Testing the Quantal Response Hypothesis*

By Emerson Melo, Kirill Pogorelskiy, and Matthew Shum[1]

Indiana University Bloomington, U.S.A.; University of Warwick, U.K.; Caltech, U.S.A.

March 22, 2018

We develop a non-parametric test for consistency of player behavior with the Quantal Response Equilibrium (QRE). The test exploits a characterization of the equilibrium choice probabilities in *any* structural QRE as the gradient of a convex function; thereby QRE-consistent choices satisfy 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 cannot reject individual-level quantal response behavior for over half of the subjects.

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

Keywords: quantal response equilibrium, behavioral game theory, cyclic monotonicity, moment inequalities, experimental economics

1. INTRODUCTION

A vast literature has demonstrated that, across a wide variety of games, behavior deviates systematically from Nash equilibrium predictions. In order to account for these deviations in a natural fashion, while preserving the idea of equilibrium, McKelvey and Palfrey (1995) introduced the notion of *Quantal Response Equilibrium (QRE)*. QRE has become a popular tool in experimental

*Manuscript received April 2017; revised October 2017.

[1]We thank Marina Agranov, Larry Blume, Kim Border, Federico Echenique, Ben Gillen, Russell Golman, Phil Haile, Jörg Stoye, Leeat Yariv and especially Tom Palfrey for support and insightful comments. We also thank the participants of the Yale conference on Heterogenous Agents and Microeconometrics and LSE theory seminar. We are grateful to Michael McBride and Alyssa Acre at UC Irvine ESSL laboratory for help with running experiments. Please address correspondence to: Matthew Shum, Division of the Humanities and Social Sciences, California Institute of Technology, 1200 E California Blvd, Pasadena CA 91125, U.S.A. Email: *mshum@caltech.edu*.

economics because typically it provides an improved fit to the experimental data; moreover, it is also a key model in behavioral game theory and serves as a benchmark for tractable alternative theories of bounded rationality.[2] While most existing work estimating structural QRE (e.g., Merlo and Palfrey (2013), Camerer et al. (2016) to name a few) utilizes the *parametric*, logit version of QRE and its variants, *nonparametric* tests of QRE have not been available.

This paper is the first to develop and implement a formal nonparametric procedure to test, using experimental data, whether subjects are indeed behaving according to QRE. Our approach is related to the econometrics literature on semiparametric discrete-choice models and moment inequalities. Our test is based on the notion of *cyclic monotonicity*, a concept from convex analysis. 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.

To illustrate the basic ideas, consider a matching pennies game (Game A) in Figure 1. The

		Column	
		Left	Right
Row	Top	320, 40	40, 80
	Bottom	40, 80	80, 40

Figure 1: Game A: An asymmetric matching pennies game.

unique Nash equilibrium has Row mixing 50/50 and Column mixing 1/8 and 7/8. This prediction is strikingly rejected by the data – in a series of lab experiments, Goeree and Holt (2001) report that Row players choose to play Top 96% of the time. To rationalize this behavior, one could add random shocks to Row's payoff as in discrete-choice models (McFadden, 1974), and assume that Row is choosing actions that maximize the sum of her expected payoff and the shock realization. However, if both players hold correct (and common) beliefs about preference shock distributions, then Row's shocks will now be also affecting Column's payoffs, and hence Column's actions. A fixed point condition is required, and this is precisely what QRE provides. Rather than playing their best responses with probability one (as in a Nash equilibrium), in a QRE the players are assumed to

[2]See Camerer (2003) and Crawford et al. (2013). The latest book-length treatment of the theory behind QRE and its numerous applications in economics and political science is Goeree et al. (2016).

play "better" responses, whereby they choose actions with *higher* expected payoffs *more often* than actions with *lower* expected payoffs, and expectations are computed under correct beliefs about other players' "better" response behavior. The "better" response property of QRE, with additional assumptions on the shock distributions, becomes a testable restriction imposed by QRE *within* a given game.

In this paper, we investigate testable restrictions imposed by QRE *between* games. Consider another matching pennies game (Game B) in Figure 2.

		Column	
		Left	Right
Row	Top	80, 40	40, 80
	Bottom	40, 80	80, 40

Figure 2: Game B: Another matching pennies game.

Suppose that the same two players play both games A and B several times and face the same preference shock distributions in each game. If a QRE is played in each game, then the following necessary condition must hold: If playing "Top" yields a higher equilibrium expected payoff in Game A than in Game B, then the Row player should *not* be observed playing "Top" in Game B more often than she does in Game A. And, if playing "Left" has a higher equilibrium expected payoff in Game A than in Game B, then the Column player should *not* be observed choosing "Left" in Game B more often than he does in Game A. In this way, QRE probabilities are sensitive to changes in expected equilibrium payoffs, and a test for QRE can be based on the response in players' choice probabilities resulting from changes in payoffs across games.

The Cyclic Monotonicity (CM) property generalizes this intuition to more than two games and more than two actions. Precisely, CM implies a set of inequality restrictions over players' quantal response equilibrium choice probabilities and corresponding differences in equilibrium expected utilities across games that only differ in the payoffs. By substituting the unobserved equilibrium choice probabilities with empirical choice frequencies, the CM inequalities can be rewritten using sample moments, thereby reducing the test of QRE to a well-studied problem in the econometrics literature: that of joint testing of a set of moment inequalities.

The classical tests for moment inequalities with estimated quantities (e.g., Wolak (1989)) have low power when the number of inequalities is large (as is typically the case). To address this problem, we employ the *generalized moment selection* (GMS) approach by Andrews and Soares (2010). GMS aims to increase power while controlling size by utlizing a form of a "pre-test" to eliminate inequalities that are sufficiently far from being violated. We provide simulation evidence suggesting that GMS controls size and is effective in detecting deviations from QRE. We also checked consistency with QRE using a more recent approach to testing moment inequalities by Romano et al. (2014).

Since we vary payoffs across games, we have to fix the distribution of players' preference shocks, i.e. we have to assume that it is unaffected by the payoffs of the particular game at play (our invariance assumption). Importantly, our test for QRE is *nonparametric* in that one need not *specify* a particular probability distribution for the shocks (only that it is the *same* between games); thus, the results are robust to a wide variety of distributions.

In addition to the invariance assumption, our test has also other limitations, discussed in detail in Section 4.2. First, in order to be able to compute CM inequalities, we need to know how players' payoffs in each game enter into their utility, so we assume that players' utility functions are known (e.g., as could be recovered from their observed game payoffs in laboratory experiments).[3] Second, we also have to assume that there is a unique equilibrium played in the data.

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. When we consider subjects individually, however, we find that the quantal response behavior 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.[4]

Our work here builds upon and extends Haile et al. (2008) (hereafter HHK) who showed how one can construct shock distributions for which QRE rationalizes any outcome *within* a game, as long as shocks are not i.i.d. across players' actions. HHK described several approaches for testing for QRE, but did not provide guidance for formal econometric implementation of such tests. We

[3]This is a straightforward exercise under risk-neutrality, but as will be explained below, the test itself does not depend on risk-neutrality – any common utility function can be plugged in, allowing the researcher to verify robustness of the test results with respect to the utility specification.

[4] While we focus on testing QRE using lab experimental data in this paper, there may be field settings in which a game is repeated among a group of players, where our test could potentially be applied. One example could be intra-household games between husband and wife; see, for instance, Del Boca et al. (2014).

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 expected 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 literature. Elsewhere, cyclic monotonicity has been studied in the context of multidimensional mechanism design. In particular, the papers by Rochet (1987), Saks and Yu (2005), Lavi and Swamy (2009), Berger et al. (2009), Ashlagi et al. (2010), and Archer and Kleinberg (2014) (summarized in Vohra (2011, Chapter 4)), relate the incentive compatibility (truthful implementation) of a mechanism to its cyclic monotonicity properties. Similarly, McFadden and Fosgerau (2012) introduce cyclic monotonicity to study revealed preference in discrete-choice models. In this paper, we apply the cyclic monotonicity to test for QRE using experimental data in which players' payoffs are known (because they are set by the experimenter). The cyclic monotonicity property may also be useful in settings where the researcher does not completely observe agents' utilities. In other work, one of the present authors (Chiong and Shum, 2016; Shi et al., 2018) has also used cyclic monotonicity for identification and estimation of semiparametric multinomial choice models, in which the payoff shocks are left unspecified (as here), but the utility functions are parametric and assumed to be known up to a finite-dimensional parameter.

