Games and Economic Behavior xxx (xxxx) xxx

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Non-parametric identification and testing of quantal response equilibrium ☆

Johannes Hoelzemann [a], Ryan Webb [b], Erhao Xie [c,*]

[a] University of Vienna, Austria
[b] University of Toronto, Canada
[c] Bank of Canada, Canada

A R T I C L E   I N F O

JEL classification:
C14
C57
C92

Keywords:
Quantal response equilibrium (QRE)
Nash equilibrium
Falsifiability
Non-parametric identification
Experiment
Matching pennies

A B S T R A C T

This paper studies the falsifiability and identification of Quantal Response Equilibrium (QRE) when each player's utility and error distribution are relaxed to be unknown non-parametric functions. Using the variation of players' choices across a series of games, we first show that both the utility function and the distribution of errors are non-parametrically over-identified. This over-identification result further suggests a straightforward testing procedure for QRE which achieves the desired type-1 error and maintains a small type-2 error. To apply this methodology, we conduct an experimental study of the matching pennies game. Our non-parametric estimates strongly reject the conventional Logit choice probability. Moreover, when the utility and the error distribution are sufficiently flexible and heterogeneous, the quantal response hypothesis cannot be rejected for 70% of participants. However, strong assumptions such as linear utility, logistically distributed errors, and homogeneity lead to substantially higher rejection rates.

1. Introduction

In many strategic settings, choice behavior systematically deviates from the canonical solution concept of Nash Equilibrium (NE). These deviations have been documented using both experimental data of individual decision-makers (Goeree and Holt, 2001) and field data of firms and managers (Goldfarb and Xiao, 2011; Aguirregabiria and Jeon, 2020). To address some of these failures, *Quantal Response Equilibrium* (QRE; McKelvey and Palfrey, 1995; Goeree et al., 2005) has been proposed as an alternative equilibrium concept. It has successfully explained many deviations from NE and has become an important benchmark in game theory (Goeree et al., 2020).[1]

This paper focuses on *structural* QRE (McKelvey and Palfrey, 1995) which extends the random utility framework to strategic settings. In particular, the expected utility of each action is randomly perturbed by an additively separable "error" that can be interpreted as either a noisy decision process or the private information of each player. A structural QRE is then defined as a fixed

☆ We thank Victor Aguirregabiria, Jose Apesteguia, Ben Greiner, Yoram Halevy, Sota Ichiba, Shengwu Li, Jean-Robert Tyran, Yuanyuan Wan, and seminar participants at the Bank of Canada, University of Toronto, Xi'an Jiaotong University, the ESA, MEG, and SBEER conferences for helpful discussions. Financial support by TD-MDAL and the University of Toronto is gratefully acknowledged. The experiment in this study was approved by the Vienna Center for Experimental Economics. The views in this paper do not necessarily represent the views of the Bank of Canada.

* Corresponding author.
URLs: http://www.johanneshoelzemann.com (J. Hoelzemann), http://ryan-webb.com (R. Webb), https://sites.google.com/view/erhaoxie (E. Xie).

[1] For recent work on axiomatizing and endogenizing QRE, see Friedman (2020); Friedman and Mauersberger (2022). Allen and Rehbeck (2021) introduce non-expected utility preference into QRE. For an order-theoretic approach to QRE and an application to coordination in networks, see Hoelzemann and Li (2022).

https://doi.org/10.1016/j.geb.2024.07.004
Received 28 February 2023

Available online 22 July 2024

point in the space of the choice probabilities implied by this error. By incorporating errors into strategic settings, yet preserving the concept of equilibrium, QRE makes predictions about behaviors in games that reduce to NE as noise vanishes.

In a structural QRE, each player's behavior (choice probability) is completely determined by two model primitives: (i) the utility function and (ii) the error distribution. Empirical applications of QRE typically impose strong, restrictive, and potentially mis-specified assumptions on these two primitives. For instance, the analyst usually assumes that each player's utility function is known, identical, and given by monetary payoffs in the experiment (henceforth, the *known utility* assumption). This assumption restricts all participants to have homogeneous risk-neutral preferences. It is problematic since heterogeneous behavior consistent with deviations from risk-neutrality (or at least, a linear aggregation of monetary payoffs) is typically observed in laboratory settings, as identified by Goeree et al. (2003), Harrison and Cox (2008), and Oprea (forthcoming), among many others. Further, most applications also assume that each player's random errors follow a common distribution and the functional form of this distribution is known by the analyst, e.g., the Logit choice rule. This *distributional* assumption is considered mainly due to its statistical convenience and it imposes strong shape restrictions that could be mis-specified, especially when the analyst fits aggregate data that consists of heterogeneous participants (Golman, 2011).[2]

This paper addresses the identification and testing of structural QRE when relaxing the above restrictions on model primitives. In particular, we specify each player's utility to be a *non-parametric function* of their monetary payoffs received in the experiment. In addition, within each player, the random errors associated with each action are jointly distributed according to a *non-parametric* function. Crucially, this distribution function allows for general error structures, where the random errors of each action may follow heterogeneous marginal distributions and exhibit arbitrary correlations with the errors of other actions.[3] Given this empirical framework, we focus on experimental settings where the analyst can design a series of games with different monetary rewards that correspond to different exogenous treatments in the experiment. Each game could be played either once or with multiple repetitions. Under an *invariance* assumption that the utility function and the distribution function remain unchanged across games, we show that each player's error distribution is *non-parametrically over-identified* when there are sufficient and independent variations of each player's monetary rewards.

