G(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$. The set of strategies available to player i is indexed by $j = 1, \ldots, J_i$, so that $S_i = \{s_1, \ldots, s_{J_i}\}$, with a generic element denoted s_{ij}. Let $\boldsymbol{s}$ denote an n-vector strategy profile; let s_i and $\boldsymbol{s}_{-i}$ denote player i's (scalar) action and the vector of actions for all players other than i. In terms of notation, all vectors are denoted by bold letters.

Let p_{ij} be the probability that player i chooses action j, and $\boldsymbol{p}_i$ denote the vector of player i's choice probabilities. Let $\boldsymbol{p} = (\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n)$ denote the vector of probabilities across all the players. Player i's utility function is given by $u_i(s_i, \boldsymbol{s}_{-i})$. At the time she chooses her action, she does not know what actions the other players will play. Define the expected utility that player i gets from playing a pure strategy s_{ij} when everyone else's joint strategy is $\boldsymbol{p}_{-i}$ as

$$u_{ij}(\boldsymbol{p}) \equiv u_{ij}(\boldsymbol{p}_{-i}) = \sum_{\boldsymbol{s}_{-i}} p(\boldsymbol{s}_{-i}) u_i(s_{ij}, \boldsymbol{s}_{-i}),$$

where $\boldsymbol{s}_{-i} = (s_{kj_k})_{k \in N_{-i}}$, and $p(\boldsymbol{s}_{-i}) = \prod_{k \in N_{-i}} p_{kj_k}$.

In the QRE framework uncertainty is generated by players' making "*mistakes*". This is modelled by assuming that, given her beliefs about the opponents' actions $\boldsymbol{p}_{-i}$, when choosing her action, player i does not choose the action j that maximizes her expected utility $u_{ij}(\boldsymbol{p})$, but rather chooses the action that maximizes $u_{ij}(\boldsymbol{p}) + \varepsilon_{ij}$, where ε_{ij} represents a preference shock at action j. For each player $i \in N$ let $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \ldots, \varepsilon_{iJ_i})$ be drawn according to an absolutely continuous distribution F_i with mean zero. Then an expected utility maximizer, player i, given beliefs $\boldsymbol{p}$, chooses action j iff

$$u_{ij}(\boldsymbol{p}) + \varepsilon_{ij} \geq u_{ij'}(\boldsymbol{p}) + \varepsilon_{ij'}, \quad \forall j' \neq j.$$

Since preference shocks are random, the probability of choosing action j given beliefs $\boldsymbol{p}$, denoted $\pi_{ij}(\boldsymbol{p})$, can be formally expressed as

$$\begin{aligned}\pi_{ij}(\boldsymbol{p}) &\equiv \mathbb{P}\left(j = \arg\max_{j' \in \{1,\ldots,J_i\}} \left\{u_{ij'}(\boldsymbol{p}) + \varepsilon_{ij'}\right\}\right) \\ &= \int_{\left\{\boldsymbol{\varepsilon}_i \in \mathbb{R}^{J_i} \mid \ u_{ij}(\boldsymbol{p}) + \varepsilon_{ij} \geq u_{ij'}(\boldsymbol{p}) + \varepsilon_{ij'} \ \forall j' \in \{1,\ldots,J_i\}\right\}} dF_i(\boldsymbol{\varepsilon}_i) \end{aligned} \tag{1}$$

Then a *Quantal Response Equilibrium* is defined as a set of choice probabilities $\left\{\pi^*_{ij}\right\}$ such that for all $(i, j) \in N \times \{1, \ldots, J_i\}$,

$$\pi^*_{ij} = \pi_{ij}(\boldsymbol{\pi}^*)$$

Throughout, we assume that all players' preference shock distributions are *invariant*; that is, the distribution does not depend on the payoffs:

Assumption 1. *(Invariant shock distribution) For all realizations $\boldsymbol{\varepsilon}_i := (\varepsilon_{i1}, \ldots, \varepsilon_{iJ_i})$ and all payoff functions $u_i(\cdot)$, we have $F_i(\boldsymbol{\varepsilon}_i | u_i(\cdot)) = F_i(\boldsymbol{\varepsilon}_i)$.*

Such an invariance assumption was also considered in HHK's study of the quantal response model, and also assumed in most empirical implementations of QRE.[1]

[1] Most applications of QRE assume that the utility shocks follow a logistic distribution, regardless of the magnitude of payoffs. One exception is McKelvey, Palfrey, & Weber (2000), who allow the logit-QRE parameter to vary across different games. This direction is further developed in Rogers, Palfrey, & Camerer (2009).

3 A test based on convex analysis

In this section we propose a test for the QRE hypothesis. We start by defining the following function:

$$\varphi^i(\boldsymbol{u}_i(\boldsymbol{\pi})) \equiv \mathbb{E}\left[\max_{j\in S_i}\{u_{ij}(\boldsymbol{\pi}) + \varepsilon_{ij}\}\right] \tag{2}$$

In the discrete choice model literature, the expression $\varphi(\boldsymbol{u})$ is known as the social surplus function.[2] Importantly, this function is smooth and convex. Now the QRE probabilities $\pi_{ij}(\boldsymbol{\pi}^*)$ can be expressed as

$$\boldsymbol{\pi}_i^* = \boldsymbol{\nabla}\varphi^i(\boldsymbol{u}_i(\boldsymbol{\pi}^*)) \tag{3}$$

This follows from the well-known Williams-Daly-Zachary theorem from discrete-choice theory (which can be considered a version of Roy's Identity for discrete choice models; see Rust (1994, p.3104)). Thus if Eq. (3) holds for all players i, then $\left\{\pi_{ij}^*\right\}$ is a quantal response equilibrium. Eq. (3) characterizes the QRE choice probabilities as the gradient of the convex function φ. It is well-known (Rockafellar, 1970, Theorem 24.8) that the gradient of a convex function satisfies a *cyclic monotonicity* property. This property is the generalization, for functions of several variables, of the fact that the derivative of a univariate convex function is monotone nondecreasing.

To define cyclic monotonicity in our setting, consider a *cycle*[3] of games $C \equiv \{G_0, G_1, G_2, \dots, G_0\}$ characterized by the same set of choices for each player, but distinguished by payoff differences, where G_m denotes the game at index m in a cycle. Let $[\boldsymbol{\pi}_i^*]^m$ denote the QRE choice probabilities for player i in game G_m, and $\boldsymbol{u}_i^m \equiv \boldsymbol{u}_i^m([\boldsymbol{\pi}^*]^m)$ the corresponding equilibrium expected payoffs. Then the cyclic monotonicity property says that

$$\sum_{m=G_0}^{G_{\mathcal{L}-1}} \left\langle [\boldsymbol{u}_i]^{m+1} - [\boldsymbol{u}_i]^m, [\boldsymbol{\pi}_i^*]^m \right\rangle \leq 0 \tag{4}$$

for all finite cycles of games of length $\mathcal{L} \geq 2$, and all players i.[4] Expanding the inner product notation, the cyclic monotonicity conditions may be written as follows:

$$\sum_{m=G_0}^{G_{\mathcal{L}-1}} \sum_{j=1}^{J_i} \left(u_{ij}^{m+1} - u_{ij}^m\right) [\pi_{ij}^*]^m \leq 0 \tag{5}$$

This property only holds under the invariance assumption.

In Appendix A we show that CM conditions can be derived from, and also imply, players' utility maximization in appropriately perturbed games. Thus CM is a characteristic feature of optimal behavior in this sense. Intuitively, this interpretation illustrates how the CM inequalities capture players' incentives compatibility across games. Namely, if the CM inequalities are violated for player i and some cycle of games C, then i is not optimally adjusting her choice probabilities in

[2] For details see McFadden (1981).

[3] A *cycle* of length $\mathcal{L}$ is just a sequence of $\mathcal{L}$ games $G_0, \dots, G_{\mathcal{L}-2}, G_{\mathcal{L}-1}$ with $G_{\mathcal{L}-1} = G_0$.

[4] Under convexity of $\varphi(\cdot)$, we have $\varphi(\boldsymbol{u}_i^{m+1}) \geq \varphi(\boldsymbol{u}_i^m) + \left\langle \nabla\varphi(\boldsymbol{u}_i^m), (\boldsymbol{u}_i^{m+1} - \boldsymbol{u}_i^m)\right\rangle$. Substituting in $\nabla\varphi(\boldsymbol{u}_i^m) = \boldsymbol{\pi}_i^m$ and summing across a cycle, we obtain the CM inequality in Eq. (4).

response to changes in the expected payoffs across the games in the cycle C.

The number of all finite game cycles times the number of players can be, admittedly, very large. To reduce it, we note that the cyclic monotonicity conditions (5) are invariant under the change of the starting game index in the cycles; for instance, the inequalities emerging from the cycles $\{G_i, \ldots, G_j, G_k, G_i\}$ and $\{G_j, \ldots, G_k, G_i, G_j\}$ are the same.

Special case: two strategies. When each player's strategy set has only two pure strategies, the Cyclic Monotonicity conditions (5) only need to be checked for cycles of length 2. Because many experiments study games where players' strategy sets consist of two elements, this observation turns out to be useful from an applied perspective.[5]

4 Moment inequalities for testing cyclic monotonicity

Consistency with QRE can be tested nonparametrically from experimental data in which the same subject i is playing a series of one-shot games with the same strategy spaces such that each game is played multiple times. In this case, the experimental data allows to estimate a vector of probabilities $[\boldsymbol{\pi}_i^*]^m \in \Delta(S_i)$ for each game m in the sample, and we can compute the corresponding equilibrium expected utilities $[\boldsymbol{u}_i]^m$ (assuming risk-neutrality).

Suppose there are $M \geq 2$ different games in the sample. We assume that we are able to obtain estimates of $\hat{\boldsymbol{\pi}}_i^m$, the empirical choice frequencies, from the experimental data, for each subject i and for each game m. Thus we compute $\hat{\pi}_{ij}^m$ from K trials for subject i in game m:

$$\hat{\pi}_{ij}^m = \frac{1}{K} \sum_{k=1}^{K} \mathbb{1}_{\{i \text{ chooses } j \text{ in trial } k \text{ of game } m\}}$$

This will be the source of the sampling error in our econometric setup. Also, let $\hat{\boldsymbol{u}}_i^m \equiv \boldsymbol{u}_i^m(\hat{\boldsymbol{\pi}}^m)$ be the estimated equilibrium expected utilities obtained by plugging in the observed choice probabilities $\hat{\boldsymbol{\pi}}^m$ into the payoffs in game m. Then the sample moment inequalities take the following form: for all cycles of length $\mathcal{L} \in \{2, \ldots, M\}$

$$\sum_{m=G_0}^{G_{\mathcal{L}-1}} \sum_{j=1}^{J_i} \left(\hat{u}_{ij}^{m+1} - \hat{u}_{ij}^m \right) \hat{\pi}_{ij}^m \leq 0 \tag{6}$$

Altogether, in an n-person game we have $n \sum_{\mathcal{L}=2}^{M} \#C(\mathcal{L})$ moment inequalities, where $\#C(\mathcal{L})$ is the number of different (up to a change in the starting game index) cycles of length $\mathcal{L}$. These inequalities make up a necessary condition for a finite sample of games to be QRE-consistent.