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 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 uses the alternative econometric procedure by Romano et al. (2014) for testing the QRE hypothesis. Appendix F contains experimental instructions.

2. QRE BACKGROUND

In this section we briefly review the main ideas behind the QRE approach.

To convey the basic intuition behind QRE first, consider again Game A in Figure 1 in the Introduction. Suppose Row believes that Column is mixing between Left and Right with probabilities $(p_L, 1-p_L) \equiv \boldsymbol{p}_C$. Then her expected payoff from playing Top is $u_{RT}(\boldsymbol{p}_C) = 320p_L + 40(1-p_L)$, and her expected payoff from playing Bottom is $u_{RB}(\boldsymbol{p}_C) = 40p_L + 80(1-p_L)$. Similarly, for Column, let $u_{CL}(\boldsymbol{p}_R)$ and $u_{CR}(\boldsymbol{p}_R)$ be expected payoffs associated with playing Left and Right respectively, computed under belief $\boldsymbol{p}_R \equiv (p_T, 1-p_T)$ about Row's mixing. A (structural) version of the quantal response behavior assumes that instead of using best responses, both players use "better" (*quantal*) responses: They choose actions that maximize the sum of their expected payoffs and the shock realization $\boldsymbol{\varepsilon}$, administered to each available action. In particular, Row plays Top if and only if $u_{RT}(\boldsymbol{p}_C) + \varepsilon_{RT} \geq u_{RB}(\boldsymbol{p}_C) + \varepsilon_{RB}$, and Column plays Left if and only if $u_{CL}(\boldsymbol{p}_R) + \varepsilon_{CL} \geq u_{CR}(\boldsymbol{p}_R) + \varepsilon_{CR}$. Assuming that shock distributions are common knowledge, it is easy to see that Row will be observed playing Top with probability

$$\pi_{RT}(\boldsymbol{p}_C) = \mathbb{P}(u_{RT}(\boldsymbol{p}_C) + \varepsilon_{RT} \geq u_{RB}(\boldsymbol{p}_C) + \varepsilon_{RB})$$

and Column will be observed playing Left with probability

$$\pi_{CL}(\boldsymbol{p}_R) = \mathbb{P}(u_{CL}(\boldsymbol{p}_R) + \varepsilon_{CL} \geq u_{CR}(\boldsymbol{p}_R) + \varepsilon_{CR}).$$

Assuming that players hold correct beliefs about each other's quantal response behavior, in a QRE the following fixed point condition on beliefs must hold:

$$\boldsymbol{p}_C = (\pi_{CL}(\boldsymbol{p}_R), 1-\pi_{CL}(\boldsymbol{p}_R)), \qquad \boldsymbol{p}_R = (\pi_{RT}(\boldsymbol{p}_C), 1-\pi_{RT}(\boldsymbol{p}_C))$$

The exact solution depends on how the preference shocks are distributed, and correspondingly, each such distribution can generate a different QRE. For example, if $\boldsymbol{\varepsilon}$ has an extreme value distribution with parameter $\lambda \in [0, \infty)$, this generates the usual logit-QRE model, in which

$$\pi_{RT}(\boldsymbol{\pi}) = e^{\lambda u_{RT}(\boldsymbol{\pi})} / \left(e^{\lambda u_{RT}(\boldsymbol{\pi})} + e^{\lambda u_{RB}(\boldsymbol{\pi})}\right),$$
$$\pi_{CL}(\boldsymbol{\pi}) = e^{\lambda u_{CL}(\boldsymbol{\pi})} / \left(e^{\lambda u_{CL}(\boldsymbol{\pi})} + e^{\lambda u_{CR}(\boldsymbol{\pi})}\right).$$

Parameter $\lambda \geq 0$ captures the level of randomness in players' behavior (i.e. deviations from Nash). If $\lambda = 0$, both Row and Column choose actions uniformly randomly (which is *not* a Nash equilibrium in Game A). But for $\lambda \to \infty$ the logit-QRE mixtures approach the unique Nash equilibrium of the underlying game. A typical fitting exercise involves estimating the value of λ from the data using maximum likelihood (e.g., McKelvey et al. (2000)).

While most existing work has used logit-QRE for its tractability, assuming extreme-value-distributed shocks, other reasonable distributions can be employed that lead to different QRE probabilities. In this paper, we present a test of the QRE hypothesis which is applicable to *any* structural QRE. Thus, if the null hypothesis is rejected, one can be sure that no amount of tweaking the shock distributions could have delivered a QRE-consistent play.

We will now present the general QRE framework, using the notation from McKelvey and Palfrey (1995). Consider a finite n-person game $G(N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$. The set of pure strategies (actions) 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 choose. 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.[5] Assume also that shock distributions $(F_i)_{i\in N}$ are independent across players. 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

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

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

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

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 $\boldsymbol{\varepsilon}_i := (\varepsilon_{i1},\dots,\varepsilon_{iJ_i})$ 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.[6] In the specific environment of our test here, the invariance assumption allows us to compare choice probabilities across different games using the property of *cyclic monotonicity*, which we turn to in the next section.

[5] As a result, each action can be chosen with a positive probability, ruling out consistency with a pure-strategy Nash equilibrium (which can be restored in the limit as the shocks go to zero.) This is never a concern for our test since the choice probabilities are estimated from the data and necessarily contain some noise ruling out pure strategies. Furthermore all games in our application have a unique totally-mixed Nash equilibrium.

[6] Most applications of QRE assume that the preference shocks follow an extreme value distribution, regardless of the magnitude of payoffs. One exception is McKelvey et al. (2000), who allow the logit-QRE parameter to vary across different games. This direction is further developed in Rogers et al. (2009). As will be described below, in some cases, one may be able to relax invariance to "weak" invariance, under which players' shock distributions may depend on their own baseline payoffs but not those of other players.

3. A TEST BASED ON CONVEX ANALYSIS

In this section we propose a test for the QRE hypothesis. The basic idea, as explained earlier, is to test the restriction imposed by QRE between games that differ in the payoffs: if an equilibrium expected payoff from an action is higher in game A than in game B, then this action should be chosen more often in game A than in game B, provided that the invariance assumption holds. 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 literature, the expression $\varphi(\boldsymbol{u})$ is known as the social surplus function.[7] In the context of QRE, and given an equilibrium $\boldsymbol{\pi}^*$, $\varphi^i(\boldsymbol{u}_i(\boldsymbol{\pi}^*))$ gives player i's expected equilibrium payoffs. This implies that the change in equilibrium behavior across games can be summarized by the social surplus function, which 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 Eq.(3) holding for all players i is necessary and sufficient for $\left\{\pi_{ij}^*\right\}$ to be a (structural) 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 (sub)gradient of a convex function is characterized by the *Cyclic Monotonicity* (CM) property: given a mapping $\rho : \mathbb{R}^{J_i} \rightarrow \mathbb{R}^{J_i}$, it is necessary and sufficient that ρ be cyclically monotone in order that there exists a convex function $\phi^i : \mathbb{R}^{J_i} \rightarrow \mathbb{R}$ such that $\rho(x) \subset \partial\phi^i(x)$ for any x (here $\partial\phi^i$ denotes the subgradient of ϕ^i). This property is the generalization, for functions of several variables, of the fact that the derivative of a univariate convex function is monotone nondecreasing.[8]

[7] For details see McFadden (1981).
[8] See Chiong et al. (2016) for additional implications of convex analysis for discrete-choice models.