The above non-parametric over-identification result has four important implications. First, to derive the QRE choice probabilities, an analyst does not have to rely on strong and potentially mis-specified distributional assumptions. Instead, the analyst can simply estimate the error distribution and this empirical estimate is robust to any preferences over own monetary rewards within the expected utility framework. Notably, a non-parametric specification, if performed at the population level, can be interpreted as a *heterogeneous* QRE that allows the error distribution to vary across participants (Golman, 2011). Such a proposal has not previously been applied to data due to a lack of identification results. The results reported here therefore provide a means to fit heterogeneous QRE at the population level. Second, since the model primitives are over-identified, it implies that QRE can be tested employing a standard over-identification test. This test addresses the non-falsifiability of QRE as raised by Haile et al. (2008), who show that when the random errors are not i.i.d. across players' actions and are non-parametrically specified, QRE can rationalize any vector of choice probabilities and is therefore non-falsifiable *within* a game. In contrast, we show that under the invariance assumption, the variations of players' choices *across* games can provide enough identification power to test QRE.[4] Third, most empirical applications of QRE assume that the mean or median of the error distribution is zero. We show that this restriction is not necessary for identification in most experimental datasets. This result allows the analyst to identify the existence of systematic errors displayed by participants. Fourth, once the distribution function has been identified, our empirical framework then reduces to a semi-parametric model where the utility function remains non-parametric but the error distribution is known by the analyst. This semi-parametric model has been studied by Bajari et al. (2010) and Aguirregabiria and Xie (2021). Based on their results, non-parametric identification of each player's utility function is indeed feasible.

To estimate and test QRE in practice, we exploit non-parametric Maximum Likelihood estimation by the method of sieves. We illustrate the finite sample property of this method in a Monte Carlo experiment to highlight the importance of relaxing both the known utility and the distributional assumptions. When either of these assumptions is mis-specified and behavior is generated by QRE, we find that the test of QRE is substantially over-rejected in typical sample sizes of laboratory studies. In particular, the type-1 error rates may exceed 90% on a purported 5% test. In contrast, under a fully non-parametric specification, our test achieves the ideal type-1 error rates and therefore guards against over-rejection of QRE. Moreover, the estimates of both the utility and the error distribution closely match their true values. Finally, we also simulate the data under alternative behavioral models such as the canonical Level-k model. In this scenario, our test has the power to reject the incorrect null hypothesis of QRE with a rejection rate close to 100%.

To assess the empirical relevance of our results, we conduct a laboratory experiment of the matching pennies game. We find that QRE under a fully non-parametric specification fits the data substantially better than existing applications, both in-sample and out-of-sample. We also observe a substantial reduction in rejections of quantal response behavior, with rejection rates dropping from 70% to 30% of participants. The estimation results also strongly reject the usual Logit choice probability. In comparison to a logistic distribution, our estimated error distribution exhibits a higher probability of both small and extremely large errors, coupled with a smaller probability of errors of moderate size. Moreover, the estimated mean of the errors is significantly positive, suggesting that

[2] In particular, suppose that all individuals' errors follow the extreme type-1 distribution (i.e., Logit) but differ in their sensitivity parameters. Golman (2011) shows that the aggregate behavior could be described by a representative player who *will not* behave according to the Logit formula. The actual error distribution of this representative player depends on the distribution of the sensitivity parameters.

[3] Of course, the correlation structure has to be permitted by a valid joint distribution function.

[4] This idea of exploiting cross-game variation was first conjectured by Haile et al. (2008).

participants in our experiment display a systematic error: they tend to mistakenly choose the action presented at the top of the screen more frequently. Finally, under the conventional Logit specification, the known utility assumption is highly rejected. In contrast, with a non-parametric error distribution, a non-linear utility function performs quantitatively similar to a linear utility function (albeit with possible curvature for higher payoffs). All these results highlight the importance of flexible specifications of model primitives and the non-parametric identification results derived in this paper are particularly useful.

This paper relates to two studies that exploit cross-game variations to test QRE. Melo et al. (2019) consider a non-parametric distribution function but impose the known utility assumption. Aguirregabiria and Xie (2021) specify a non-parametric utility function but maintain the distributional assumption. This paper jointly relaxes both assumptions and derives a test of QRE. Moreover, we obtain non-parametric identification results of both model primitives. In contrast, Melo et al. (2019) do not study the identification problem and Aguirregabiria and Xie (2021) only report a semi-parametric result.[5]

In contrast to testing structural QRE through cross-game variation, Goeree et al. (2005) propose a reduced-form approach known as *regular* QRE. This approach avoids modeling random errors and directly defines a reduced-form quantal response function that maps from expected utilities to choice probabilities. By imposing four specific conditions on this function, Goeree et al. (2005) derive testable implications of regular QRE *within* repetitions of the same game.[6] Notably, these testable restrictions rely on the known utility assumption, which we relax in this paper. Moreover, when the random errors are i.i.d. across players' actions and satisfy some other weak conditions, structural QRE is nested in regular QRE. However, since we relax the i.i.d. restriction and allow for general error structures, regular QRE does not nest the class of structural QRE considered in this paper, and vice versa.