[5] Formally, this fact follows from the observation that for games with pure strategy sets having two actions, we can rewrite the functions $\varphi_i(u_i(\boldsymbol{\pi}))$ as $\varphi_i(u_{i1}(\boldsymbol{\pi})) = \varphi_i(u_{i1}(\boldsymbol{\pi}) - u_{i2}(\boldsymbol{\pi}), 0) + u_{i2}(\boldsymbol{\pi})$. Noting that without loss of generality we can normalize $u_{i2}(\boldsymbol{\pi})$ to be constant, we get that $\varphi_i(u(\boldsymbol{\pi}))$ is a univariate function. Using Rochet (1987, Proposition 2) we conclude that if φ_i satisfies (5) for all cycles of length 2, then (5) is also satisfied for cycles of arbitrary length $\mathcal{L} > 2$.

4.1 "Cumulative rank" test as a special case of Cyclic Monotonicity

HHK propose alternative methods of testing the QRE model based on cumulative rankings of choice probabilities across perturbed games,[6] which also imply stochastic equalities or inequalities involving estimated choice probabilities from different games. We will show here that, in fact, our CM conditions are directly related to HHK's rank-cumulative probability conditions in the special case when there are only two games (i.e. all cycles are of length 2), and under a certain non-negativity condition on utility differences between the games.

Formally, HHK consider two perturbed games with the same strategy spaces and re-order strategy indices for each player i such that

$$\tilde{u}^1_{i1} - \tilde{u}^0_{i1} \geq \tilde{u}^1_{i2} - \tilde{u}^0_{i2} \geq \ldots \geq \tilde{u}^1_{iJ_i} - \tilde{u}^0_{iJ_i}$$

where $\tilde{u}^m_{ij} \equiv u_i(s_{ij}, \boldsymbol{\pi}^m_{-i}) - \frac{1}{J_i}\sum_{j=1}^{J_i} u_i(s_{ij}, \boldsymbol{\pi}^m_{-i})$ for $m = 0, 1$. These inequalities can be equivalently rewritten as

$$u^1_{i1} - u^0_{i1} \geq u^1_{i2} - u^0_{i2} \geq \ldots \geq u^1_{iJ_i} - u^0_{iJ_i} \tag{7}$$

HHK's Theorem 2 states that given the indexing in (7) and assuming Invariance (see Assumption 1), QRE consistency implies the following *cumulative rank property:*

$$\sum_{j=1}^{k} (\pi^1_{ij} - \pi^0_{ij}) \geq 0 \quad \text{for all } k = 1, \ldots, J_i. \tag{8}$$

This property is related to our test as the following proposition demonstrates (the proof is in Appendix B):

Proposition 1. *Let $M = 2$. If all expected utility differences in (7) are non-negative, then HHK rank cumulative condition (8) implies the CM inequalities (5). Conversely, the CM inequalities (5) imply HHK condition (8) (without additional assumptions on expected utility differences).*

4.2 Limitations and extensions of the test

Our test makes use of expected payoffs and choice probabilities in a fixed set of $M \geq 2$ games.

To estimate expected payoffs we had to assume that players are risk-neutral. This assumption might be too strong a priori (e.g., Goeree, Holt, & Palfrey (2000) argue that risk aversion can help explain QRE inconsistencies). Notice, however, that the test itself does not depend on risk-neutrality: it only requires that we know the form of the utility function. Thus under additional assumptions about the utility, we can also investigate how risk aversion affects the test results. See Section 6.2 for details.

Furthermore, even when players are risk-neutral, we compute the expected payoffs using observed choice frequencies, which may be different from true choice probabilities used in the CM

[6] HHK also describe another testing approach which utilizes Block-Marschak polynomials involving choice probabilities across games with different choice sets, but this is unrelated to our approach in this paper.

conditions. It is therefore interesting to investigate the set of admissible utility indices/probability distributions for which QRE is not rejected by the CM test. We will consider these possibilities in our simulations and empirical implementation below.

Our test also assumes that for each of the games considered, there is only one unique QRE. Note that since we do not specify the distribution of the random utility shocks, this uniqueness assumption is not verifiable. However, as in much of the recent empirical games literature in industrial organization[7], our testing procedure strictly speaking only assumes that *there is a unique equilibrium played in the data*.[8] Several considerations make us feel that this is a reasonable assumption in our application. First, in our experiments, the subjects are randomly matched across different rounds of each game, so that playing multiple equilibria in the course of an experiment would require a great deal or coordination. Second, as we discuss below, and in Appendix C, all the experimental games that we apply our test to in this paper have a unique QRE under an additional regularity assumption.

Notwithstanding the above discussion, in some potential applications our test may wrongly reject (i.e., it is biased) the QRE null hypothesis when there are multiple quantal response equilibria played in the data. Given the remarks here, our test of QRE should be generally considered a joint test of the QRE hypothesis along with those of risk neutrality of the subjects and unique equilibrium in the data.

5 Econometric implementation of the test

In this section we consider the formal econometric properties of our test. Let $\boldsymbol{\nu} \in \mathbb{R}^P$ denote the vector of the left hand sides of the cyclic monotonicity inequalities (5), written out for all cycle lengths and all players. Here $P \equiv n \sum_{\mathcal{L}=2}^{M} \#C(\mathcal{L})$ and $\#C(\mathcal{L})$ is the number of different (up to the starting game index) cycles of length $\mathcal{L}$. Let us order all players and all different cycles of length $\mathcal{L}$ from 2 to M in a single ordering, and for $\ell \in \{1, \dots, P\}$, let $\mathcal{L}(\ell)$ refer to the cycle length at coordinate number ℓ in this ordering, $m_0(\ell)$ refer to the first game in the respective cycle, and $\iota(\ell)$ refer to the corresponding player at coordinate number ℓ. Then we can write $\boldsymbol{\nu} \equiv (\nu_1, \dots, \nu_\ell, \dots, \nu_P)$ where each generic component ν_ℓ is given by (6), i.e.

$$\nu_\ell = \sum_{m=m_0(\ell)}^{\mathcal{L}(\ell)-1} \sum_{j=1}^{J_{\iota(\ell)}} \sum_{\boldsymbol{s}_{-\boldsymbol{\iota}(\boldsymbol{\ell})}} \pi_{\iota(\ell)j}^{m} \left(\left(\prod_{k \in N_{-\iota(\ell)}} \pi_{k_{j_k}}^{m+1} \right) u_{\iota(\ell)}^{m+1}(s_{\iota(\ell)j}, \boldsymbol{s}_{-\boldsymbol{\iota}(\boldsymbol{\ell})}) - \left(\prod_{k \in N_{-\iota(\ell)}} \pi_{k_{j_k}}^{m} \right) u_{\iota(\ell)}^{m}(s_{\iota(\ell)j}, \boldsymbol{s}_{-\boldsymbol{\iota}(\boldsymbol{\ell})}) \right) \tag{9}$$

Define $\boldsymbol{\mu} \equiv -\boldsymbol{\nu}$, then cyclic monotonicity is equivalent to $\boldsymbol{\mu} \geq \mathbf{0}$. Let $\hat{\boldsymbol{\mu}}$ denote the estimate of

[7] See, e.g., Aguirregabiria & Mira (2007), Bajari et al. (2007).

[8] Indeed, practically all of the empirical studies of experimental data utilizing the quantal response framework assume that a unique equilibrium is played in the data, so that the observed choices are drawn from a homogeneous sampling environment. For this reason, our test may not be appropriate for testing for QRE using field data, which were not generated under these controlled laboratory experimental conditions. See De Paula & Tang (2012) for a test of multiple equilibria presence in the data.

$\boldsymbol{\mu}$ from our experimental data. In our setting, the sampling error is in the choice probabilities π's. Using the Delta method, we can derive that, asymptotically (when the number of trials of each game out of a fixed set of M games goes to infinity),

$$\hat{\boldsymbol{\mu}} \overset{a}{\sim} N(\boldsymbol{\mu}_{\mathbf{0}}, \Sigma) \quad \text{and} \quad \Sigma = JVJ'$$

where V denotes the variance-covariance matrix for the $Mn \times 1$-vector $\boldsymbol{\pi}$ and J denotes the $P \times Mn$ Jacobian matrix of the transformation from $\boldsymbol{\pi}$ to $\boldsymbol{\mu}$. Since $P >> Mn$, the resulting matrix Σ is singular.[9]

We want to perform the hypothesis test:

$$H_0 : \ \boldsymbol{\mu}_{\mathbf{0}} \geq \mathbf{0} \quad vs. \quad H_1 : \ \boldsymbol{\mu}_{\mathbf{0}} \ngeq \mathbf{0},$$

where $\mathbf{0} \in \mathbb{R}^P$. Given the large number of moment inequalities in our test (in the application, $P = 40$), we utilize the Generalized Moment Selection (GMS) procedure of Andrews & Soares (2010). Letting $\hat{\Sigma}$ denote an estimate of Σ, we define the test statistic

$$S(\hat{\mu}, \hat{\Sigma}) := \sum_{\ell=1}^{P} \left[\hat{\mu}^{\ell} / \hat{\sigma}_{\ell}\right]_{-}^{2} \tag{10}$$

where $[x]_{-}$ denotes $x \cdot \mathbb{1}(x < 0)$, and $\hat{\sigma}_1^2, \ldots, \hat{\sigma}_P^2$ denote the diagonal elements of $\hat{\Sigma}$. The test statistic is sum of squared violations across the moment inequalities.

The general intuition of the GMS procedure is to evaluate the asymptotic distribution of the test statistic under a sequence of parameters under the null hypothesis which resemble the sample moment inequalities, and are drifting to zero. By doing this, moment inequalities which are far from binding in the sample (i.e. the elements of $\hat{\mu}$ which are $>> 0$) will not contribute to the asymptotic distribution of the test statistic, leading to a (stochastically) smaller distribution and hence smaller critical values.[10]

To obtain valid critical values for S under H_0, we use the following procedure from Andrews & Soares (2010):

1. Let $D \equiv Diag^{-1/2}(\hat{\Sigma})$ denote the diagonal matrix with elements $1/\hat{\sigma}_1, \ldots, 1/\hat{\sigma}_P$. Compute $\Omega \equiv D \cdot \hat{\Sigma} \cdot D$.