To define cyclic monotonicity in our setting, consider a *cycle*[9] of games $C \equiv (G_0, G_1, G_2, \ldots, G_0)$ where G_m denotes the game at index m in a cycle. These games are characterized by the same set of choices for each player, and the same distribution of payoff shocks (i.e. satisfying Assumption 1), but distinguished by payoff differences. 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.[10] Expanding the inner product notation, the CM 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. Without it, the payoff shock distributions, and hence the social surplus functions, may be different in each game, implying that, under a QRE play, the corresponding choice probabilities are gradients of different social surplus functions; thus the cyclic monotonicity property need not hold.

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, G_j, G_k, G_i) and (G_j, G_k, G_i, G_j) are the same.

Intuitively CM conditions can be derived from, and also imply, players' utility maximization in appropriately perturbed games. In particular, in Appendix A we prove the following result:

Proposition 1. *Consider a cycle of length $\mathcal{L} \geq 2$. The cyclic monotonicity conditions hold if and only if each player i's choice probabilities in each game along the cycle maximize the difference between her expected utility and the convex conjugate of the social surplus function* (2) *over the set of admissible distributions from $\Delta(S_i)$.*

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

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

Proposition 1 establishes the equivalence between the cyclic monotonicity condition and utility maximization among a family of perturbed games that only vary in the payoffs. The equivalence in proposition critically depends on the fact that equilibrium probabilities are given by Eq. (3). Formally, Eq. (3) allows us to use *conjugate duality* arguments to show that CM is equivalent to players' utility maximization. Thus Proposition 1 helps us interpret our test directly in terms of the main QRE property of positive responsiveness, where each action's probability is increasing in the action's expected payoff.

Remark: Cyclic monotonicity and incentive compatibility. Proposition 1 also implies that the CM inequalities can be interpreted as "incentive compatibility" conditions on players' choices 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 response to changes in the expected payoffs across the games in the cycle C. To see this, consider a family of games with only two actions, which vary in the expected payoff difference $\mathbb{E}(u_2 - u_1)$ between actions 2 and 1. Obviously, utility maximization should imply that the choice probabilities of playing action 2 across these games should be nondecreasing in the expected payoff differences; alternatively, monotonicity in expected payoff differences is an "incentive compatibility" condition on choice probabilities across these games.

In games with more than two actions, the expected payoff differences (relative to a benchmark action) form a vector, as do the choice probabilities for each action. Proposition 1 and the discussion above show that utility maximization across a family of games varying in expected payoff differences imply that the vector of corresponding choice probabilities in each game can be represented as a gradient of the social surplus function; that is, cyclic monotonicity is an incentive compatibility condition on choice probabilities in a family of games with more than two actions. ■

Special case: Two actions. As the previous remark pointed out, in a family of games with only two actions, cyclic monotonicity reduces to the usual monotonicity. That is, in these games, the cyclic monotonicity conditions (5) only need to be checked for cycles of length two.[11] Because many experiments study games where players' strategy sets consist of two elements, this

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

observation turns out to be useful from an applied perspective. ■

4. 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 one to estimate a vector of probabilities $[\boldsymbol{\pi}_i^*]^m \in \Delta(S_i)$ for each game m in the sample and compute the corresponding equilibrium expected utilities $[\boldsymbol{u}_i]^m$ (assuming risk-neutrality, or more generally, known utility functions).

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. 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 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 $n \sum_{\mathcal{L}=2}^{M} \#C(\mathcal{L})$ 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.

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,[12] which also imply stochastic inequalities involving estimated choice probabilities

[12] HHK also consider testing the QRE hypothesis using the so called Block-Marschak polynomials. A QRE-consistent play implies linear inequalities involving the Block-Marschak polynomials which may be tested formally using similar methods as we describe in this paper. However, the two approaches are qualitatively quite different as

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 two), 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

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

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

$$u_{i1}^1 - u_{i1}^0 \geq u_{i2}^1 - u_{i2}^0 \geq \ldots \geq u_{iJ_i}^1 - u_{iJ_i}^0. \tag{7}$$

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:*

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

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

Proposition 2. *Let $M = 2$. If all expected utility differences between corresponding pairs of actions in two games in* (7) *are non-negative, then HHK's cumulative rank property* (8) *implies the CM inequalities* (5)*. Conversely, the CM inequalities* (5) *imply the cumulative rank property* (8) *(without additional assumptions on expected utility differences).*

Proposition 2 shows that for the special case of comparing two games, HHK's cumulative rank property is related to, but distinct from the cyclic monotonicity inequalities. When there are more

the Block-Marschak approach involves comparing games which vary in the actions available to players and, as HHK point out, can be feasibly tested in the lab only for special families of games (such as Stackelberg games or games with some nonstrategic players). In contrast, our approach, based on cyclic monotonicity, compares games which vary in players' payoffs. For these reasons, a full exploration of the Block-Marschak approach is beyond the scope of this paper.

than two games, with more than two strategies in each game, the pairwise comparisons do not exhaust the restrictions in the cyclic monotonicity inequalities. (See Saks and Yu (2005), Ashlagi et al. (2010), and Vohra (2011, Chapter 4).)

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 et al. (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.

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[13], our testing procedure, strictly speaking, only assumes that *there is a unique equilibrium played in the data.*[14] 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 of 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, invariant shock distribution (Assumption 1), and unique equilibrium in the data.

Remark: Applicability outside experimental settings. Our test relies on the invariance assumption, which implies a setting with different games played by the same set of players with the

[13] See, e.g., Aguirregabiria and Mira (2007), Bajari et al. (2007).

[14] 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. See de Paula and Tang (2012) for a test of multiple equilibria presence in the data.

same joint strategy space. It may appear too restrictive to be applicable outside an experimental setting. However, there are field settings that may potentially satisfy these restrictions. As a simple example, consider auctions held at regular time intervals with essentially the same set of participants, like car dealer auctions. The payoffs (i.e. cars on sale today) differ across *some* games (sometimes very similar cars are sold at different auctions, sometimes not; in general, there will be different car models), but the players and their bid strategies likely change very little over a reasonably short time horizon. Another example, already mentioned earlier, could be intra-household games between husband and wife (Del Boca et al., 2014).

One may also wonder whether repeating some games several times may be incompatible with a static equilibrium concept like QRE. The same concern applies to testing for consistency with Nash equilibrium – repetitions are needed to estimate player choice probabilities in any game. Whether or not repeated-game effects matter is an empirical question. With quantal response, however, repeated strategies are stochastic, so a repeated equilibrium may be more difficult to sustain. ■

5. ECONOMETRIC IMPLEMENTATION: GENERALIZED MOMENT SELECTION PROCEDURE

In this section we consider the formal econometric properties of our test, and the application of the generalized model selection procedure of Andrews and Soares (2010). In Appendix E, we describe the implementation of an alternative, more recent moment selection procedure due to Romano et al. (2014). Both approaches show similar performance when applied to our experimental data.