Structural QRE shares an identical mathematical structure with Bayesian Nash Equilibrium (BNE) in incomplete information games where private information is independent across players. The identification of the latter framework, mainly using field data, has been extensively studied. In particular, Bajari et al. (2010) consider a non-parametric utility function but maintain the distributional assumption. Liu et al. (2017) focus on binary choice games and further relax the distributional assumption, achieving the fully non-parametric identification for both model primitives. In a recent paper, Xie (2022) extends these fully non-parametric results to multinomial choice games and attains identification even when players occasionally deviate from equilibrium. He also derives a testable implication of BNE. Our results advance these results in four directions. First, Xie (2022) considers a restrictive class of error distributions. In contrast, our experimental setting allows for general error structures, where the errors of each action can follow heterogeneous marginal distributions and exhibit arbitrary correlations across actions. Second, the testable implication in Xie (2022) relies on an "equal choice probabilities" condition which is extremely difficult to construct in empirical applications, and thus has not been applied to an actual dataset. Conversely, our results do not require constructing this condition explicitly and are straightforward to implement in practice. We illustrate this using a dataset from a laboratory experiment of the matching pennies game. Third, our test includes not only all the testable restrictions derived by Xie (2022), but our experimental setting allows us to derive many other additional restrictions imposed by QRE. As such, our test has higher statistical power.[7] Finally, Xie's testable implication is a restriction on players' choice probabilities, which are multivariate functions of all players' monetary payoffs. In contrast, our test is a restriction on *model primitives* which are single-variable functions. This dimension reduction ensures precise estimation, especially under a fully non-parametric specification. It consequently improves the finite sample performance of the test of QRE.

There are several plausible explanations for why QRE might not be satisfied in an experimental dataset, such as departures from expected-utility, other-regarding preference, and incorrect / biased beliefs about the other player's behavior. Xie (2022) attributes non-QRE behavior solely to biased beliefs, allowing identification of each player's belief about the other player's choice. In contrast, this paper remains agnostic regarding the underlying factors that lead to violations of QRE. Instead, we aim to derive a test that has the power to reject QRE in the presence of any potential factor that may cause its violation.

The rest of the paper proceeds as follows. Section 2 reviews QRE in 2×2 games, and Section 3 presents the identification results and our test. Generalizations to games with more players and / or more actions require extra notation, and are in the appendix. A Monte Carlo exercise is presented in Section 4 and the laboratory experiment is discussed in Section 5. We conclude in Section 6. Proofs and other extensions are delegated to the appendix.

2. QRE in 2×2 games

Players are indexed by $i \in \{1, 2\}$ and $-i$ represents the other player. Each player i simultaneously chooses an action, denoted as a_i, from their action set $\mathcal{A}_i = \{0, 1\}$. Moreover, let $\mathbf{a} = (a_i, a_{-i}) \in \mathcal{A} = \mathcal{A}_i \times \mathcal{A}_{-i}$ be an action profile of this game. In an experiment, when $\mathbf{a}$ is the chosen profile, player i will receive a monetary payoff that equals $m_i(\mathbf{a})$ in experimental currency units.

We define $u_i(m) : \mathbb{R} \to \mathbb{R}$ as player i's utility function, so that their preference depends on, but is not necessarily equal to, their monetary reward m. Given this function, player i will receive a utility $u_i[m_i(\mathbf{a})]$ when the realized action profile is $\mathbf{a}$. We consider a *non-parametric* specification of $u_i(m)$ which allows for any form of self-regarding preferences for money within the expected-utility

[5] Goeree et al. (2003) is an early attempt to relax the known utility assumption. They assume a parametric utility function and also impose the distributional assumption.

[6] Another important example that exploits this reduced-form approach is the rank-dependent choice equilibrium by Goeree et al. (2019).

[7] In our setting, each player's utility is a function that depends only on their received monetary reward, and this reward varies across action profiles and games. The structure of monetary payoffs leads to model restrictions beyond those typical in field data. For instance, suppose that one player receives the same monetary reward in two action profiles (either in the same game or in different games), our setting implies that this player must receive the same utility. This restriction (and others) allow us to identify a more general error structure and obtain a test with higher statistical power.

Table 1
Monetary payoff matrix of the matching pennies ($m_1 > 8$, $m_2 > 8$) game.

MP		Player 2: 0	Player 2: 1
Player 1	0	m_1, 8	8, 16
	1	8, m_2	16, 8

framework imposed by QRE. Moreover, we allow $u_i(\cdot)$ to be heterogeneous across players or participants in economic experiments. We only impose weak restrictions on $u_i(m)$, as stated by Assumption 1.

Assumption 1. Each player i's utility function $u_i(m)$ is strictly increasing and continuous in m.

In this 2×2 game, let $\mathbf{m}_i$ be a 4×1 vector that equals $\left[m_i(a_i=0,a_{-i}=0), m_i(a_i=0,a_{-i}=1), m_i(a_i=1,a_{-i}=0), m_i(a_i=1,a_{-i}=1)\right]'$. Each element in $\mathbf{m}_i$ represents player i's monetary reward for the corresponding action profile. In our notation of $m_i(\mathbf{a})$, the term m_i is not interpreted as a function that depends on $\mathbf{a}$. Instead, m_i is treated as an observed variable that may vary across games, with $\mathbf{a}$ serving as an index representing the $|\mathbf{a}|^{th}$ variable.[8] In summary, the results in this paper are applicable to any experimental dataset where the analyst can observe each action profile's *outcome variable* $m_i(\mathbf{a})$ and where each player's preferences are defined over the space of such an outcome variable.