2. Compute the vector $\xi = \kappa_K^{-1} \cdot D \cdot \hat{\mu}$ which is equal to $\frac{1}{\kappa_K} \cdot \left[\frac{\hat{\mu}^1}{\hat{\sigma}_1}, \frac{\hat{\mu}^2}{\hat{\sigma}_2}, \cdots, \frac{\hat{\mu}^P}{\hat{\sigma}_P}\right]'$ where $\kappa_K =$

[9] Note also that Σ is the approximation of the *finite-sample* covariance matrix, so that the square-roots of its diagonal elements correspond to the standard errors; i.e. the elements are already "divided through" by the sample size, which accounts for the differences between the equations below and the corresponding ones in Andrews & Soares (2010).

[10] In contrast, other inequality based testing procedure (e.g., Wolak (1989)) evaluate the asymptotic distribution of the test statistic under the "least-favorable" null hypothesis $\mu = 0$, which may lead to very large critical values and low power against alternative hypotheses of interest.

$(\log K)^{1/2}$.[11] Here $K = N/M$ is the number of trials in each of the M games, and N is the sample size.

3. For $r = 1, \ldots, R$, we generate $Z_r \sim N(0, \Omega)$ and compute $s_r \equiv S(Z_r + [\xi]_+, \Omega)$, where $[x]_+ = \max(x, 0)$.

4. Take the critical value $c_{1-\alpha}$ as the $(1 - \alpha)$-th quantile among $\{s_1, s_2, \ldots, s_R\}$.

Essentially, the asymptotic distribution of S is evaluated at the null hypothesis $[\xi]_+ \geq 0$ which, because of the normalizing sequence κ_K, is drifting towards zero. In finite samples, this will tend to increase the number of rejections relative to evaluating the asymptotic distribution at the zero vector. This is evident in our simulations, which we turn to next, after describing the test set of games.

5.1 Test set of two-player games: "Joker" games

As test games, we used a series of four card-matching games where each player has three choices. These games are so-called "Joker" games which have been studied in the previous experimental literature (cf. O'Neill (1987) and Brown & Rosenthal (1990)), and can be considered generalizations of the familiar "matching pennies" game in which each player calls out one of three possible cards, and the payoffs depend on whether the called-out cards match or not. Since these games will also form the basis for our laboratory experiments below, we will describe them in some detail here.[12] Table 1 shows the payoff matrices of the four games which we used in our simulations and experiments.

Each of these games has a unique mixed-strategy Nash equilibrium, the probabilities of which are given in bold font in the margins of the payoff matrices. Note that the four games in Table 1 differ only by Row player's payoff. Nash equilibrium logic, hence, dictates that the Row player's equilibrium choice probabilities never change across the four games, but that the Column player should change her mixtures to maintain the Row's indifference amongst choices.

Joker games have unique regular QRE. An important advantage of using Games 1–4 for our application is that in each of our games there is a unique QRE for *any* regular quantal response function (see Appendix C for details). Regular QRE is an extremely important class of QRE, so let us briefly describe the additional restrictions regularity imposes on the admissible quantal response functions.[13] In this paper we focus on testing QRE via its implication of cyclic monotonicity, which involves checking testable restrictions on QRE probabilities when the shock distribution is fixed in a series of games that only differ in the payoffs. These are comparisons *across games*. In the QRE literature, there are typically additional restrictions imposed on quantal response functions *within a*

[11] Andrews & Soares (2010) mention several alternative choices for κ_K. We investigate their performance in the Monte Carlo simulations reported below.

[12] In our choice of games, we wanted to use simple games comparable to games from the previous literature. Moreover, we wanted our test to be sufficiently powerful. This last consideration steered us away from games for which we know that some structural QRE (in particular, the Logit QRE) performs very well so that the chance to fail the CM conditions is pretty low.

[13] See Appendix C for formal definitions and additional details.

Table 1: Four 3×3 games inspired by the Joker Game of O'Neill (1987).

			1	2	J
Game 1 (Symmetric Joker)			**[1/3]** (.325)	**[1/3]** (.308)	**[1/3]** (.367)
	1	**[1/3]** (.273)	10, 30	30, 10	10, 30
	2	**[1/3]** (.349)	30, 10	10, 30	10, 30
	J	**[1/3]** (.378)	10, 30	10, 30	30, 10
Game 2 (Low Joker)			1	2	J
			[9/22] (.359)	**[9/22]** (.439)	**[4/22]** (.202)
	1	**[1/3]** (.253)	10, 30	30, 10	10, 30
	2	**[1/3]** (.304)	30, 10	10, 30	10, 30
	J	**[1/3]** (.442)	10, 30	10, 30	55, 10
Game 3 (High Joker)			1	2	J
			[4/15] (.258)	**[4/15]** (.323)	**[7/15]** (.419)
	1	**[1/3]** (.340)	25, 30	30, 10	10, 30
	2	**[1/3]** (.464)	30, 10	25, 30	10, 30
	J	**[1/3]** (.196)	10, 30	10, 30	30, 10
Game 4 (Low 2)			1	2	J
			[2/5] (.487)	**[1/5]** (.147)	**[2/5]** (.366)
	1	**[1/3]** (.473)	20, 30	30, 10	10, 30
	2	**[1/3]** (.220)	30, 10	10, 30	10, 30
	J	**[1/3]** (.307)	10, 30	10, 30	30, 10

Notes. For each game, the unique Nash equilibrium choice probabilities are given in bold font within brackets, while the probabilities in regular font within parentheses are aggregate choice probabilities from our experimental data, described in Section 6.1.

fixed game. In particular, the quantal response functions studied in Goeree, Holt, & Palfrey (2005), in addition to the assumptions we impose in Section 2, also satisfy the *rank-order property*[14], which states that actions with higher expected payoffs are played with higher probability than actions with lower expected payoffs. Formally, a quantal response function $\pi_i : \mathbb{R}^{J_i} \to \Delta(S_i)$ satisfies the *rank-order property* if for all $i \in N, j, k \in \{1, \dots, J_i\}$:

$$u_{ij}(\mathbf{p}) > u_{ik}(\mathbf{p}) \Rightarrow \pi_{ij}(\mathbf{p}) > \pi_{ik}(\mathbf{p}). \tag{11}$$

We stress here that while the rank-order property is not assumed for our test, it is nevertheless a reasonable assumption for QRE, and the inequalities (11) can be tested using the same formal statistical framework we described in the previous sections. See Appendix C for details.

Monte Carlo simulations. For the Monte Carlo simulations, we first considered artificial data generated under the QRE hypothesis (specifically, under a logit QRE model). Table 2 shows the results of the GMS test procedure applied to our setup in terms of the number of rejections. From Table 2, we see that the test tends to (slightly) underreject under the QRE null for most values of the tuning parameter κ_K. The results appear relatively robust to changes in κ_K; a reasonable choice appears to be $\kappa_K = 5(\log(K))^{\frac{1}{4}}$, which we will use in our experimental results below. In

[14] We borrow the name of this property from Fox (2007). Goeree, Holt, & Palfrey (2005) call this a "monotonicity" property, but we chose not to use that name here to avoid confusion with "cyclic monotonicity", which is altogether different.

a second set of simulations, we generated artificial data under a non-QRE play (specifically, we generated a set of choice probabilities that generate violations of all of the CM inequalities for both players). The results here are quite stark: in all our simulations, and for all the tuning parameters that we checked, we find that the QRE hypothesis is rejected in every single replication. Thus our proposed test appears to have very good power properties.

Table 2: Monte Carlo simulation results under QRE-consistent data

		Tuning parameter κ_K			
N	# rejected[a] at	$5(\log(K))^{\frac{1}{2}}$	$5(\log(K))^{\frac{1}{4}}$	$5(\log(K))^{\frac{1}{8}}$	$5(2\log\log(K))^{\frac{1}{2}}$
1000	5%	7	9	13	7
	10%	13	17	26	14
	20%	33	54	68	46
5000	5%	13	32	43	21
	10%	26	55	77	41
	20%	61	102	124	85
9000	5%	18	40	51	32
	10%	33	61	79	53
	20%	65	114	143	97

Notes. $K = \frac{N}{4}$ is the total number of rounds of each of the four games. All numbers in columns 3–6 are observed rejections out of 500 replications. All computations use $R = 1000$ to simulate the corresponding critical values.
[a]: # rejected out of 500 replications for each significance level.

6 Experimental evidence

In this section we describe an empirical application of our test to data generated from laboratory experiments. Lab experiments appear ideal for our test because the invariance of the distribution of utility shocks across games (Assumption 1) may be more likely to hold in a controlled lab setting than in the field. This consideration also prevented us from using the experimental data from published 3×3 games, as those usually do not have the same subjects participating in several (we require up to 4) different games in one session.

Our testing procedure can be applied to the experimental data from Games 1–4 as follows. As defined previously, let the P-dimensional vector $\boldsymbol{\nu}$ contain the value of the CM inequalities evaluated at the choice frequencies observed in the experimental data. Using our four games, we can construct cycles of length $2, 3$, and 4. Thus we have 12 possible orderings of 2-cycles, 24 possible orderings of 3-cycles, and 24 possible orderings of 4-cycles. Since CM inequalities are invariant to the change of the starting game index, it is sufficient to consider the following 20 cycles of Games $1 - 4$:

$$\begin{aligned}
&121, 131, 141, 232, 242, 343\\
&1231, 1241, 1321, 1341, 1421, 1431, 2342, 2432\\
&12341, 12431, 13241, 13421, 14231, 14321 \qquad (12)
\end{aligned}$$

Moreover, these 20 cycles are distinct depending on whether we are considering the actions facing the Row or Column player (which involve different payoffs); thus the total number of cycles across the four games and the two player roles is $P = 40$. This is the number of coordinates of vector $\boldsymbol{\nu}$, defined in (9). Additional details on the implementation of the test, including explicit expressions for the variance-covariance matrix of ν, are provided in Appendix D.

6.1 Experimental design

The subjects in our laboratory experiments were undergraduate students at the University of California, Irvine, and all experiments were conducted at the ESSL lab there. We have conducted a total of 3 sessions where in each session subjects played one the following sequences of games from Table 1: 12, 23, and 3412. In the first two sessions, subjects played 20 rounds per game; in the last session subjects played 10 rounds per game.[15] Across all three sessions, there was a total of 96 subjects. To reduce repeated game effects, subjects were randomly rematched each round. To reduce framing effects, the payoffs for every subject were displayed as payoffs for the Row player, and actions were abstractly labelled A, B, and C for the Row player, and D, E, and F for the Column player.

In addition to recording the actual choice frequencies in each round of the game, we periodically also asked the subjects to report their beliefs regarding the likelihood of their current opponent playing each of the three strategies. Each subject was asked this question once s/he had chosen her action but before the results of the game were displayed. To simplify exposition, we used a two-thumb slider which allowed subjects to easily adjust the probability distribution among three choices. Thus we were able to compare the CM tests based on subjective probability estimates with the ones based on actual choices.[16]

Subjects were paid the total sum of payoffs from all rounds, exchanged into U.S. dollars using the exchange rate of 90 cents for 100 experimental currency units, as well as a show-up fee of \$7. The complete instructions of the experiment are provided in Appendix E.