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 $\ell \in \{1, \ldots, P\}$, 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 $\boldsymbol{\nu} \equiv (\nu_1, \ldots, \nu_\ell, \ldots, \nu_P)$ where each generic component ν_ℓ is given by (6), i.e.

$$\begin{aligned} \nu_\ell &= \sum_{m=m_0(\ell)}^{\mathcal{L}(\ell)-1} \sum_{j=1}^{J_{\iota(\ell)}} \sum_{\boldsymbol{s}_{-\iota(\ell)}} \pi^m_{\iota(\ell)j} \left(\left(\prod_{k \in N_{-\iota(\ell)}} \pi^{m+1}_{k_{j_k}} \right) u^{m+1}_{\iota(\ell)}(s_{\iota(\ell)j}, \boldsymbol{s}_{-\iota(\ell)}) \right. \\ &\quad \left. - \left(\prod_{k \in N_{-\iota(\ell)}} \pi^m_{k_{j_k}} \right) u^m_{\iota(\ell)}(s_{\iota(\ell)j}, \boldsymbol{s}_{-\iota(\ell)}) \right). \end{aligned} \tag{9}$$

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 $\boldsymbol{\mu}$ 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.[15]

We perform the following hypothesis test:

$$H_0: \ \boldsymbol{\mu}_{\mathbf{0}} \geq \mathbf{0} \quad vs. \quad H_1: \ \boldsymbol{\mu}_{\mathbf{0}} \ngeq \mathbf{0}, \tag{10}$$

where $\mathbf{0} \in \mathbb{R}^P$. Letting $\hat{\Sigma}$ denote an estimate of Σ, we utilize the following test statistic

$$S(\hat{\mu}, \hat{\Sigma}) := \sum_{\ell=1}^{P} \left[\hat{\mu}^\ell / \hat{\sigma}_\ell \right]_-^2 \tag{11}$$

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 a sum of squared violations across the inequalities, so that larger values of the statistic indicate evidence against the null hypothesis (10). Since there is a large number of inequalities in (10) (in the application, $P = 40$), we utilize the Generalized Moment Selection (GMS) procedure

[15]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 and Soares (2010).

of Andrews and Soares (2010) (hereafter "AS"). This procedure combines moment selection along with hypothesis testing, and is especially useful when there are many moment conditions. Typically, hypothesis tests involving many inequalities have low power, since "redundant" moment conditions, which are far from binding, tend to shift the asymptotic null distribution of the test statistic higher (in a stochastic sense), making it harder to reject. The AS procedure, which combines moment selection (that is, eliminating redundant moment conditions), with hypothesis testing, increases power and yields uniformly asymptotically critical values.

From the above description, we see that the AS procedure is to evaluate the asymptotic distribution of the test statistic under a sequence of parameters under the null hypothesis which resemble the sample inequalities, and are drifting to zero. By doing this, 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 null distribution of the test statistic, leading to a (stochastically) smaller distribution and hence smaller critical values.[16]

Using the AS procedure requires specifying an appropriate test statistic for the inequalities, and also specifying a drifting sequence of null hypotheses converging to zero. In doing both we follow the suggestions in AS. The test statistic in Eq. (11) satisfies the requirements for the AS procedure.[17] The drifting sequence $\{\kappa_K, K \to \infty\}$ also follows AS's suggestions and is described immediately below.

To obtain valid critical values for S under H_0, we use the following procedure:

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 $\kappa_K = (\log K)^{1/2}$.[18] 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

[16] In contrast, other inequality-based testing procedures (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.

[17] Cf. Eq. (3.6) and Assumption 1 in AS.

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

$[x]_+ = \max(x, 0)$.

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.

Interpretation of test results. As discussed in the remarks after Proposition 1, our test, in essence, checks for violations of the cyclic monotonicity inequalities. If such violations are substantial, at least one of the players must violate "incentive compatibility" conditions on her choices in at least one game cycle, making the joint behavior non-rationalizable by *any* structural QRE satisfying our assumptions. This is much more general than a bad fit of a logit QRE. Moreover, by estimating the left-hand sides $\hat{\mu}$ of CM inequalities from the data, we can pin down the exact game cycles that violate CM.

On the other hand, if our test does not reject the QRE hypothesis, our conclusions are less crisp. Similarly to the results from the revealed preference literature, consistency with CM conditions indicates a possibility result, i.e. in our case, the possibility of existence of a QRE rationalizing the data, rather than yielding estimates of a particular quantal response function or shock distribution.

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 studied in the previous experimental literature (cf. O'Neill (1987) and Brown and Rosenthal (1990)), which 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.[19] Table 1 shows the payoff matrices of the four games which we used in our simulations and experiments.

[19]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.

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.

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 payoffs. 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 his 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.[20] 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 *between games*. In the QRE literature, there are typically additional restrictions imposed on quantal response functions

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

within a fixed game. In particular, the quantal response functions studied in Goeree et al. (2005), in addition to the assumptions we impose in Section 2, also satisfy the *rank-order property*[21], which states that actions with higher expected payoffs are played with higher probability than actions with lower expected payoffs *within a game.* Formally, a quantal response function $\pi_i : \mathbb{R}^{J_i} \to \Delta(S_i)$ satisfies the *rank-order property* if for all $i \in N, j, k \in \{1, \ldots, J_i\}$:

$$u_{ij}(\mathbf{p}) > u_{ik}(\mathbf{p}) \Rightarrow \pi_{ij}(\mathbf{p}) > \pi_{ik}(\mathbf{p}). \tag{12}$$

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 (12) 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) under-reject 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 a second set of simulations, we generated artificial data under a non-QRE play (specifically, we generated a set of choice probabilities resulting in violation of all of the CM inequalities for both players).[22] 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.

To illustrate the power properties further, we took the set of QRE-consistent probabilities π_i and added a noise term $\lambda\varepsilon_i$, with ε_i distributed uniformly on $[-0.25, 0.25]$ and adjusted such that choice probabilities for each player in each game sum up to one, and varied the relative noise weight, λ, from 0 to 1. Thus at $\lambda = 0$ we have a fully QRE-consistent distribution, and $0 < \lambda \leq 1$ spans a "mixed" distribution in which some cyclic monotonicity inequalities are violated, and some

[21] We borrow the name of this property from Fox (2007). Goeree et al. (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.

[22] These probabilities are as follows: For Row mix $[0.3334, 0.3324, 0.3342]$, $[0.4923, 0.5057, 0.0020]$, $[0.3326, 0.3332, 0.3343]$, $[0.3326, 0.3335, 0.3339]$, in Games 1–4, respectively. For Column, mix $[0.3325, 0.3327, 0.3349]$, $[0.3334, 0.3334, 0.3333]$, $[0.3331, 0.3321, 0.3348]$, $[0.3331, 0.3325, 0.3344]$ in Games 1–4, respectively. The resulting test statistic is about 6,663.

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%	8	11	14	10
	10%	14	26	32	19
	20%	35	49	63	43
5000	5%	10	23	30	17
	10%	22	37	53	30
	20%	39	74	96	58
9000	5%	19	44	53	30
	10%	34	64	85	55
	20%	66	114	140	94

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 $R = 5000$ to simulate the corresponding critical values.
[a]: # rejected out of 500 replications for each significance level.

are not. The corresponding rejection probabilities are graphed in Figure 3, and show that our test appears keenly sensitive to violations of QRE; rejection probabilities are quite high once the mixing parameter λ exceeds 0.2 (for large $N = 9000$) and 0.4 (for small $N = 1000$), and for values of λ exceeding 0.4 (0.6 for small N), the test rejection probability exceeds 90 %. This demonstrates the good power properties of our testing procedure.

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

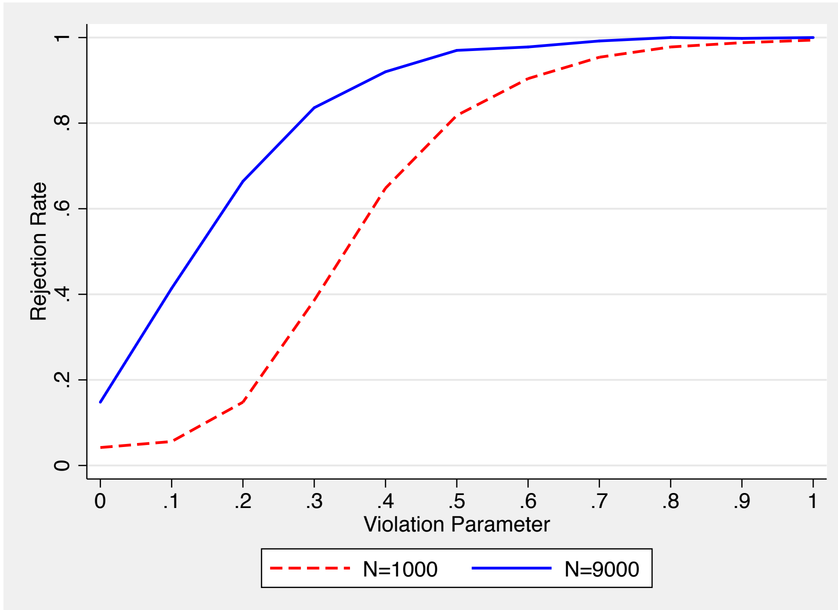

Figure 3: Rejection rate at 10% level as a function of violation parameter λ in Monte Carlo simulations for different N ($\kappa_K = 5(\log(K))^{\frac{1}{4}}$.)