Our empirical setting is an experiment that exogenously and independently varies each player's monetary payoffs. Suppose that we fix $\mathbf{m}_i$ and only consider the variation of $\mathbf{m}_{-i}$. Under QRE, player $-i$'s choice probability will exogenously vary due to the variation of $\mathbf{m}_{-i}$. Moreover, the structure of QRE implies that such probability can be viewed as a regressor for player i's decision rule with player i's utility as the corresponding coefficient. Notably, this coefficient remains unchanged due to the fixation of $\mathbf{m}_i$. As we will demonstrate in Section 3, such a regression-based interpretation leads to the identification of the error distribution.

To formally state the above condition of exogenous variation, define $\mathcal{M}_i(\mathbf{a}) \subseteq \mathbb{R}$ as the support of $m_i(\mathbf{a})$ (i.e., the set of all possible values that $m_i(\mathbf{a})$ can take). In addition, let $\mathcal{M}_i \subseteq \times_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ denote the support of player i's own monetary rewards $\mathbf{m}_i$. The exogenous condition is summarized by Assumption 2.

Assumption 2. The following conditions are satisfied for each player i:

(a) For each $\mathbf{a} \in \mathcal{A}$, $\mathcal{M}_i(\mathbf{a})$ is either a singleton or an interval.
(b) There exists at least one $\mathbf{a} \in \mathcal{A}$ such that $\mathcal{M}_i(\mathbf{a})$ is an interval.
(c) Conditional on each $\mathbf{m}_{-i} \in \mathcal{M}_{-i}$, $\mathbf{m}_i$ has exogenous variation over its support $\mathcal{M}_i$.

The structure of $\mathcal{M} \equiv \mathcal{M}_i \times \mathcal{M}_{-i}$ determines the type of game. Assumption 2 allows for general structures. Specifically, it includes experiments that vary every action profile's payoffs across games as well as experiments that only vary some profiles' payoffs (i.e., $\mathcal{M}_i(\mathbf{a})$ is an interval) while holding the payoffs of other profiles constant (i.e., $\mathcal{M}_i(\mathbf{a})$ is a singleton). Assumption 2(b) only requires that, for each player i, there are variations of monetary payoffs for at least one action profile. In addition, across games, player i's monetary rewards of any two profiles could be either independent or exhibit arbitrary correlations.

For instance, Table 1 represents a matching pennies game that independently varies each player's monetary rewards for only one action profile, as represented by variables m_1 and m_2. In contrast, the payoffs of all other profiles remain unchanged across games. This matching pennies game has been studied by Goeree and Holt (2001) among others. We also study this game in our Monte Carlo exercise and empirical application.

Table 2 represents a coordination game that is generated by a different structure of $\mathcal{M}$. Specifically, player i's payoff of the safe action (i.e., $a_i = 0$) does not depend on the other player's choice and varies by the same magnitude across games. In this example, the payoffs of two action profiles are perfectly positively correlated.

To define QRE in this environment, let $p_{-i}(\mathbf{m})$ denote player $-i$'s choice probability of action $a_{-i} = 0$. Since we consider strategic settings, a player's decision depends on all players' monetary rewards $\mathbf{m} = (\mathbf{m}_i', \mathbf{m}_{-i}')'$. Given $p_{-i}(\mathbf{m})$, the expected utility of action a_i for player i is as follows:

$$EU_i[\mathbf{m}_i, a_i, p_{-i}(\mathbf{m})] = u_i[m_i(a_i, a_{-i}=0)] \cdot p_{-i}(\mathbf{m}) + u_i[m_i(a_i, a_{-i}=1)] \cdot [1 - p_{-i}(\mathbf{m})]. \tag{1}$$

This expected utility $EU_i(\cdot, a_i)$ is a *function* that depends on player i's own monetary rewards $\mathbf{m}_i$ and their opponent's $\mathbf{m}_{-i}$ via the other player's choice probabilities $p_{-i}(\mathbf{m})$. We focus on structural QRE that places an additively separable error on this expected utility. Specifically, let $\varepsilon_i(a_i)$ denote the error on player i's expected utility of action a_i. Consequently, player i will choose $a_i = 0$ if and only if the following condition holds:

[8] We define $|\mathbf{a}| = a_i \cdot |\mathcal{A}_{-i}| + a_{-i} + 1$. With this definition, $m_i(\mathbf{a})$ can be equivalently represented as $m_{i,|\mathbf{a}|}$ and $\mathbf{m}_i$ is a vector of four variables in the form of $(m_{i,1}, m_{i,2}, m_{i,3}, m_{i,4})'$. We decided against this alternative representation as it is cumbersome for the proofs of some results.

Table 2
Monetary payoff matrix of the coordination ($0 \leq m_i \leq 15$) game.

CO		Player 2	
		0	1
Player 1	0	m_1, m_2	m_1, 0
	1	0, m_2	15, 15

$$\begin{aligned} & EU_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 0) \geq EU_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 1) \\ \Leftrightarrow \; & \underbrace{\varepsilon_i(a_i = 1) - \varepsilon_i(a_i = 0)}_{=\tilde{\varepsilon}_i} \leq \underbrace{EU_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})] - EU_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})]}_{=\widetilde{EU}_i[\mathbf{m}_i, p_{-i}(\mathbf{m})]}. \end{aligned} \tag{2}$$

To derive the QRE choice probabilities, let $F_i(\cdot)$ be the cumulative distribution function (C.D.F.) of $\tilde{\varepsilon}_i = [\varepsilon_i(a_i = 1) - \varepsilon_i(a_i = 0)]$. We specify $F_i(\tilde{\varepsilon}_i)$ to be a *non-parametric* function that is unknown to the analyst and allow this distribution function to be heterogeneous across players and experimental participants with the following restrictions summarized by Assumption 3.