6.2 Results

We start analyzing the experimental data by reporting the aggregate choice frequencies in Games 1–4 in Table 1 alongside Nash equilibrium predictions. Comparing theory with the data, we see that there are a lot of deviations from Nash for both Column and Row players.

Table 3 is our main results table. It shows the test results of checking the cyclic monotonicity conditions with our experimental data.

Based on our current dataset, we find that QRE is soundly rejected for the pooled data (with test statistic 68.194 and 5% critical value 29.985). This may not be too surprising, since in our design subjects experience both player roles (Row and Column), and so this pooled test imposes

[15] We had to adjust the number of rounds because of the timing constraints.

[16] We chose not to incentivize belief elicitation rounds largely to avoid imposing extra complexity on the subjects. Thus our results using elicited belief estimates should be taken with some caution. On the other hand, if what we elicited was completely meaningless, we would not observe as much QRE consistency as we do in our subject-by-subject results below.

Table 3: Testing for Cyclic Monotonicity in Experimental Data: Generalized Moment Selection

Data sample	*AS* **test stat**	$c^R_{0.95}$				
All subjects pooled:						
All cycles	68.194	29.985				
Row cycles	68.194	26.265				
Col cycles	0.000	6.835				
Subject-by-subject:	**Avg** *AS*	**Avg** $c^R_{0.95}$	**# rejected**			**Avg CM violations**
			at 5%	at 10%	at 20%	(% of total)
Subj. v. self[a]						
(Total subj.: 96)						
All cycles	212.570	18.538	29	37	41	41.59
Row cycles	203.108	12.421	20	26	31	44.53
Column cycles	9.462	11.849	18	19	21	38.65
Subj. v. others[b]						
(Total subj.: 96)						
Row cycles	3.936	10.920	7	11	15	35.00
(Total subj.: 96)						
Col cycles	103.872	11.734	16	18	22	38.70
Subj. v. beliefs[c]						
(Total subj.: 59)						
Row cycles	3.681	6.883	5	5	5	33.051
(Total subj.: 61)						
Col cycles	9.622	8.732	10	13	14	35.000

Notes. All computations use $R = 1,000$. In subject-by-subject computations some subjects in some roles exhibited zero choice variance, so in those cases we replaced the corresponding (ill-defined) elements of $Diag^{-1/2}(\hat{\Sigma})$ with ones and when computing the test statistic, left out the corresponding components of $\hat{\mu}$. The tuning parameter in AS procedure was set equal to $\kappa_z = 5(\log(z))^{\frac{1}{4}}$.

[a]v. self: the opponent's choice frequencies are obtained from the same subject playing the respective opponent's role.

[b]v. others: the opponent's choice frequencies are averages over the subject's actual opponents' choices when the subject was playing her respective role in column (ii).

[c]v. beliefs: the opponent's choice frequencies are averages over the subject's elicited beliefs about the opponent choices when the subject was playing the respective role (since belief elicitation rounds were fixed at the session level, subjects' beliefs may not be elicited in some roles and some games. We dropped them from the analysis).

the auxiliary assumption on all subjects being homogeneous across roles in that their utility shocks are drawn from identical distributions.

Therefore, in the remaining portion of Table 3, we test the QRE hypothesis separately for different subsamples of the data. First, we consider separately the CM inequalities pertaining to Row players and those pertaining to Column players.[17] By doing this, we allow the utility shock distributions to differ depending on a subject's role (but conditional on role, to still be identical across subjects).

We find that while we still reject QRE[18] for the Row players, we cannot do so for Column players. Thus overall QRE-inconsistency is largely due to the behavior of the Row players. Seeing that Row players' inequalities are violated more often than Column players' inequalities suggests that Row players' choice probabilities do not always adjust toward higher-payoff strategies. That the violations come predominantly from the choices of Row players is interesting because, as we discussed above, the Column players' payoffs are the same across all the games in our experiment, but vary across games for the Row player.

Continuing in this vein, the lower panel of Table 3 considers tests of the QRE hypothesis for each subject individually. Obviously, this allows the distributions of the utility shocks to differ across subjects. For these subject-by-subject tests, there is a question about how to determine a given's subject beliefs about her opponents' play. We consider three alternatives: (i) set beliefs about opponents equal to the subject's own play in the opponent's role; (ii) set beliefs about opponents equal to opponents' actual play (i.e., as if the subject was playing against an average opponent); and (iii) set beliefs about opponents equal to the subject's elicited beliefs regarding the opponents' play.

The results appear largely robust across these three alternative ways of setting subjects' beliefs. We see that we are not able to reject the QRE hypothesis for most of the subjects, for significance levels going from 5% to 20%. When we further break down each subject's observations depending on his/her role (as Column or Row player), thus allowing the utility shock distributions to differ not only across subjects but also for each subject in each role, the number of rejections decreases even more. Curiously, we see that in the subject vs. self results, the Row inequalities generate more violations, while the Column inequalities generate more violations in the subject vs. others results.

One caveat here, is that when we are testing on a subject-by-subject basis, we are, strictly speaking, no longer testing an equilibrium hypothesis, because we are not testing – and indeed, *cannot* test given the randomized pairing of subjects in the experiments – whether the given subject's opponents are playing optimally according to a QRE. Hence, our tests should be interpreted as tests of subjects' "best response" behavior given beliefs about how their opponents' play.

[17]Note that the sum of the test statistics corresponding to the Column and Row inequalities sum up to the overall test statistic; this is because the Row and Column inequalities are just subsets of the full set of inequalities.

[18]Strictly speaking, this is no longer a test of QRE, because by restricting attention to cycles pertaining to only one player role, we essentially consider only one-player equilibrium version of QRE, which is more akin to a discrete choice problem.

The general trend of these findings – that the QRE hypothesis appears more statistically plausible once we allow for sufficient heterogeneity across subjects and across roles – confirms existing results in McKelvey, Palfrey, & Weber (2000) who, within the parametric logit QRE framework and 2×2 asymmetric matching pennies, also found evidence increasing for the QRE hypothesis once subject-level heterogeneity was accommodated.

Robustness check: Nonlinear utility and risk aversion. Our test results above are computed under the assumption of risk-neutrality. Goeree, Holt, & Palfrey (2000) have shown that allowing for nonlinear utility (i.e. risk aversion) greatly improves the fit of QRE to experimental evidence. Since our test can be applied under quite general specification of payoff functions, to see the effects of risk aversion on the test results we recomputed the test statistics under an alternative assumption that for each player, utility from a payoff of x is $u(x) = x^{1-r}$, where $r \in [0, 1)$ is a constant relative risk aversion factor.[19] Here, we computed the test statistics and critical values for values of r ranging from 0 to 0.99.

When we pool all the subjects together, we find results very similar to what is reported in Table 3: QRE is rejected when all cycles are considered; it is also rejected when only the Row cycles are considered; it cannot be rejected when only the Column cycles are considered, for all values of $r \in [0, 0.99]$. Thus we do not observe any risk effects in the pooled data.[20]

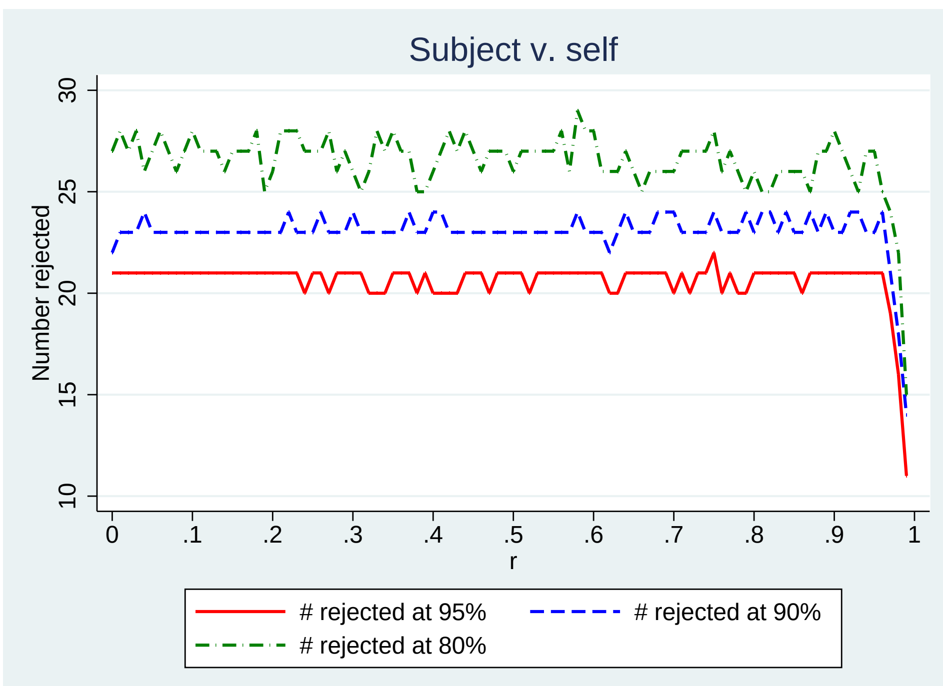

Figure 1: Effects of risk aversion on subject-by-subject rejections

Breaking down these data on a subject-by-subject basis, we once again see that allowing for

[19]For $r = 1$ the log-utility form is used. In our computations, we restrict the largest value of r to 0.99 to avoid dealing with this issue.

[20]For space reasons, we have not reported all the test statistics and critical values, but they are available from the authors upon request.

risk aversion does not change our previous results obtained under the assumption of linear utility. Specifically, as graphed in Figure 1, the number of rejections of the QRE hypothesis for the "subject vs. self" specification is relatively stable for all $r < 0.99$, staying at about 21 rejections at 5% level, at about 23 rejections at 10% level, and between 25 and 30 rejections at 20% level. Thus our analysis here suggests the our test results are not driven by risk aversion.

More generally, risk aversion might be an important factor in other games, so checking for the potential effects of risk aversion on test results might be a necessary post-estimation step.

7 Conclusions and Extensions

In this paper we present a new approach for testing the QRE hypothesis in finite normal form games. The testing approach is based on moment inequalities derived from the *cyclic monotonicity* condition, which is in turn derived from the convexity of the random utility model underlying the QRE hypothesis. We investigate the performance of our test using a lab experiment where subjects play a series of generalized matching pennies games.

While we primarily focus on developing a test of the QRE hypothesis in games involving two or more players, our procedure can also be applied to situations of stochastic individual choice. Thus our test can be viewed more generally as a semiparametric test of quantal response, and, in particular, discrete choice models. Moreover, in finite action games as considered here, QRE has an identical structure to Bayesian Nash equilibria in discrete games of incomplete information which have been considered in the empirical industrial organization literature (e.g. Bajari et al. (2010), De Paula & Tang (2012), or Liu, Vuong, & Xu (2013)). Our approach can potentially be useful for specification testing in those settings; however, as we remarked above, one hurdle to implementing such tests on field data is the possibility of multiple equilibria played in the data. Adapting these tests to allow for multiple equilibria is a challenging avenue for future research.