$1-4$:

$$\begin{aligned} &121, 131, 141, 232, 242, 343 \\ &1231, 1241, 1321, 1341, 1421, 1431, 2342, 2432 \\ &12341, 12431, 13241, 13421, 14231, 14321 \end{aligned} \tag{13}$$

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 conducted a total of three 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.[23] 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 she 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.[24]

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 F.

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. Our data deviate substantially from the Nash equilibrium predictions 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 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 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 explore how different subsamples of the data contribute to rejection of QRE in the pooled data. Formally, this is no longer a test of QRE, because

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

[24]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 quantal response behavior 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	28.798
Row cycles	68.194	26.953
Col cycles	0.000	7.760

Subject-by-subject:	**Avg** AS	**Avg** $c^R_{0.95}$	**# rejected** at 5%	at 10%	at 20%	**Avg CM violations** (% of total)
Subj. v. self[a]						
(Total subj.: 96)						
All cycles	212.570	18.896	28	36	41	41.59
Row cycles	203.108	11.950	20	27	31	44.53
Column cycles	9.462	11.734	18	20	21	38.65
Subj. v. others[b]						
(Total subj.: 96)						
Row cycles	3.936	11.011	7	13	16	35.00
(Total subj.: 96)						
Col cycles	103.872	11.732	16	18	21	38.70
Subj. v. beliefs[c]						
(Total subj.: 59)						
Row cycles	3.681	6.861	4	5	5	33.05
(Total subj.: 61)						
Col cycles	9.622	8.653	11	13	14	35.00

Notes. All computations use $R = 5,000$. 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.

[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).

by restricting attention to subsets of cycles, pertaining to one player role or one individual subject, in subject-by-subject tests below, we essentially consider a "one-player equilibrium" version of QRE, which is more akin to a discrete-choice problem, and informs us about consistency of subject behavior with the use of quantal responses. 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.

First, we consider separately the CM inequalities pertaining to Row players and those pertaining to Column players.[25] 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 quantal responses 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-player inequalities are violated more often than Column-player inequalities suggests that Row-player 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 payoffs are the same across all the games in our experiment for the Column player, but vary across games for the Row player.

Some intuition for this may come from considering the nature of the Nash equilibria in these games. The (unique) mixed-strategy Nash equilibrium prescribes mixing probabilities which are the same across games for the Row player, but vary across games for the Column player – since in equilibrium, each player's mixed-strategy probabilities are chosen so as to make the *opponent* indifferent between their actions (cf. Goeree and Holt (2001)). This is clear from Table 1, where the Nash equilibrium probabilities are given. The greater degree of violations observed for the Row players may reflect an intrinsic misapprehension of this somewhat paradoxical logic of Nash equilibrium play, and sensitivity to change in own payoffs that carries over from Nash to noisier quantal response equilibria.[26]

[25] 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.

[26] From a theoretical point of view, this observation is also consistent with Golman (2011) who showed that in heterogeneous population games a representative agent for a pool of individuals may not be described by a structural quantal response model even if all individuals use quantal response functions (in the sense that with at least four pure strategies there is no i.i.d payoff shock structure that generates the representative quantal response function). Games with three pure strategies like in our experiment usually fail to have a representative agent, too, indicating a need to take into account heterogeneity of the player roles.

Our experimental design, which consists of a sequence of games in which the Row-player payoffs vary, but not those of the Column player, allows us to consider situations which are more general than those allowed for under Assumption 1 (Invariance). Specifically, such a setting (suggested also by HHK) allows for a weaker notion of invariance, in which a player's shock distribution can vary with her own payoffs, but not with those of her opponent. We call this *weak invariance*. Since the Column-player baseline payoffs are the same in all games, the Row-player shock distribution is trivially invariant to any change to Column's payoffs (but may depend on the change to Row's payoffs). Therefore, for the Row player, weak invariance always holds. For the same reason, the Column-player shock distribution is trivially invariant to any change in his own payoffs (but may depend on the change to the Row's payoffs). Thus, for the Column player, weak invariance must be imposed as an extra assumption. Combining these observations, with our design, only weak invariance for the Column player is required.

The lower panel of Table 3 considers tests of quantal responses for each subject individually; 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 subject's 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 opponent's play.

The results appear largely robust across these three alternative ways of accounting for subjects' beliefs. We see that we are not able to reject the quantal response behavior 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 than the Column inequalities, while the Column inequalities generate more violations in the subject vs. others results. This pattern is consistent with our findings for the pooled data: Since it is the Row-player inequalities that are predominantly violated in the pooled data, these violations only intensify in subject vs. self for the Row cycles. In subject vs. others, however, the inequalities are computed under beliefs corresponding to the average opponent's actual behavior in the *opponent role*, e.g., for Row players the average Column behavior and vice versa. The set of all cyclic monotonicity inequalities is then further partitioned into Row and Column cycles. Thus for Row players the Row cycles reflect the Column behavior, which is more consistent with QRE and so shows less violations than the Column cycles for the Column player, which reflect the more erratic Row-player behavior.

The general trend of these findings – that the quantal responses appear more statistically plausible once we allow for sufficient heterogeneity across subjects and across roles – confirms existing results in McKelvey et al. (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.

Non-equilibrium beliefs. In the previous results, we made the assumption that subjects' beliefs were "in equilibrium", and computed these beliefs as functions of their opponents' (or their own) observed choice frequencies. However, there are alternative behavioral models of non-equilibrium beliefs, which have been prominent in the existing experimental and theoretical literature. Perhaps the most well-known of these is the "level-k" model (and a closely related "cognitive hierarchy" model), which posits that players best respond to beliefs that their opponents are playing at lower levels of rationality (Stahl and Wilson, 1995; Nagel, 1995; Camerer et al., 2004; Crawford et al., 2013). Specifically, the simplest version of level-k starts from a "zero-rationality" $k = 0$ level at which players are assumed to choose actions at random, and proceeds to successively higher rationality levels $k = 1, 2, \ldots$, which are assumed to play best-responses to beliefs that all their opponents are one level less rational $(k - 1)$.

Here we consider testing a "noisy" version of the level-k model applied to our experimental games, in which players choose noisy best responses (according to the additive quantal response framework), believing that their opponents have lower rationality level than they have. We consider two levels: $k = 1$ (so that all subjects view their opponents as randomizing non-rational types); and $k = 2$ (so that all subjects view their opponents as level $k = 1$). In Table 4 we list the beliefs for each level, for each of the four games in our experiment.

The results are reported in Table 5. At the top of the table, we see that, when we pool results across all subjects, we soundly reject the model that all players are $k = 2$. When we break down the test statistic into Row vs. Column players, we see that the rejection is driven by the play of

Table 4: Level-k beliefs for the four games

Role	Own Level	Believes Opponent is	Believes Opponent's playing as follows in Game 1:	Game 2:	Game 3:	Game 4:
Row	$k = 1$	$k = 0$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$
–	$k = 2$	$k = 1$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$
Col	$k = 1$	$k = 0$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$	$1/3, 1/3, 1/3$
–	$k = 2$	$k = 1$	$1/3, 1/3, 1/3$	$0, 0, 1$	$1/2, 1/2, 0$	$1, 0, 0$

Notes. Numbers in the cells are level-k *beliefs* about their *opponent's choice probability* for each of their three available actions. For Row, both level-1 and level-2 believe their opponents randomize uniformly in all games. For Column, this is the case only for Game 1.

the Row players, which is reminiscent of the results from the previous section. However, when we consider the model that all players have rationality level $k = 1$, we are unable to reject: in this case, the test statistic is identically zero for both Row and Column cycles. This implies that a model where each player believes that her rivals are choosing their actions completely randomly cannot be rejected by our data.