Assumption 3.

(a) For each player i, $F_i(\tilde{\varepsilon}_i)$ is continuously differentiable and strictly increasing over the real line.
(b) For each player i, $F_i(\tilde{\varepsilon}_i)$ is independent of $(\mathbf{m}_i, \mathbf{m}_{-i})$.

Assumption 3(a) is the standard regularity condition that connects structural QRE to regular QRE (Goeree et al., 2005). In particular, Assumption 3(a) implies that the choice probability of any action a_i is strictly positive (i.e., the interiority condition)[9] and is strictly increasing in this action's expected utility (i.e., the responsiveness condition). It also implies that the quantal response function is continuous and differentiable (i.e., the continuity condition). Most applications of structural QRE further assume that $\varepsilon_i(a_i)$ is i.i.d. across actions so that the difference of errors, denoted as $\tilde{\varepsilon}_i$, has a median of zero. This zero median restriction implies that a_i is chosen more frequently than a'_i if a_i has a higher expected utility than a'_i (i.e., monotonicity condition). Equivalently, Assumption 3(a) and the i.i.d. restriction imply all four conditions of the regular quantal response function so that structural QRE is nested in regular QRE. In contrast, we relax the i.i.d. restriction to allow for systematic errors. For instance, players may consistently over-estimate their expected utility of action a. In this case, the choice probability of a can be higher than a' even when a has a lower *actual* expected utility (i.e., violation of monotonicity). Since we allow for systematic errors, regular QRE does not nest the class of structural QRE considered in this paper, and vice versa.

Assumption 3(b) is known as the *invariance assumption* and requires the error distribution to remain unchanged across games. It is commonly maintained in empirical applications of QRE, including formal tests of QRE (Melo et al., 2019; Goeree et al., 2020; Aguirregabiria and Xie, 2021). In an extension offered in the appendix, we relax this invariance assumption by allowing $F_i(\cdot)$ to depend on player i's own $\mathbf{m}_i$ but to be independent of the other player's $\mathbf{m}_{-i}$.

Given Equation (2) and $F_i(\cdot)$, player i's choice probability of action $a_i = 0$ takes the form of a quantal response function, as presented below:

$$p_i(\mathbf{m}) = F_i\Big[\underbrace{EU_i\big(\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})\big) - EU_i\big(\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})\big)}_{=\widetilde{EU}_i\left(\mathbf{m}_i, p_{-i}(\mathbf{m})\right)}\Big]. \tag{3}$$

In QRE, each player forms correct beliefs about other players' choice probabilities. Consequently, QRE is a fixed point in the space of choice probabilities defined by:

Definition 1. For any $\mathbf{m} \in \mathcal{M}$, the vector $\left(p_i(\mathbf{m}), p_{-i}(\mathbf{m})\right)'$ is a QRE if and only if the following condition holds:

$$p_i(\mathbf{m}) = F_i\Big[\underbrace{EU_i\big(\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})\big) - EU_i\big(\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})\big)}_{=\widetilde{EU}_i\left(\mathbf{m}_i, p_{-i}(\mathbf{m})\right)}\Big], \forall i \in \{1, 2\}, \tag{4}$$

where $F_i(\cdot)$ satisfies Assumption 3.

[9] This interiority condition is the result of $\tilde{\varepsilon}_i$ having full support over $\mathbb{R}$. The full support condition also implies that the variance of $\tilde{\varepsilon}_i$ is strictly positive, though can be arbitrarily small. Therefore our framework does not include NE (i.e., $Var(\tilde{\varepsilon}_i) = 0$) as a special case, but can well approximate it with an arbitrarily small $Var(\tilde{\varepsilon}_i)$. We maintain the full support condition since it is necessary to avoid the zero-likelihood problem.

Since $F_i(\cdot)$ is continuous by Assumption 3(a), any game has at least one QRE according to Brouwer's fixed-point theorem. When multiple QRE exist, we assume that there is a *deterministic* equilibrium selection mechanism that chooses one of these equilibria, as stated in Assumption 4. Therefore, even though there may be multiple equilibria in the model, the analyst only observes a single equilibrium for each game in the data. Except for this deterministic condition, we do not impose any further restrictions on the selection mechanism.

Assumption 4. For any $\mathbf{m} \in \mathcal{M}$, if there are multiple vectors $\left(p_i(\mathbf{m}), p_{-i}(\mathbf{m})\right)'$ that satisfy the conditions in Definition 1 (i.e., multiple QRE), there exists a mechanism that selects one of the vectors.

3. Identification results and the over-identification test

In this section, we demonstrate that both the error distribution and utility function can be non-parametrically over-identified under the QRE restrictions (Equation (4)). This result further implies that the null hypothesis of QRE can be tested employing the standard over-identification test.

To derive these results, we introduce a continuity condition on $p_i(\mathbf{m})$.

Assumption 5. For each player i, the choice probability function $p_i(\mathbf{m})$ varies with both $\mathbf{m}_i$ and $\mathbf{m}_{-i}$. Moreover, $p_i(\mathbf{m})$ is continuous over its support $\mathcal{M}$ with probability 1. If there are points of discontinuity in $\mathcal{M}$, the number of all discontinuous points is finite.