References

V. Aguirregabiria, and P. Mira (2007). Sequential Estimation of Dynamic Discrete Games, *Econometrica*, 75(1): 1–53.

D. Andrews and G Soares (2010). Inference for Parameters Defined by Moment Inequalities using Generalized Moment Selection, *Econometrica*, 78(1): 119-157.

A. Archer and R. Kleinberg (2014). Truthful germs are contagious: A local to global characterization of truthfulness, *Games and Economic Behavior*, 86: 340–366.

I. Ashlagi, M. Braverman, A. Hassidim, and D. Monderer (2010). Monotonicity and Implementability, *Econometrica*, 78(5): 1749–1772

P. Bajari, L. Benkard, and J. Levin (2007). Estimating Dynamic Models of Imperfect Competition, *Econometrica*, 75: 1331-1370.

P. Bajari, H. Hong, J. Krainer, and D. Nekipelov (2010). Estimating Static Models of Strategic Interactions, *Journal of Business and Economic Statistics*, 28: 469-482.

J. N. Brown and R.W. Rosenthal (1990). Testing the Minimax Hypothesis: A Re-Examination of O'Neill's Game Experiment, *Econometrica*, 58: 1065-1081.

R. Cominetti, E. Melo, and S. Sorin, A payoff-based learning procedure and its application to traffic games. *Games and Economic Behavior*, 70: 71–83.

A. de Paula and X. Tang (2012). Inference of Signs of Interaction Effects in Simulataneous Games with Incomplete Information, *Econometrica*, 80(1): 143-172.

I. Erev and A.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.

D. Fudenberg and D. Levine. The theory of learning in games, MIT Press: Cambridge, MA, 1998.

M. Fosgerau & A. de Palma (2015). Demand Systems for Market Shares. *MPRA Paper No. 62106*.

J. Fox (2007). Semiparametric Estimation of Multinomial Discrete-Choice Models using a Subset of Choices. *RAND Journal of Economics*, 38: 1002–1029.

J. Goeree and C. Holt (2001). Ten Little Treasures of Game Theory and Ten Intuitive Contradictions. *American Economic Review*, 91(5):1402–1422.

J. Goeree, C. Holt, and T. Palfrey (2000). Risk Averse Behavior in Asymmetric Matching Pennies Games. Working paper, University of Virginia.

J. Goeree, C. Holt, and T. Palfrey (2003). Risk Averse Behavior in Generalized Matching Pennies Games, *Games and Economic Behavior*, 45:97-113.

J. Goeree, C. Holt, and T. Palfrey (2005). Regular Quantal Response Equilibrium, *Experimental Economics*, 8(4): 347-67.

P. Haile, A. Hortacsu, and G. Kosenok (2008). On the Empirical Content of Quantal Response Models. *American Economic Review*, 98:180–200.

R. Lavi and C.Swamy (2009), Truthful mechanism design for multidimensional scheduling via cycle monotonicity, *Games and Economic Behavior*, 67: 99-124.

N. Liu, Q. Vuong, and H. Xu (2013). Rationalization and Identification of Discrete Games with Correlated Types. Working paper.

D. McFadden (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In: P. Zarembka (Ed), Frontiers in econometrics, NY: Academic Press, 105–142.

D. McFadden (1981). Econometric Models of Probabilistic Choice. In: C.Manski and D. McFadden (Eds), Structural Analysis of Discrete Data with Economic Applications, Cambridge, MA: MIT Press, 198–272.

D. McFadden and M. Fosgerau (2012). A Theory of the Perturbed Consumer with General Budgets. *NBER Working paper No. 17953.*

R. D. McKelvey and T. R. Palfrey (1995). Quantal response equilibria in normal form games, *Games and Economic Behavior*, vol. 7, pp. 6–38.

R. D. McKelvey, T. R. Palfrey, and R. 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: 523-548.

B. O'Neill (1987). Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences*, 84: 2106-2109.

P. Mertikopoulos and W. Sandholm (2014). Regularized best response and reinforcement learning in games, *Working Paper*.

J-C. Rochet (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*,16: 191-200

R. Rockafellar (1970). *Convex Analysis*. Princeton University Press.

J. Rust (1994). Structural Estimation of Markov Decision Processes. In *Handbook of Econometrics*, Vol. 4, edited R. Engle and D. McFadden, pp. 3082, 3146. North-Holland.

B. W. Rogers, T. R. Palfrey, and C.C. Camerer (2009). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144: 1440-1467.

M. Saks and L. Yu (2005). Weak monotonicity suffices for truthfulness on convex domains. in *ACM Conference on Electronic Commerce*. Vancouver, Canada: ACM, 286–293.

R. Vohra (2011). *Mechanism Design: a Linear Programming Approach*. Cambridge University Press.

F. Wolak (1989). Testing Inequality Constraints in Linear Econometric Models, *Journal of Econometrics*, 41: 205–235.

Appendix: supplemental material (not for publication)

A Cyclic Monotonicity and Utility Maximization

In this section we show that the CM inequalities are equivalent to players' utility maximization. In order to establish this result we exploit the fact that the set of QRE can be seen as the set of NE of a perturbed game. In particular, it can be shown[21] that the set of QRE corresponds to the set of NE of a game where players' payoffs are given by

$$\mathcal{G}^i(\boldsymbol{\pi}) := \left\langle \boldsymbol{\pi}_i, \boldsymbol{u}_i \right\rangle - \tilde{\varphi}^i(\boldsymbol{\pi}_i), \quad \forall i \in N,$$

with $\tilde{\varphi}^i(\boldsymbol{\pi}_i)$ corresponding to the Fenchel-Legendre conjugate of the function $\varphi^i(\boldsymbol{u}_i)$ defined in (2).

Utility maximization implies CM: Consider a cycle of length $\mathcal{L}-1$ with $[\boldsymbol{u}_i]^m$ and $[\boldsymbol{\pi}_i^*]^m$ denoting the expected payoffs and equilibrium probabilities respectively. By utility maximization, it is easy to see that for each player i the following inequalities must hold:

$$\left\langle [\boldsymbol{u}_i]^{m+1}, [\boldsymbol{\pi}_i^*]^m \right\rangle - \tilde{\varphi}^i([\boldsymbol{\pi}_{\boldsymbol{i}}^*]^m) \leq \left\langle [\boldsymbol{u}_i]^{m+1}, [\boldsymbol{\pi}_i^*]^{m+1} \right\rangle - \tilde{\varphi}^i([\boldsymbol{\pi}_i^*]^{m+1}), \quad \forall m. \tag{13}$$

Rewriting as

$$\left\langle [\boldsymbol{u}_i]^{m+1} - [\boldsymbol{u}_i]^m, [\boldsymbol{\pi}_i^*]^m \right\rangle \leq \left\langle [\boldsymbol{u}_i]^{m+1}, [\boldsymbol{\pi}_i^*]^{m+1} \right\rangle - \left\langle [\boldsymbol{u}_i]^m, [\boldsymbol{\pi}_i^*]^m \right\rangle + \tilde{\varphi}^i([\boldsymbol{\pi}_{\boldsymbol{i}}^*]^m) - \tilde{\varphi}^i([\boldsymbol{\pi}_i^*]^{m+1}),$$

and adding up over the cycle, we get the CM inequalities (4).

CM Implies utility maximization: Suppose that CM holds. This implies that $\boldsymbol{\pi}_i \in \partial\varphi^i(\boldsymbol{u}_i)$. Let $\tilde{\mathcal{G}}^i$ denote the convex conjugate of $\mathcal{G}^i$. Noting that $\varphi^i(\boldsymbol{u}_i)$ is the convex conjugate of $\mathcal{G}^i(\boldsymbol{\pi})$, we obtain that $\pi_i \in \partial\tilde{\mathcal{G}}^i(\boldsymbol{u}_i)$. Thanks to Corollary 23.5.1 in Rockafellar (1970) we know that $\boldsymbol{\pi}_i \in \partial\tilde{\mathcal{G}}^i(\boldsymbol{u}_i)$ iff $\boldsymbol{u}_i \in \partial\mathcal{G}^i(\boldsymbol{\pi})$. The latter fact is equivalent to saying that $\boldsymbol{\pi}_i$ maximizes $\mathcal{G}^i(\boldsymbol{\pi})$. Then the conclusion follows. □

Intuitively, the set of inequalities (13) can be seen as set of incentive compatibility constraints across games that only differ in the payoffs. This means that our CM conditions capture players' optimization behavior with respected to changes in the expected payoffs across such games.

B Proof of Proposition 1

Suppose there are two games that differ only in the payoffs. For $M = 2$, the cyclic monotonicity condition (5) reduces to

$$\sum_{j=1}^{J_i}(u_{ij}^1 - u_{ij}^0)\pi_{ij}^0 + \sum_{j=1}^{J_i}(u_{ij}^0 - u_{ij}^1)\pi_{ij}^1 \leq 0$$

or, equivalently,

$$\sum_{j=1}^{J_i}(u_{ij}^1 - u_{ij}^0)(\pi_{ij}^0 - \pi_{ij}^1) \leq 0 \tag{14}$$

Suppose that the RHS of (7) is non-negative. Then HHK condition (8) implies CM. To see this, notice that for non-negative utilities differences in (7)

$$(u_{i1}^1 - u_{i1}^0)(\pi_{i1}^0 - \pi_{i1}^1) \leq 0$$

by HHK condition for $k = 1$. Then

$$\begin{aligned}
&(u_{i2}^1 - u_{i2}^0)(\pi_{i2}^0 - \pi_{i2}^1) + (u_{i1}^1 - u_{i1}^0)(\pi_{i1}^0 - \pi_{i1}^1) \leq \\
&(u_{i2}^1 - u_{i2}^0)(\pi_{i2}^0 - \pi_{i2}^1) + (u_{i2}^1 - u_{i2}^0)(\pi_{i1}^0 - \pi_{i1}^1) = \\
&\qquad (u_{i2}^1 - u_{i2}^0)((\pi_{i1}^0 + \pi_{i2}^0) - (\pi_{i1}^1 + \pi_{i2}^1)) \leq 0
\end{aligned}$$

[21] See, e.g., Cominetti, Melo, & Sorin (2010) and Mertikopoulos & Sandholm (2014).

where the last inequality follows from HHK condition for $k = 2$ and $u^1_{i2} - u^0_{i2} \geq 0$. Repeating the same procedure for $k = 3, \ldots, J_i$, we obtain the CM condition (14) for $M = 2$.