These results with pooled data are robust when we test on a subject-by-subject basis. At the bottom portion of Table 5, we see that we can reject the $k = 2$ model for 44 of the 96 subjects (at a 5% significance level), but reject only for 15 subjects for the $k = 1$ model. At face value, these results echo other results in the experimental literature, which likewise suggest that rationality levels calibrated empirically from experimental data are not very high.

Robustness check: Nonlinear utility and risk aversion. Our test results above are computed under the assumption of risk-neutrality. Goeree et al. (2000) showed that allowing for risk aversion greatly improves the fit of logit 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.[27] 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; quantal responses are also rejected when only

[27]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.

Table 5: Testing for Cyclic Monotonicity under Level-k beliefs

Data sample	***AS* test stat**	$c^R_{0.95}$	Rationality level				
All subjects pooled:							
All cycles	1857.220	36.368	k=2				
Row cycles	1857.220	34.444	–				
Col cycles	0.000	6.585	–				
All cycles	0.000	11.787	k=1				
Row cycles	0.000	7.789	–				
Col cycles	0.000	6.549	–				
Subject-by-subject:	**Avg *AS***	**Avg** $c^R_{0.95}$	Rationality level	**# rejected** at 5%	at 10%	at 20%	**Avg CM violations** (% of total)
Subj. v. level-k beliefs[a]							
(Total subj.: 96)							
Row cycles	26.057	12.706	k=2	36	45	51	66.927
Col cycles	5.453	10.260	–	7	11	16	25.938
Row cycles	4.725	11.096	k=1	7	11	15	31.615
Col cycles	5.453	10.226	–	8	11	16	25.938

Notes. All computations use $R = 5,000$. 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. level-k beliefs: the opponent's choice frequencies are level-k beliefs. We checked the two cases: one in which everyone is assumed to be level-1, and the other in which everyone is assumed to be level-2.

the Row cycles are considered; quantal responses 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.[28]

Breaking down these data on a subject-by-subject basis, we once again see that allowing for risk aversion does not change our previous results obtained under the assumption of linear utility. Specifically, as graphed in Figure 4, the number of rejections of the quantal response hypothesis for the "subject v. self" specification is relatively stable for all $r < 0.99$, staying between 28 and 33 rejections at 5% level, at about 38 rejections at 10% level, and between 41 and 44 rejections at 20% level (out of the total of 96 subjects). Thus our analysis here suggests the our test results are not driven by risk aversion.

This robustness to risk aversion might seem surprising at first (e.g., the analysis in Goeree et al. (2003) shows that risk aversion might be an important factor in fitting logit QRE in several bimatrix games). However, introducing constant relative risk aversion has limited effects on the

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

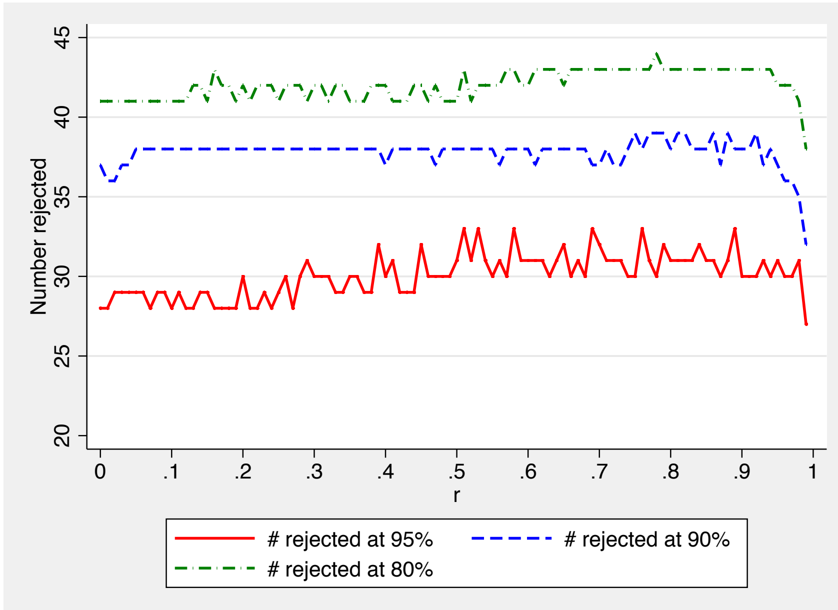

Figure 4: Risk aversion effects on quantal response rejections ("subject v. self", total subjects: 96).

cyclic monotonicity inequalities as it only changes the scale of payoffs, with little effects on the relative difference between payoffs across games in a cycle, whenever the relative differences are not large to begin with, as in our games. Similarly, our test results depend less on the exact form of the utility function as long as differences in payoffs across the games are not too dramatic.

Of course, 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 nonparametric approach for testing the QRE hypothesis in finite normal form games. We go far beyond consistency with the usual logit QRE by allowing players to use any structural quantal response function satisfying mild regularity conditions. This flexibility comes at the cost of requiring the payoff shocks to be fixed across a series of games. The testing approach is based on 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 and Tang (2012), or Liu et al. (2017)). 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

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

Andrews, D. W. K. and G. Soares, "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection," *Econometrica* 78 (2010), 119–157.

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

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

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

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

Berger, A., R. Müller and S. H. Naeemi, "Characterizing Incentive Compatibility for Convex Valuations," in M. Mavronicolas and V. G. Papadopoulou, eds., *SAGT: International Symposium on Algorithmic Game Theory* (Springer Berlin Heidelberg, 2009), 24–35.

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

Camerer, C. F., *Behavioral Game Theory: Experiments in strategic interaction* (Massachusetts: Cambridge: Princeton University Press, 2003).

Camerer, C. F., T.-H. Ho and J.-K. Chong, "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics* 119 (2004), 861–898.

Camerer, C. F., S. Nunnari and T. R. Palfrey, "Quantal Response and Nonequilibrium Beliefs Explain Overbidding in Maximum-Value Auctions," *Games and Economic Behavior* 98 (2016), 243–263.

Chiong, K. X., A. Galichon and M. Shum, "Duality in Dynamic Discrete-Choice Models," *Quantitative Economics* 7 (2016), 83–115.

Chiong, K. X. and M. Shum, "Random Projection Estimation of Discrete-Choice Models with Large Choice Sets," *Forthcoming, Management Science* (2016).

Cominetti, R., E. Melo and S. Sorin, "A payoff-based learning procedure and its application to traffic games," *Games and Economic Behavior* 70 (2010), 71–83.

Crawford, V. P., M. A. Costa-Gomes and N. Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature* 51 (2013), 5–62.

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

Del Boca, D., C. Flinn and M. Wiswall, "Household Choices and Child Development," *Review of Economic Studies* 81 (2014), 137–185.

Fox, J. T., "Semiparametric Estimation of Multinomial Discrete-Choice Models using a Subset of Choices," *RAND Journal of Economics* 38 (2007), 1002–1019.

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

Goeree, J. K., C. A. Holt and T. R. Palfrey, "Risk Averse Behavior in Asymmetric Matching Pennies Games," Working paper, University of Virginia, 2000.

———, "Risk Averse Behavior in Generalized Matching Pennies Games," *Games and Economic Behavior* 45 (2003), 97–113.

———, "Regular Quantal Response Equilibrium," *Experimental Economics* 8 (2005), 347–67.

———, *Quantal Response Equilibrium: A Stochastic Theory of Games* (New Jersey: Princeton: Princeton University Press, 2016).

Golman, R., "Quantal response equilibria with heterogeneous agents," *Journal of Economic Theory* 146 (2011), 2013–2028.

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

Hofbauer, J. and W. H. Sandholm, "On the global convergence of stochastic fictitious play," *Econometrica* 70 (2002), 2265–2294.

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

Liu, N., Q. Vuong and H. Xu, "Rationalization and Identification of Discrete Games with Correlated Types," *Journal of Econometrics* 201 (2017), 249–268.