This condition is quite weak. When there is a unique equilibrium for each $\mathbf{m}$ (e.g., matching pennies in Table 1), Assumption 5 trivially holds under the restrictions of QRE by Definition 1 (Aguirregabiria and Mira, 2019). In scenarios where there are multiple QRE (for example, the coordination game in Table 2), Aguirregabiria and Mira (2019) show that every QRE can be classified into a finite number of types. Within each type, QRE choice probabilities are continuous in $\mathbf{m}$. Consequently, $p_i(\mathbf{m})$ is discontinuous only when players switch between equilibrium types. Importantly, Assumption 5 allows players to select an equilibrium in any arbitrary way, and only restricts the number of equilibrium switching points to be finite.[10] Assumption 5 also holds in alternative behavioral models such as Level-k (Nagel, 1995), cognitive hierarchy (Camerer et al., 2004), and more general iterative reasoning models (Halevy et al., 2023). In these models, the continuity of the error distribution and the utility function directly implies the continuity of $p_i(\mathbf{m})$.

To derive the identification results, we focus directly on each player's choice probability and assume that $p_i(\mathbf{m})$ and $p_{-i}(\mathbf{m})$ are observed by the analyst. This assumption is innocuous since these probabilities can be consistently estimated using choice data. In practice, this approach is applicable to datasets for which only a single choice is observed from each $(\mathbf{m}_i, \mathbf{m}_{-i})$ pair (as in our experiment).[11] For notation, we use pure letters (e.g., $\mathbf{m}_i$) to denote random variables and add superscripts to the letters (e.g., $\mathbf{m}_i^1$) to denote their realizations.

3.1. Over-identification of the error distribution

Given Assumption 3(a), we can invert the QRE conditions in Equation (4).[12] This inversion expresses the difference of expected utilities for player i, which is linear in player $-i$'s choice probability:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i})] &= EU_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i})] - EU_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i})] \\ &= \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}) \; \forall i. \end{aligned} \tag{5}$$

To derive the second line of Equation (5), $EU_i(\cdot)$ needs to be replaced with its definition in Equation (1) and we define $\tilde{\pi}_i(\mathbf{m}_i, a_{-i}) = u_i[m_i(a_i = 0, a_{-i})] - u_i[m_i(a_i = 1, a_{-i})]$; where $\tilde{\pi}_i(\mathbf{m}_i, a_{-i})$ represents the difference of the utilities between player i's two actions given the other player's choice a_{-i}.

Equation (5) contains all model restrictions that are imposed on player i's behavior and, importantly, implies that the error distribution is over-identified. To see why, first define $\mathcal{P}_i(\mathbf{m}_i^1) \subset [0, 1]$ as the image of the choice probability function $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$ when the analyst fixes $\mathbf{m}_i$ at $\mathbf{m}_i^1$ but varies $\mathbf{m}_{-i}$ over its support. Similarly, let $\mathcal{P}_i \subset [0, 1]$ denote the image of $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$ when the analyst varies both $\mathbf{m}_i$ and $\mathbf{m}_{-i}$. Given Assumptions 2 and 5, these images are either an interval or a union of finite numbers of disjoint intervals. Proposition 1 describes the over-identification of $F_i(\cdot)$.

[10] Notably, we do not require the analyst to know these switching points. For instance, consider the coordination game in Table 2. It is reasonable to expect that players may choose the equilibrium with a low probability of the safe action (i.e., action 0) when the payoff for the safe action—as represented by m_i—is relatively low. As m_i increases, players may switch to the equilibrium with a higher probability of action 0. Assumption 5 includes this reasonable equilibrium selection mechanism as a special case.

[11] In more detail, suppose that the analyst has a dataset of a series of games with different $(\mathbf{m}_i, \mathbf{m}_{-i})$. If $p_i(\mathbf{m})$ is continuous $\forall \mathbf{m} \in \mathcal{M}$, the analyst could use the Nadaraya-Waston estimator or the method of sieves to consistently estimate $p_i(\mathbf{m})$. If this choice probability function contains finite discontinuous points, consistent estimation could be also attained using the methods developed by Müller (1992) and Delgado and Hidalgo (2000).

[12] In games where players have more than two actions, we utilize the bijective results from discrete choice literature (Hotz and Miller, 1993; Sørensen and Fosgerau, 2022) to establish the inversion of QRE conditions. See the appendix for more details.

Proposition 1. *Suppose that Assumptions 1 to 5 and the QRE restrictions hold. Suppose further that the analyst fixes $\mathbf{m}_i$ at an arbitrary value $\mathbf{m}_i^1$ and only considers the variation of $\mathbf{m}_{-i}$. If there exist two distinct values of probabilities, denoted as $p^1, p^2 \in \mathcal{P}_i(\mathbf{m}_i^1)$, such that the values of $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ are known by the analyst, then the quantile function $F_i^{-1}(p)$ is point identified $\forall p \in \mathcal{P}_i(\mathbf{m}_i^1)$.*

Proof. See the appendix. □

Proposition 1 identifies the quantile function $F_i^{-1}(p)$ and consequently the distribution function $F_i(\tilde{\varepsilon})$ due to the inverse relationship between the two functions. Moreover, by applying Proposition 1 for each $\mathbf{m}_i \in \mathcal{M}_i$, the analyst can identify $F_i^{-1}(p)$ over its entire support, i.e., $\forall p \in \mathcal{P}_i$.