Conversely, suppose that (14) holds. For the case of two games, (14) holding for all players is necessary and sufficient to generate QRE-consistent choices. All premises are satisfied for HHK's Theorem 2, so condition (8) follows. One can also show it directly. Clearly, given (14), we can always re-label strategy indices so that (7) holds. Let $k = 1$ and by way of contradiction, suppose that (8) is violated, i.e. $\pi^1_{i1} - \pi^0_{i1} < 0$. Since (14) holds, the probabilities in both games are generated by a QRE. Due to indexing in (7),

$$u^1_{i1} - u^1_{ij} \geq u^0_{i1} - u^0_{ij}$$

for all $j > 1$. But then by definition of QRE in (1), $\pi^1_{i1} \geq \pi^0_{i1}$. Contradiction, so (8) holds for $k = 1$. By induction on the strategy index, one can show that (8) holds for all $k \in \{1, \ldots, J_i\}$. This completes the proof. □

C Uniqueness of Regular QRE in Experimental Joker Games

In this section we show formally that any regular QRE in Games 1–4 from Table 1 is unique. We start by recalling the necessary definitions.

A quantal response function $\pi_i : \mathbb{R}^{J_i} \to \Delta(S_i)$ is *regular*, if it satisfies Interiority, Continuity, Responsiveness, and Rank-order[22] axioms (Goeree, Holt, & Palfrey, 2005, p.355). Interiority, Continuity, and Responsiveness are satisfied automatically under the *structural approach* to quantal response[23] that we pursue in this paper as long as the shock distributions have full support. Importantly, for some shock distributions this approach may fail to satisfy the Rank-order Axiom, i.e. the intuitive property of QRE saying that actions with higher expected payoffs are played with higher probability than actions with lower expected payoffs. For the sake of convenience, we repeat the axiom here.

A quantal response function $\pi_i : \mathbb{R}^{J_i} \to \Delta(S_i)$ satisfies *Rank-order Axiom* if for all $i \in N, j, k \in \{1, \ldots, J_i\}$ $u_{ij}(\mathbf{p}) > u_{ik}(\mathbf{p}) \Rightarrow \pi_{ij}(\mathbf{p}) > \pi_{ik}(\mathbf{p})$.

Notice that the Rank-order Axiom involves comparisons of expected payoffs from choosing different pure strategies *within* a fixed game. As briefly discussed in Section 4.2, consistency of the data with the Rank-order Axiom can be tested: the Axiom is equivalent to the following inequality for each player i and pair of i's strategies $j, k \in \{1, \ldots, J_i\}$:

$$(u_{ij}(\mathbf{p}) - u_{ik}(\mathbf{p}))(\pi_{ij}(\mathbf{p}) - \pi_{ik}(\mathbf{p})) \geq 0 \tag{15}$$

Thus a modified test for consistency with a *regular* QRE involves two stages: first, check if the data are consistent with a structural QRE using the cyclic monotonicity inequalities (which compare choices across games) as described in Section 5, and second, if the test does not reject the null hypothesis of consistency, check if the Rank-order Axiom (which compares choices within a game) holds by estimating (15) for each game.[24]

Alternatively, the Rank-order Axiom can be imposed from the outset by making an extra assumption about the shock distributions. In particular, Goeree, Holt, & Palfrey (2005, Proposition 5) shows that under the additional assumption of exchangeability, the quantal response functions derived under the structural approach are regular.

Notice that Assumption 1 (Invariance) is not required for the test of the Rank-order Axiom. In theory, we may have cases where the data can be rationalized by a structural QRE that fails Rank-order, by a structural QRE that satisfies it (i.e., by a regular QRE), or by quantal response functions that satisfy Rank-order in each game but violate the assumption of fixed shock distributions across games. In the latter case, checking consistency with other boundedly rational models (e.g., Level-k or Cognitive Hierarchy) becomes a natural follow-up step.

We can now turn to the uniqueness of the regular QRE in our test games.

[22] This property is called Monotonicity in Goeree, Holt, & Palfrey (2005).

[23] In this approach, the quantal response functions are derived from the primitives of the model with additive payoff shocks, as described in Section 2.

[24] The test procedure is similar to the one in Section 5, with an appropriately modified Jacobian matrix.

Proposition 2. *For each player $i \in N \equiv \{Row, Col\}$ fix a regular quantal response function $\pi_i : \mathbb{R}^{J_i} \to \Delta(S_i)$ and let $Q \equiv (\pi_i)_{i\in N}$. In each of Games 1–4 from Table 1 there is a unique quantal response equilibrium (σ^R, σ^C) with respect to Q. Moreover, in Game 1, the unique quantal response equilibrium is $\sigma^R = \sigma^C = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. In Game 2, $\sigma_1^R = \sigma_2^R \in \left(0, \frac{1}{3}\right)$ and $\sigma_1^C = \sigma_2^C \in \left(\frac{1}{3}, \frac{9}{22}\right]$. In Game 3, $\sigma_1^R = \sigma_2^R \in \left(\frac{1}{3}, 1\right)$ and $\sigma_1^C = \sigma_2^C \in \left[\frac{4}{15}, \frac{1}{3}\right)$. In Game 4, $\sigma_2^R = \sigma_J^R \in \left(0, \frac{1}{3}\right)$ and $\sigma_1^C = \sigma_J^C \in \left(\frac{1}{3}, \frac{2}{5}\right]$.*

Notice that the equilibrium probability constraints in Proposition 2 hold in any regular QRE, not only logit QRE. For the logit QRE they hold for any scale parameter $\lambda \in [0, \infty)$.

Proof. In order to prove uniqueness and bounds on QRE probabilities we will be mainly using Rank-order and Responsiveness properties of a regular QRE.

Suppose Row plays $\sigma^R = (\sigma_1^R, \sigma_2^R, \sigma_J^R)$, Col plays $\sigma^C = (\sigma_1^C, \sigma_2^C, \sigma_J^C)$, and (σ^R, σ^C) is a regular QRE.[25] Expected utility of Col from choosing each of her three pure strategies in any of Games 1–4 (see payoffs in Table 1) is

$$u_{C1}(\sigma^R) = 30 - 20\sigma_2^R$$
$$u_{C2}(\sigma^R) = 30 - 20\sigma_1^R$$
$$u_{CJ}(\sigma^R) = 10 + 20\sigma_1^R + 20\sigma_2^R$$

Consider Game 1. Expected utility of Row in this game from choosing each of her three pure strategies is

$$u_{R1}(\sigma^C) = 10 + 20\sigma_2^C$$
$$u_{R2}(\sigma^C) = 10 + 20\sigma_1^C$$
$$u_{RJ}(\sigma^C) = 30 - 20\sigma_1^C - 20\sigma_2^C$$

Consider Row's equilibrium strategy. There are two possibilities: 1) $\sigma_1^R > \sigma_2^R$. Then Rank-order applied to Col implies $\sigma_2^C < \sigma_1^C$. Now Rank-order applied to Row implies $\sigma_1^R < \sigma_2^R$. Contradiction. 2) $\sigma_1^R < \sigma_2^R$. Then Rank-order applied to Col implies $\sigma_2^C > \sigma_1^C$. Now Rank-order applied to Row implies $\sigma_1^R > \sigma_2^R$. Contradiction. Therefore, in *any* regular QRE in Game 1, $\sigma_1^R = \sigma_2^R$, and consequently, $\sigma_1^C = \sigma_2^C$. Suppose $\sigma_1^R > \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_J^C > \sigma_1^C$, and since $\sigma_J^C = 1 - 2\sigma_1^C$, we have $\frac{1}{3} > \sigma_1^C$. Then Rank-order applied to Row implies $\sigma_1^R < \sigma_J^R$, and so $\sigma_1^R < \frac{1}{3}$. Contradiction. Suppose $\sigma_1^R < \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_J^C < \sigma_1^C$, and so $\frac{1}{3} < \sigma_1^C$. Then Rank-order applied to Row implies $\sigma_1^R > \sigma_J^R$, and so $\sigma_1^R > \frac{1}{3}$. Contradiction. Therefore $\sigma_1^R = \frac{1}{3}$, and hence $\sigma^R = \sigma^C = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ in *any* regular QRE, so the equilibrium is unique.

Consider Game 2. Expected utility of Row in this game from choosing each of her three pure strategies is

$$u_{R1}(\sigma^C) = 10 + 20\sigma_2^C$$
$$u_{R2}(\sigma^C) = 10 + 20\sigma_1^C$$
$$u_{RJ}(\sigma^C) = 55 - 45\sigma_1^C - 45\sigma_2^C$$

The previous analysis immediately implies that in *any* regular QRE, $\sigma_1^R = \sigma_2^R$, and $\sigma_1^C = \sigma_2^C$. Suppose $\sigma_1^R \geq \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_1^C \leq \sigma_J^C$, and since $\sigma_J^C = 1 - 2\sigma_1^C$, we have $\sigma_1^C \leq \frac{1}{3}$. Then Rank-order applied to Row implies $\sigma_1^R < \sigma_J^R$, so $\sigma_1^R < \frac{1}{3}$. Contradiction. Therefore, $\sigma_1^R < \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_1^C > \sigma_J^C$, and since $\sigma_J^C = 1 - 2\sigma_1^C$, we have $\sigma_1^C > \frac{1}{3}$. If we also had $\sigma_1^C > \frac{9}{22}$, then Rank-order applied to Row would imply $\sigma_1^R > \sigma_J^R$, hence $\sigma_1^R > \frac{1}{3}$, contradiction. Thus in any regular QRE, $\frac{1}{3} < \sigma_1^C \leq \frac{9}{22}$ and $\sigma_1^R < \frac{1}{3}$. It remains to prove that σ_1^R and σ_1^C are uniquely defined. Applying Responsiveness to Col implies that σ_1^C is strictly increasing in U_{C1}, and therefore is strictly decreasing in σ_1^R. Using the same argument for Row, σ_1^R is strictly increasing in σ_1^C. Therefore any regular QRE in Game 2 is unique.