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

———, "Econometric Models of Probabilistic Choice," in C. Manski and D. McFadden, eds., *Structural Analysis of Discrete Data with Economic Applications* (Massachusetts, Cambridge: The MIT Press, 1981), 198–272.

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

McKelvey, R. D. and T. R. Palfrey, "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior* 10 (1995), 6–38.

McKelvey, R. D., T. R. Palfrey and R. A. Weber, "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 (2000), 523–548.

Merlo, A. and T. R. Palfrey, " External Validation of Voter Turnout Models by Concealed Parameter Recovery," Working paper, 2013.

Nagel, R., "Unraveling in Guessing Games: An Experimental Study," *American Economic Review* 85 (1995), 1313–1326.

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

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

Rockafellar, T. R., *Convex Analysis* (New Jersey: Princeton: Princeton University Press, 1970).

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

Romano, J. P., A. M. Shaikh and M. Wolf, "A practical two-step method for testing moment inequalities," *Econometrica* 82 (2014), 1979–2002.

Rust, J., "Structural Estimation of Markov Decision Processes," in R. F. Engle and D. McFadden, eds., *Handbook of Econometrics* (Elsevier, North-Holland, 1994), 3082–3146.

Saks, M. and L. Yu, "Weak Monotonicity Suffices for Truthfulness on Convex Domains," in *Proceedings of the 6th ACM conference on Electronic commerce* (Vancouver, Canada: ACM, 2005), 286–293.

Shi, X., M. Shum and W. Song, " Estimating Semi-parametric Panel Multinomial Choice Models using Cyclic Monotonicity," *Forthcoming, Econometrica* (2018).

Stahl, D. O. and P. W. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10 (1995), 218–254.

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

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

ONLINE APPENDIX FOR "TESTING THE QUANTAL RESPONSE HYPOTHESIS"

BY EMERSON MELO, KIRILL POGORELSKIY, AND MATTHEW SHUM

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 specially perturbed game. It can be shown[29] that the set of QRE corresponds to the set of NE of a game where players' payoffs are given by

$$\mathcal{G}^i(\boldsymbol{\pi}_i;\boldsymbol{u}_i) := \langle \boldsymbol{\pi}_i, \boldsymbol{u}_i \rangle - \tilde{\varphi}^i(\boldsymbol{\pi}_i), \quad \forall i \in N, \tag{14}$$

with $\tilde{\varphi}^i(\boldsymbol{\pi}_i)$ corresponding to the Fenchel-Legendre conjugate (hereafter convex conjugate) of the function $\varphi^i(\boldsymbol{u}_i)$ defined in (2).[30] By fundamental properties of the conjugacy relationship (cf. (Rockafellar, 1970, pp. 103–104)), we also have

$$\varphi^i(\boldsymbol{u}_i) = \sup_{\boldsymbol{\pi_i}} \mathcal{G}^i(\boldsymbol{\pi}_i;\boldsymbol{u}_i) = \sup_{\boldsymbol{\pi_i}} \left[\langle \boldsymbol{\pi}_i, \boldsymbol{u}_i \rangle - \tilde{\varphi}^i(\boldsymbol{\pi}_i)\right], \quad \forall i \in N, \tag{15}$$

hence $\forall i \in N$ and $\forall \boldsymbol{\pi}_i \in \Delta(S_i)$, Fenchel's inequality holds:

$$\varphi^i(\boldsymbol{u}_i) \geq \langle \boldsymbol{\pi}_i, \boldsymbol{u}_i \rangle - \tilde{\varphi}^i(\boldsymbol{\pi}_i)$$

with equality holding at $\boldsymbol{\pi}_i^* = \operatorname{argsup}_{\boldsymbol{\pi}_i} \mathcal{G}^i(\boldsymbol{\pi}_i;\boldsymbol{u}_i)$.

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

$$\langle [\boldsymbol{u}_i]^{m+1}, [\boldsymbol{\pi}_i^*]^m \rangle - \tilde{\varphi}^i([\boldsymbol{\pi}_{\boldsymbol{i}}^*]^m) \leq \langle [\boldsymbol{u}_i]^{m+1}, [\boldsymbol{\pi}_i^*]^{m+1} \rangle - \tilde{\varphi}^i([\boldsymbol{\pi}_i^*]^{m+1}), \quad \forall m. \tag{16}$$

Rewriting as

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

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

CM implies utility maximization: Suppose that CM holds and let $[\boldsymbol{u}_i]^m$ denote expected utility in an unperturbed game m. By (Rockafellar, 1970, Theorems 24.8 and 24.9) we know that there exists a closed convex function ζ^i such that $[\boldsymbol{\pi}_i^*]^m = \partial\zeta^i([\boldsymbol{u}_i]^m)$. By (Rockafellar, 1970, Theorem 23.5) since $[\boldsymbol{\pi}_i^*]^m$ is in the subdifferential of ζ^i evaluated at $[\boldsymbol{u}_i]^m$, the Fenchel's inequality is satisfied as an equation at $([\boldsymbol{\pi}_i^*]^m, [\boldsymbol{u}_i]^m)$, i.e.,

$$\langle [\boldsymbol{\pi}_i^*]^m, [\boldsymbol{u}_i]^m \rangle - \tilde{\zeta}^i([\boldsymbol{\pi}_i^*]^m) = \zeta^i([\boldsymbol{u}_i]^m).$$

Condition b^*) of (Rockafellar, 1970, Theorem 23.5) implies that $[\boldsymbol{\pi}_i^*]^m \in \arg\sup_{\boldsymbol{\pi}_i}\{\langle \boldsymbol{\pi}_i, [\boldsymbol{u}_i]^m \rangle - \tilde{\zeta}^i(\boldsymbol{\pi}_i)\}$, i.e.,

$$\zeta^i([\boldsymbol{u}_i]^m) = \sup_{\boldsymbol{\pi_i}} \left[\langle \boldsymbol{\pi}_i, [\boldsymbol{u}_i]^m \rangle - \tilde{\zeta}^i(\boldsymbol{\pi}_i)\right]$$

Comparing this with Eq. (15) above, by (Rockafellar, 1970, Theorem 24.9), we conclude that, up to an additive constant, the function ζ (resp. $\tilde{\zeta}$) is identical to φ (resp. $\tilde{\varphi}$), therefore $[\boldsymbol{\pi}_i^*]^m \in \arg\sup_{\boldsymbol{\pi}_i}\{\mathcal{G}^i(\boldsymbol{\pi}_i;[\boldsymbol{u}_i]^m)\}$.

[29] In particular, Cominetti et al. (2010, Prop. 3) establish a way of representing a Logit QRE (and more generally, a structural QRE) as a Nash equilibrium of the game with payoffs (14). See also Hofbauer and Sandholm (2002, Thm. 2.1) for a related result in the single-agent case.

[30] Formally, for a convex function f the Fenchel-Legendre conjugate is defined by $\tilde{f}(\boldsymbol{y}) = \sup_{\boldsymbol{x}}\{\langle \boldsymbol{y}, \boldsymbol{x} \rangle - f(\boldsymbol{x})\}$.

That is, CM implies utility maximization in the perturbed game for each m. □

Intuitively, the set of inequalities (16) can be seen as set of incentive compatibility constraints across the series of games that only differ in the payoffs. This means that our CM conditions capture players' optimization behavior with respect to changes in expected payoffs across such games.

B. PROOF OF PROPOSITION 2

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,

$$\sum_{j=1}^{J_i}(u_{ij}^1 - u_{ij}^0)(\pi_{ij}^0 - \pi_{ij}^1) \leq 0. \tag{17}$$

Note that cyclic monotonicity inequality in Eq. (17) is distinct from HHK's cumulative rank condition (8). Suppose that the right-hand side 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

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

where the last inequality follows from HHK condition for $k = 2$ and $u_{i2}^1 - u_{i2}^0 \geq 0$. Repeating the same procedure for $k = 3, \ldots, J_i$, we obtain the CM condition (17) for $M = 2$.