The over-identification result arises from a combination of the restrictions imposed by QRE. Specifically, player i's decision rule (Equation (2)) can be interpreted as a discrete choice model where each action a_i has a *deterministic component* $EU_i[\mathbf{m}_i, a_i, \mathbf{p}_{-i}(\mathbf{m})]$ and a *perturbed error* $\varepsilon_i(a_i)$. In single agent models without uncertainty, Norets and Takahashi (2013) show that without additional restrictions on the deterministic component, even partial identification of the error distribution is impossible. However, the expected-utility preference (which is assumed by QRE) places additional structure that allows the exogenous variation of $\mathbf{m}$ to point identify the error distribution. Specifically, player i's expected utility function $EU_i(\cdot)$ is linear in the other player's choice probability, despite the non-parametric utility function $u_i(\cdot)$.

It is this linearity combined with the QRE choice rule that leads to Equation (5), which is the key equation to establish our identification results. Consider fixing $\mathbf{m}_i$ at some value $\mathbf{m}_i^1$ while varying the other player's $\mathbf{m}_{-i}$. Since player i's monetary rewards are fixed, their utility difference remains unchanged at $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i})$. In contrast, the choice probability of the other player $p_{-i}(\mathbf{m})$ depends on both $\mathbf{m}_i$ and $\mathbf{m}_{-i}$ and will shift due to variation in $\mathbf{m}_{-i}$. Consequently, the right-hand side of Equation (5) can be viewed as a linear regression where $p_{-i}(\mathbf{m})$ is the independent variable with a coefficient $[\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=1)]$ and an intercept $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=1)$. This linear structure is akin to the semi-parametric binary choice models studied by Klein and Spady (1993) and Lewbel (2000), who established the non-parametric identification of the error distribution. Importantly, our framework remains fully non-parametric in both the utility and the distribution functions.

Proposition 1 requires the analyst to know *ex ante* the values of the quantile function at two probabilities p^1 and p^2. In the context of discrete choice games with field data, Liu et al. (2017) show that this requirement is innocuous and is equivalent to the standard location and scale normalizations required by discrete choice models (Train, 2009).[13] In Subsection 3.3, we show that in experimental settings, the values of $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ can be identified under weaker assumptions and are therefore not required to be known.

We refer to Proposition 1 as the over-identification result. Intuitively, with only two values of $\mathbf{m}_i$, Proposition 1 implies that $F_i^{-1}(\cdot)$ is over-identified. In the next subsection, we build on this intuition and construct an over-identification test for the hypothesis of QRE for all $\mathbf{m}_i$.

3.2. Test of QRE

Let $\hat{F}_i^{-1}(p|\mathbf{m}_i^1)$ be the quantile function that satisfies the QRE restrictions when the analyst fixes $\mathbf{m}_i$ at $\mathbf{m}_i^1$. Specifically, $\hat{F}_i^{-1}(p|\mathbf{m}_i^1)$ satisfies the following equation:

$$\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i})|\mathbf{m}_i^1] = \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}=1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}). \tag{6}$$

Proposition 1 shows that $\hat{F}_i^{-1}(p|\mathbf{m}_i)$ can be identified for each $\mathbf{m}_i$. This implies an over-identifying restriction such that $\hat{F}_i^{-1}(p|\mathbf{m}_i^1) = \hat{F}_i^{-1}(p|\mathbf{m}_i^2)\ \forall \mathbf{m}_i^1 \neq \mathbf{m}_i^2$. This restriction can be used to test the null hypothesis of QRE as stated in Proposition 2.

Proposition 2. *Suppose that Assumptions 1 to 5 hold and consider any two realizations of $\mathbf{m}_i$ denoted as $\mathbf{m}_i^1$ and $\mathbf{m}_i^2$ such that $\mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$ includes an interval. If there exist two distinct probabilities denoted as $p^1, p^2 \in \mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$ such that the values of $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ are known by the analyst, then the null hypothesis of QRE implies the following testable restriction:*

$$\hat{F}_i^{-1}(p|\mathbf{m}_i^1) = \hat{F}_i^{-1}(p|\mathbf{m}_i^2),\ \forall p \in \mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2). \tag{7}$$

Proof. A direct implication of Proposition 1. □

Equation (7) is a testable implication of player i's quantal response behavior, testing whether they quantal respond to other players' choice probabilities. To perform our test of QRE, we examine whether Equation (7) jointly holds for every player, therefore we employ the equilibrium correspondence approach as described in Sections 4 and 5.

Proposition 2 extends and generalizes previous tests of QRE in the existing literature. In particular, Xie (2022) also derives a non-parametric testable implication of BNE and equivalently of QRE. His result can be summarized by the following lemma:

[13] For instance, let $p^1 = 1/2$, then setting $F_i^{-1}(p^1) = 0$ is equivalent to the location normalization that sets the median value of $\tilde{\varepsilon}_i$ to zero. Moreover, let p^2 be the cumulative probability at one standard deviation above the median (e.g., approximately 68% for the normal distribution and roughly 86% for the Logit specification), then setting $F_i^{-1}(p^2) = 1$ is equivalent to the scale normalization that sets the standard deviation of $\tilde{\varepsilon}_i$ to one.