Consider Game 3. Expected utility of Row in this game from choosing each of her three pure strategies

[25] Obviously, $\sigma_J^R = 1 - \sigma_1^R - \sigma_2^R$ and $\sigma_J^C = 1 - \sigma_1^C - \sigma_2^C$.

is

$$u_{R1}(\sigma^C) = 10 + 15\sigma_1^C + 20\sigma_2^C$$
$$u_{R2}(\sigma^C) = 10 + 20\sigma_1^C + 15\sigma_2^C$$
$$u_{RJ}(\sigma^C) = 30 - 20\sigma_1^C - 20\sigma_2^C$$

As before, it is easy to show that in *any* regular QRE, $\sigma_1^R = \sigma_2^R$, and therefore $\sigma_1^C = \sigma_2^C$. Applying Responsiveness to Col implies that σ_1^C is strictly increasing in U_{C1}, and therefore is strictly decreasing in σ_1^R. Using the same argument for Row, σ_1^R is strictly increasing in σ_1^C. Therefore any regular QRE in Game 3 is unique. To prove the bounds on QRE probabilities, suppose $\sigma_1^R \leq \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_1^C \geq \sigma_J^C$, and since $\sigma_J^C = 1 - 2\sigma_1^C$, we have $\sigma_1^C \geq \frac{1}{3}$. Then Rank-order applied to Row implies $\sigma_1^R > \sigma_J^R$, so $\sigma_1^R > \frac{1}{3}$. Contradiction. Therefore, $\sigma_1^R > \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_1^C < \sigma_J^C$, and since $\sigma_J^C = 1 - 2\sigma_1^C$, we have $\sigma_1^C < \frac{1}{3}$. If we also had $\sigma_1^C < \frac{4}{15}$, then Rank-order applied to Row would imply $\sigma_1^R < \sigma_J^R$, hence $\sigma_1^R < \frac{1}{3}$, contradiction. Thus in any regular QRE, $\frac{4}{15} \leq \sigma_1^C < \frac{1}{3}$ and $\sigma_1^R > \frac{1}{3}$.

Finally, consider Game 4. We will now write $\sigma_2^C = 1 - \sigma_1^C - \sigma_J^C$, then the expected utility of Row in this game from choosing each of her three pure strategies is

$$u_{R1}(\sigma^C) = 30 - 10\sigma_1^C - 20\sigma_J^C$$
$$u_{R2}(\sigma^C) = 10 + 20\sigma_1^C$$
$$u_{RJ}(\sigma^C) = 10 + 20\sigma_J^C$$

Consider Col's equilibrium strategy. There are two possibilities: 1) $\sigma_1^C > \sigma_J^C$. Then Rank-order applied to Row implies $\sigma_2^R > \sigma_J^R$, hence $\sigma_1^R + 2\sigma_2^R > 1$. Then Rank-order applied to Col implies $\sigma_1^C < \sigma_J^C$. Contradiction. 2) $\sigma_1^C < \sigma_J^C$. Then Rank-order applied to Row implies $\sigma_2^R < \sigma_J^R$, hence $\sigma_1^R + 2\sigma_2^R < 1$. Then Rank-order applied to Col implies $\sigma_1^C > \sigma_J^C$. Contradiction. Therefore, in *any* regular QRE in Game 4, $\sigma_1^C = \sigma_J^C$, and consequently, $\sigma_2^R = \sigma_J^R$ (or, equivalently, $\sigma_1^R + 2\sigma_2^R = 1$). Applying Responsiveness to Row, σ_J^R is strictly increasing in U_{RJ}, and therefore is strictly increasing in $\sigma_J^C \equiv \sigma_1^C$. Using the same argument for Col, σ_1^C is strictly decreasing in $\sigma_2^R \equiv \sigma_J^R$. Therefore any regular QRE in Game 4 is unique. To prove the bounds on QRE probabilities, suppose $\sigma_2^R \geq \sigma_1^R$. Then by Rank-order applied to Col, $\sigma_1^C \leq \sigma_2^C$ and so $\sigma_1^C \leq \frac{1}{3}$. But then Rank-order applied to Row implies $\sigma_2^R < \sigma_1^R$. Contradiction. Hence $\sigma_2^R < \sigma_1^R$, and so $\sigma_2^R < \frac{1}{3}$. Then Rank-order applied to Col implies $\sigma_1^C > \sigma_2^C$, and so $\sigma_1^C > \frac{1}{3}$. If $\sigma_1^C > \frac{2}{5}$, then by Rank-order $\sigma_1^R < \sigma_2^R$. Contradiction. Therefore $\frac{1}{3} < \sigma_1^C \leq \frac{2}{5}$ and $\sigma_2^R < \frac{1}{3}$. □

D Additional Details for Computing the Test Statistic

As defined in the main text, the P-dimensional vector $\boldsymbol{\nu}$ contains the value of the CM inequalities evaluated at the choice frequencies observed in the experimental data. Specifically, the ℓ-th component of $\boldsymbol{\nu}$, corresponding to a given cycle $G_0, \ldots, G_{\mathcal{L}}$ of games is given by

$$\begin{aligned}
\nu_\ell = \sum_{m=G_0}^{G_{\mathcal{L}}} \pi_{i1}^m & \left[\pi_{k1}^{m+1} u_i^{m+1}(s_{i1}, s_{k1}) - \pi_{k1}^m u_i^m(s_{i1}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{i1}, s_{k2}) - \pi_{k2}^m u_i^m(s_{i1}, s_{k2}) \right. \\
& + \left. (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{i1}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{i1}, s_{kJ})\right] \\
& + \pi_{i2}^m \left[\pi_{k1}^{m+1} u_i^{m+1}(s_{i2}, s_{k1}) - \pi_{k1}^m u_i^m(s_{i2}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{i2}, s_{k2}) - \pi_{k2}^m u_i^m(s_{i2}, s_{k2}) \right. \\
& + \left. (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{i2}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{i2}, s_{kJ})\right] \\
& + (1 - \pi_{i1}^m - \pi_{i2}^m) \left[\pi_{k1}^{m+1} u_i^{m+1}(s_{iJ}, s_{k1}) - \pi_{k1}^m u_i^m(s_{iJ}, s_{k1}) \right. \\
& + \pi_{k2}^{m+1} u_i^{m+1}(s_{iJ}, s_{k2}) - \pi_{k2}^m u_i^m(s_{iJ}, s_{k2}) \\
& + \left. (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{iJ}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{iJ}, s_{kJ})\right]
\end{aligned}$$

where we use i to denote the Row player, k to denote the Column player, and ℓ changes from 1 to 20. For the Column player and $\ell \in [21, 40]$ the analogous expression is as follows:

$$\begin{aligned}
\nu_\ell = \sum_{m=G_0}^{G_{\mathcal{L}}} & \pi_{k1}^m \left[\pi_{i1}^{m+1} u_k^{m+1}(s_{i1}, s_{k1}) - \pi_{i1}^m u_k^m(s_{i1}, s_{k1}) + \pi_{i2}^{m+1} u_k^{m+1}(s_{i2}, s_{k1}) - \pi_{i2}^m u_k^m(s_{i2}, s_{k1})\right. \\
& + \left.(1 - \pi_{i1}^{m+1} - \pi_{i2}^{m+1}) u_k^{m+1}(s_{iJ}, s_{k1}) - (1 - \pi_{i1}^m - \pi_{i2}^m) u_k^m(s_{iJ}, s_{k1})\right] \\
& + \pi_{k2}^m \left[\pi_{i1}^{m+1} u_k^{m+1}(s_{i1}, s_{k2}) - \pi_{i1}^m u_k^m(s_{i1}, s_{k2}) + \pi_{i2}^{m+1} u_k^{m+1}(s_{i2}, s_{k2}) - \pi_{i2}^m u_k^m(s_{i2}, s_{k2})\right. \\
& + \left.(1 - \pi_{i1}^{m+1} - \pi_{i2}^{m+1}) u_k^{m+1}(s_{iJ}, s_{k2}) - (1 - \pi_{i1}^m - \pi_{i2}^m) u_k^m(s_{iJ}, s_{k2})\right] \\
& + (1 - \pi_{k1}^m - \pi_{k2}^m) \left[\pi_{i1}^{m+1} u_k^{m+1}(s_{i1}, s_{kJ}) - \pi_{i1}^m u_k^m(s_{i1}, s_{kJ})\right. \\
& + \pi_{i2}^{m+1} u_k^{m+1}(s_{i2}, s_{kJ}) - \pi_{i2}^m u_k^m(s_{i2}, s_{kJ}) \\
& + \left.(1 - \pi_{i1}^{m+1} - \pi_{i2}^{m+1}) u_k^{m+1}(s_{iJ}, s_{kJ}) - (1 - \pi_{i1}^m - \pi_{i2}^m) u_k^m(s_{iJ}, s_{kJ})\right]
\end{aligned}$$

We differentiate the above expressions with respect to π^m to obtain a $P \times 16$ estimate of the Jacobian $\hat{J} = \frac{\partial}{\partial \boldsymbol{\pi}} \boldsymbol{\mu}(\hat{\boldsymbol{\pi}})$ in order to compute an estimate of the variance-covariance matrix $\hat{\Sigma}_{[P \times P]} = \hat{J} \hat{V} \hat{J}'$ by the Delta method. For the case of four games, the partial derivatives form the 40×16 matrix $\hat{J}$. The first 20 rows correspond to the differentiated LHS of the cycles for the Row player, and the last 20 rows correspond to the differentiated LHS of the cycles for the Column player. The first 8 columns correspond to the derivatives with respect to π_{i1}^m, π_{i2}^m, and the last 8 columns correspond to the derivatives with respect to π_{k1}^m, π_{k2}^m, $m \in \{1, \ldots, 4\}$.[26]

Let $S_0^m \equiv \{\ell \in \{1, \ldots, 40\} | m \notin C_\ell\}$ be the set of cycle indices such that corresponding cycles (in the order given in (12)) do not include game m. E.g., for $m = 1$, $S_0^m = \{4, 5, 6, 13, 14, 24, 25, 26, 33, 34\}$. Let $S_i^m \equiv \{\ell \in \{1, \ldots, 20\} | \ell \notin S_0^m\}$ be a subset of cycle indices that include game m and pertain to the Row player, and let $S_k^m \equiv \{\ell \in \{21, \ldots, 40\} | \ell \notin S_0^m\}$ be a subset of cycle indices that include game m and pertain to the Column player. Finally, for a cycle of length $\mathcal{L}$, denote $\ominus \equiv - \mod \mathcal{L}$ subtraction modulus $\mathcal{L}$.

We can now express the derivatives with respect to π_{i1}^m and π_{i2}^m, $m \in \{1, \ldots, 4\}$, in the following general form. The partial derivatives wrt π_{i1}^m are

$$\begin{aligned}
\frac{\partial \nu_\ell}{\partial \pi_{i1}^m} &= 0 && \text{for } \ell \in S_0^m \\
\frac{\partial \nu_\ell}{\partial \pi_{i1}^m} &= \pi_{k1}^{m+1} u_i^{m+1}(s_{i1}, s_{k1}) - \pi_{k1}^m u_i^m(s_{i1}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{i1}, s_{k2}) \\
&- \pi_{k2}^m u_i^m(s_{i1}, s_{k2}) + (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{i1}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{i1}, s_{kJ}) \\
&- \left[\pi_{k1}^{m+1} u_i^{m+1}(s_{iJ}, s_{k1}) - \pi_{k1}^m u_i^m(s_{iJ}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{iJ}, s_{k2}) - \pi_{k2}^m u_i^m(s_{iJ}, s_{k2})\right. \\
&+ \left.(1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{iJ}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{iJ}, s_{kJ})\right] && \text{for } \ell \in S_i^m \\
\frac{\partial \nu_\ell}{\partial \pi_{i1}^m} &= \pi_{k1}^m \left[-u_k^m(s_{i1}, s_{k1}) + u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^m \left[-u_k^m(s_{i1}, s_{k2}) + u_k^m(s_{iJ}, s_{k2})\right] \\
&+ (1 - \pi_{k1}^m - \pi_{k2}^m) \left[-u_k^m(s_{i1}, s_{kJ}) + u_k^m(s_{iJ}, s_{kJ})\right] \\
&+ \pi_{k1}^{m\ominus 1} \left[u_k^m(s_{i1}, s_{k1}) - u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^{m\ominus 1} \left[u_k^m(s_{i1}, s_{k2}) - u_k^m(s_{iJ}, s_{k2})\right] \\
&+ (1 - \pi_{k1}^{m\ominus 1} - \pi_{k2}^{m\ominus 1}) \left[u_k^m(s_{i1}, s_{kJ}) - u_k^m(s_{iJ}, s_{kJ})\right] && \text{for } \ell \in S_k^m
\end{aligned}$$

[26]Clearly, the probability to choose Joker can be expressed via the probabilities to choose 1 and 2, using the total probability constraint.