Conversely, suppose that (17) holds. For the case of two games, (17) 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 (17), 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. $\pi_{i1}^1 - \pi_{i1}^0 < 0$. Since (17) holds, the probabilities in both games are generated by a QRE. Due to indexing in (7),

$$u_{i1}^1 - u_{ij}^1 \geq u_{i1}^0 - u_{ij}^0$$

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

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[31] axioms (Goeree et al., 2005, p.355). Interiority, Continuity, and Responsiveness are satisfied automatically under the *structural approach* to quantal response[32] 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

[31] This property is called Monotonicity in Goeree et al. (2005).

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

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 $j, k \in \{1, \ldots, J_i\}$:

$$(u_{ij}(\mathbf{p}) - u_{ik}(\mathbf{p}))(\pi_{ij}(\mathbf{p}) - \pi_{ik}(\mathbf{p})) \geq 0 \tag{18}$$

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 (18) for each game.[33]

Alternatively, the Rank-order Axiom can be imposed from the outset by making an extra assumption about the shock distributions. In particular, Goeree et al. (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.

Proposition 3. *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 3 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.[34] Expected utility of Col from choosing each of her three pure strategies in any of Games 1–4 (see payoffs in Table 1) is

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

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

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

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

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

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. Row's expected utility in this game from choosing each of her three pure strategies is

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

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. Row's expected utility in this game from choosing each of her three pure strategies is

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

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

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

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, $m \in \{1, \ldots, 4\}$.[35]

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 (13)) 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

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

form. The partial derivatives wrt π_{i1}^m are

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

The partial derivatives wrt π_{i2}^m are

$$
\begin{aligned}
\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} &= 0 && \text{for } \ell \in S_0^m \\
\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} &= \big[\pi_{k1}^{m+1} u_i^{m+1}(s_{i2}, s_{k1}) - \pi_{k1}^m u_i^m(s_{i2}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{i2}, s_{k2}) - \pi_{k2}^m u_i^m(s_{i2}, s_{k2}) \\
&+ (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{i2}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{i2}, s_{kJ})\big] \\
&- \big[\pi_{k1}^{m+1} u_i^{m+1}(s_{iJ}, s_{k1}) - \pi_{k1}^m u_i^m(s_{iJ}, s_{k1}) + \pi_{k2}^{m+1} u_i^{m+1}(s_{iJ}, s_{k2}) - \pi_{k2}^m u_i^m(s_{iJ}, s_{k2}) \\
&+ (1 - \pi_{k1}^{m+1} - \pi_{k2}^{m+1}) u_i^{m+1}(s_{iJ}, s_{kJ}) - (1 - \pi_{k1}^m - \pi_{k2}^m) u_i^m(s_{iJ}, s_{kJ})\big] && \text{for } \ell \in S_i^m \\
\frac{\partial \nu_\ell}{\partial \pi_{i2}^m} &= \pi_{k1}^m \left[-u_k^m(s_{i2}, s_{k1}) + u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^m \left[-u_k^m(s_{i2}, s_{k2}) + u_k^m(s_{iJ}, s_{k2})\right] \\
&+ (1 - \pi_{k1}^m - \pi_{k2}^m) \left[-u_k^m(s_{i2}, s_{kJ}) + u_k^m(s_{iJ}, s_{kJ})\right] \\
&+ \pi_{k1}^{m\ominus 1} \left[u_k^m(s_{i2}, s_{k1}) - u_k^m(s_{iJ}, s_{k1})\right] + \pi_{k2}^{m\ominus 1} \left[u_k^m(s_{i2}, s_{k2}) - u_k^m(s_{iJ}, s_{k2})\right] \\
&+ (1 - \pi_{k1}^{m\ominus 1} - \pi_{k2}^{m\ominus 1}) \left[u_k^m(s_{i2}, s_{kJ}) - u_k^m(s_{iJ}, s_{kJ})\right] && \text{for } \ell \in S_k^m
\end{aligned}
$$

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 completeness we list the cycle index subsets for each game $m \in \{1, .., 4\}$ in Table 6.

E. TESTING CYCLIC MONOTONICITY USING THE TWO-STEP APPROACH BY ROMANO, SHAIKH, AND WOLF (2014).

As an extra robustness check, we tested cyclic monotonicity conditions using another recent procedure developed by Romano et al. (2014) (henceforth referred to as RSW) for testing moment inequalities. Here, we describe the simplified version from the supplemental appendix of RSW, which applies to our setting. This is the case where we have a random vector $\boldsymbol{\mu} \in \mathbb{R}^P$ which is (asymptotically) distributed $N(\boldsymbol{\mu_0}, \Sigma)$, where $\boldsymbol{\mu_0}$ is unknown. We want to test

$$H_0: \ \boldsymbol{\mu_0} \geq \mathbf{0} \quad vs. \quad H_1: \ \boldsymbol{\mu_0} \ngeq \mathbf{0},$$

where $\mathbf{0} \in \mathbb{R}^P$. The relevant algorithm is described in section S.1.2 in the supplemental appendix of RSW.

First, for ease of comparison with RSW, we redefine $\boldsymbol{\mu} = -\boldsymbol{\mu}$, because RSW consider the complementary hypothesis

$$H_0: \ \boldsymbol{\mu_0} \leq \mathbf{0} \quad vs. \quad H_1: \ \boldsymbol{\mu_0} \nleq \mathbf{0}.$$

Letting $\hat{\Sigma}$ denote an estimate of Σ, we use the test statistic $T(\hat{\boldsymbol{\mu}}, \hat{\Sigma}) = \sum_{j=1}^P \left[\frac{\hat{\mu}_j}{\hat{\sigma}_j}\right]_+^2$, where $[x]_-$ denotes $x \cdot \mathbb{1}(x < 0)$, and $\hat{\sigma}_j$ is the square-root of the j-th diagonal entry in $\hat{\Sigma}$.

Table 6: 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.

There are now two choice variables: α, the size of the test, which would be 0.05; and β, the size of the first-step "moment selection" procedure, $0 \leq \beta < \alpha$. β should be small – we use $\beta = 0.01$. There are three main steps (see S.1.2 in RSW for notation):

1. Compute the quantity $K^{-1}(1-\beta)$ by simulation: draw vectors $Z_s \sim N(0, \Sigma)$, for $s = 1, \ldots S$, and for each one, compute the maximal element $Z_s^* = \max(Z_{s,1}, \ldots, Z_{s,P})$. Then $K^{-1}(1-\beta)$ is the $(1-\beta)$-th quantile among $\{Z_s^*\}_{s=1}^S$.
2. Compute the vector $\tilde{\boldsymbol{\mu}}$ where, for $j = 1, \ldots P$:
$$\tilde{\mu}_j = \min\left\{\mu_j + K^{-1}(1-\beta), 0\right\}$$
3. Evaluate the $(1-\alpha+\beta)$ quantile of the test statistic $T(\cdot, \hat{\Sigma})$. We do this by simulation: draw vectors $X_s \sim N(\tilde{\mu}, \hat{\Sigma})$, for $s = 1, \ldots S$, and for each one, compute the test statistic $T_s \equiv T(X_s, \hat{\Sigma})$. Then the critical value is the $(1-\alpha+\beta)$-th quantile among $\{T_s\}_{s=1}^S$.

Table 7 shows the test results of checking the cyclic monotonicity conditions with our experimental data using the RSW approach. We see that the results are unchanged from those reported in Section 6.2 using the Andrews-Soares (2010) procedure.

Table 7: Testing for Cyclic Monotonicity in Experimental Data: Using Romano-Shaikh-Wolf (2014) Procedure

Data (all subjects)	**Test statistic**	$c^S_{1-\alpha+\beta}$
All cycles	68.194	39.436
Row cycles	68.194	31.594
Col cycles	0.000	5.119

Notes. All computations use $S = 5,000$. $\alpha = 0.05, \beta = 0.01$.

F. EXPERIMENT INSTRUCTIONS

The instructions in the experiment, given below, largely follow McKelvey et al. (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.