The partial derivatives wrt π_{i2}^m are

$$\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} = 0 \qquad \text{for } \ell \in S_0^m$$

$$\begin{aligned}\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} &= \big[\pi_{k1}^{m+1} u_i^{m+1}(s_{i2}, s_{k1}) - \pi_{k1}^m u_i^m(s_{i2}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{i2}, s_{k2}) - \pi_{k2}^m u_i^m(s_{i2}, s_{k2}) \\ &+ (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{i2}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{i2}, s_{kJ})\big] \\ &- \big[\pi_{k1}^{m+1} u_i^{m+1}(s_{iJ}, s_{k1}) - \pi_{k1}^m u_i^m(s_{iJ}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{iJ}, s_{k2}) - \pi_{k2}^m u_i^m(s_{iJ}, s_{k2}) \\ &+ (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{iJ}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{iJ}, s_{kJ})\big] \qquad \text{for } \ell \in S_i^m\end{aligned}$$

$$\begin{aligned}\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} &= \pi_{k1}^m \left[-u_k^m(s_{i2}, s_{k1}) + u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^m \left[-u_k^m(s_{i2}, s_{k2}) + u_k^m(s_{iJ}, s_{k2})\right] \\ &+ (1 - \pi_{k1}^m - \pi_{k2}^m) \left[-u_k^m(s_{i2}, s_{kJ}) + u_k^m(s_{iJ}, s_{kJ})\right] \\ &+ \pi_{k1}^{m\ominus 1} \left[u_k^m(s_{i2}, s_{k1}) - u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^{m\ominus 1} \left[u_k^m(s_{i2}, s_{k2}) - u_k^m(s_{iJ}, s_{k2})\right] \\ &+ (1 - \pi_{k1}^{m\ominus 1} - \pi_{k2}^{m\ominus 1}) \left[u_k^m(s_{i2}, s_{kJ}) - u_k^m(s_{iJ}, s_{kJ})\right] \qquad \text{for } \ell \in S_k^m\end{aligned}$$

To obtain the derivatives with respect to π_{k1}^m and π_{k2}^m, one just needs to use the corresponding partial derivatives wrt π_{i1}^m and π_{i2}^m, and exchange everywhere the subscripts i and k, so we omit the derivation. For the sake of completenes, though, we list the indices subsets for each game $m \in \{1, .., 4\}$ in Table 4.

Table 4: Sets of cycle indices for each game.

m	S_0^m	S_i^m	S_k^m
1	4, 5, 6, 13, 14, 24, 25, 26, 33, 34	1, 2, 3, 7, 8, 9, 10, 11, 12, 15, 16, 17, 18, 19, 20	21, 22, 23, 27, 28, 29, 30, 31, 32, 35, 36, 37, 38, 39, 40
2	2, 3, 6, 10, 12, 22, 23, 26, 30, 32	1, 4, 5, 7, 8, 9, 11, 13, 14, 15, 16, 17, 18, 19, 20	21, 24, 25, 27, 28, 29, 31, 33, 34, 35, 36, 37, 38, 39, 40
3	1, 3, 5, 8, 11, 21, 23, 25, 28, 31	2, 4, 6, 7, 9, 10, 12, 13, 14, 15, 16, 17, 18, 19, 20	22, 24, 26, 27, 29, 30, 32, 33, 34, 35, 36, 37, 38, 39, 40
4	1, 2, 4, 7, 9 21, 22, 24, 27, 29	3, 5, 6, 8, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20	23, 25, 26, 28, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40

Notes. The cycle indices are for the Row player. To obtain the corresponding Column player cycle indices, swap the last two columns.

E Experiment Instructions

The instructions in the experiment, given below, largely follow McKelvey, Palfrey, & Weber (2000).

This is an experiment in decision making, and you will be paid for your participation in cash. Different subjects may earn different amounts. What you earn depends partly on your decisions and partly on the decisions of others.

The entire experiment will take place through computer terminals, and all interaction between subjects will take place through the computers. It is important that you do not talk or in any way try to communicate with other subjects during the experiment. If you violate the rules, we may ask you to leave the experiment.

We will start with a brief instruction period. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

This experiment consists of several periods or matches and will take between 30 to 60 minutes. I will now describe what occurs in each match.

[Turn on the projector]

First, you will be randomly paired with another subject, and each of you will simultaneously be asked to make a choice.

Each subject in each pair will be asked to choose one of the three rows in the table which will appear on the computer screen, and which is also shown now on the screen at the front of the room. Your choices will be always displayed as rows of this table, while your partner's choices will be displayed as columns. It will be the other way round for your partner: for them, your choices will be displayed as columns, and their choices as rows.

You can choose the first, the second, or the third row. Neither you nor your partner will be informed of what choice the other has made until after all choices have been made.

After each subject has made his or her choice, payoffs for the match are determined based on the choices made. Payoffs to you are indicated by the red numbers in the table, while payoffs to your partner are indicated by the blue numbers. Each cell represents a pair of payoffs from your choice and the choice of your partner. The units are in francs, which will be exchanged to US dollars at the end of the experiment.

For example, if you choose 'A' and your partner chooses 'D', you receive a payoff of 10 francs, while your partner receives a payoff of 20 francs. If you choose 'A' and your partner chooses 'F', you receive a payoff of 30 francs, while your partner receives a payoff of 30 francs. If you choose 'C' and your partner chooses 'E', you receive a payoff of 10 francs, while your partner receives a payoff of 20 francs. And so on.

Once all choices have been made the resulting payoffs and choices are displayed, the history panel is updated and the match is completed.

[show the slide with a completed match]

This process will be repeated for several matches. The end of the experiment will be announced without warning. In every match, you will be randomly paired with a new subject. The identity of the person you are paired with will never be revealed to you. The payoffs and the labels may change every match.

After some matches, we will ask you to indicate what you think is the likelihood that your current partner has made a particular choice. This is what it looks like.

[show slide with belief elicitation]

Suppose you think that your partner has a 15% chance of choosing 'D' and a 60% chance of choosing 'E'. Indicate your opinion using the slider, and then press 'Confirm'. Once all subjects have indicated their opinions and confirmed them, the resulting payoffs and choices are displayed, the history panel is updated and the match is completed as usual.

Your final earnings for the experiment will be the sum of your payoffs from all matches. This amount in francs will be exchanged into U.S. dollars using the exchange rate of 90 cents for 100 francs. You will see your total payoff in dollars at the end of the experiment. You will also receive a show-up fee of $7. Are there any questions about the procedure?

[wait for response]

We will now start with four practice matches. Your payoffs from the practice matches are not counted in your total. In the first three matches you will be asked to choose one of the three rows of a table. In the fourth match you will be also asked to indicate your opinion about the likelihood of your partner's choices for each of three actions. Is everyone ready?

[wait for response]

Now please double click on the 'Client Multistage' icon on your desktop. The program will ask you to type in your name. Please type in the number of your computer station instead.

[wait for subjects to connect to server]

We will now start the practice matches. Do not hit any keys or click the mouse button until you are told to do so.

[start first practice match]

You see the experiment screen. In the middle of the screen is the table which you have previously seen up on the screen at the front of the room. At the top of the screen, you see your subject ID number, and your computer name. You also see the history panel which is currently empty.

We will now start the first practice match. Remember, do not hit any keys or click the mouse buttons until you are told to do so. You are all now paired with someone from this class and asked to choose one of the three rows. Exactly half of your see label 'A' at the left hand side of the top row, while the remaining half now see label 'D' at the same row.

Now, all of you please move the mouse so that it is pointing to the top row. You will see that the row is highlighted in red. Move the mouse to the bottom row and the highlighting goes along with the mouse. To choose a row you just click on it. Now please click once anywhere on the bottom row.

[Wait for subjects to move mouse to appropriate row]

After all subjects have confirmed their choices, the match is over. The outcome of this match, 'C'-'F', is now highlighted on everybody's screen. Also, note that the moves and payoffs of the match are recorded in the history panel. The outcomes of all of your previous matches will be recorded there throughout the experiment so that you can refer back to previous outcomes whenever you like. The payoff to the subject who chose 'C' for this match is 20, and the payoff to the subject who chose F is '10'.

You are not being paid for the practice session, but if this were the actual experiment, then the payoff you see on the screen would be money (in francs) you have earned from the first match. The total you earn over all real matches, in addition to the show-up fee, is what you will be paid for your participation in the experiment.

We will now proceed to the second practice match.

[Start second match]

For the second match, you have been randomly paired with a different subject. You are not paired with the same person you were paired with in the first match. The rules for the second match are exactly like for the first. Please make your choices.

[Wait for subjects]

We will now proceed to the third practice match. The rules for the third match are exactly like the first. Please make your choices.

[Start third match]

We will now proceed to the fourth practice match. The rules for the fourth match are exactly like the first. Please make your choices.

[Wait for subjects]

Now that you have made your choice, you see that a slider appears asking you to indicate the relative likelihood of your partner choosing each of the available actions. There is also a confirmation button. Please indicate your opinion by adjusting the thumbs and then press 'Confirm'.

[wait for subjects]

This is the end of the practice match. Are there any questions?

[wait for response]

Now let's start the actual experiment. If there are any problems from this point on, raise your hand and an experimenter will come and assist you. Please pull up the dividers between your cubicles.

[start the actual session]

The experiment is now completed. Thank you all very much for participating in this experiment. Please record your total payoff from the matches in U.S. dollars at the experiment record sheet. Please add your show-up fee and write down the total, rounded up to the nearest dollar. After you are done with this, please remain seated. You will be called by your computer name and paid in the office at the back of the room one at a time. Please bring all your things with you when you go to the back office. You can leave the experiment through the back door of the office.

Please refrain from discussing this experiment while you are waiting to receive payment so that privacy regarding individual choices and payoffs may be maintained.