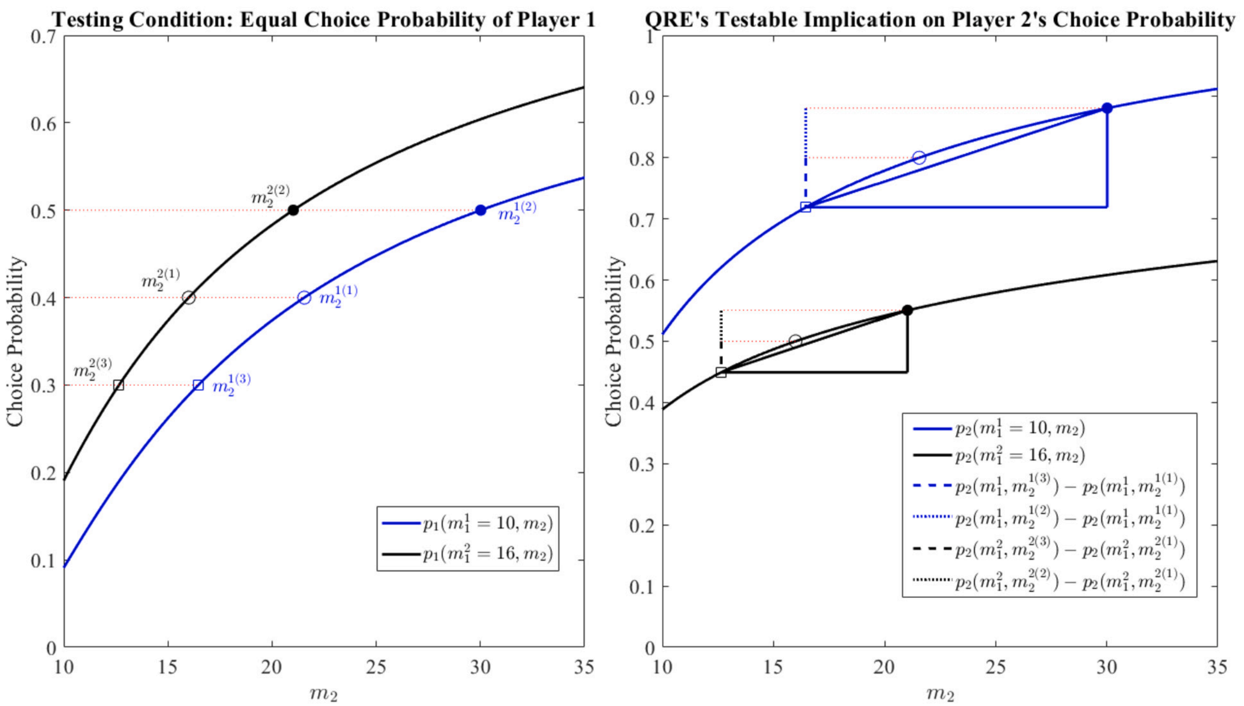

Fig. 1. Testable implication of QRE. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Lemma 1. (Xie 2022) *Suppose that Assumptions 1 to 5 hold. For any three pairs of realizations of* $(\mathbf{m}_i, \mathbf{m}_{-i})$ *that satisfy the condition* $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ *for* $l = 1, 2, 3$*; QRE implies the following testable restriction:*

$$\frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} = \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}, \tag{8}$$

when $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) \neq p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})$.

Our over-identification test illustrated in Proposition 2 advances Xie (2022)'s test in three important directions. First, as shown in Lemma 1, Xie (2022) requires an *equal choice probability* condition for at least three pairs of games.[14] This condition is difficult to construct in an actual dataset because it requires equating two estimated quantities and has yet to be implemented.[15] In contrast, our test in Proposition 2 circumvents the equal choice probability condition. It is, essentially, a by-product of the non-parametric estimation of the model primitives. Such estimation methods have been well developed in the econometrics literature. For our experimental analysis in Section 5, we apply non-parametric MLE by the method of sieves (Chen, 2007) to obtain the non-parametric estimate of the error distribution and implement the test.

Second, to implement Xie (2022)'s test, the analyst has to estimate and test the restrictions on each player's $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$, which is a multi-variate function. In contrast, our test requires the estimation of just two single-dimensional functions: $F_i(\tilde{\epsilon}_i)$ and $u_i(m)$. This dimension reduction improves estimation precision in finite samples, especially with a fully non-parametric specification. The benefit of dimension reduction is even more pronounced as the number of actions and / or players increases.[16]

Finally, our over-identification test includes not only all the testable implications derived by Xie (2022) but also many additional restrictions imposed by QRE. Consequently, our test has higher statistical power. While the detailed descriptions and proofs of these results are left to the appendix, we conduct a simulation of the matching pennies game in Table 1 to explain how our test nests Xie (2022). Fig. 1 shows simulated choice probabilities under QRE restrictions for various monetary payoffs.[17] The left panel plots Player 1's choice probability as a function of m_2. We hold m_1 constant at two distinct values: $m_1^1 = 10$ ($m_1^2 = 16$) is depicted using the blue

[14] Precisely, a pair consists of two games with different realizations of $(\mathbf{m}_i, \mathbf{m}_{-i})$ and player i's choice probability must remain constant within each pair.

[15] Specifically, it requires equating two functions of the estimates $p_i(\mathbf{m})$. This process incurs estimation error that substantially complicates the derivation of the finite sample property of the test.

[16] The dimension of $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$ grows in the order of $|\mathcal{A}|^N$, where $|\mathcal{A}|$ is the number of action profiles and N is the number of players. In contrast, the dimension of the error distribution only increases in the order of $|\mathcal{A}_i|$, which is merely the number of player i's actions. Additionally, the single dimension of the utility function remains constant.

[17] For illustrative purposes only, our simulation simply assumes players' utilities equal their monetary payoffs (i.e., $u_i(m) = m$) and employs the Logit choice probability.

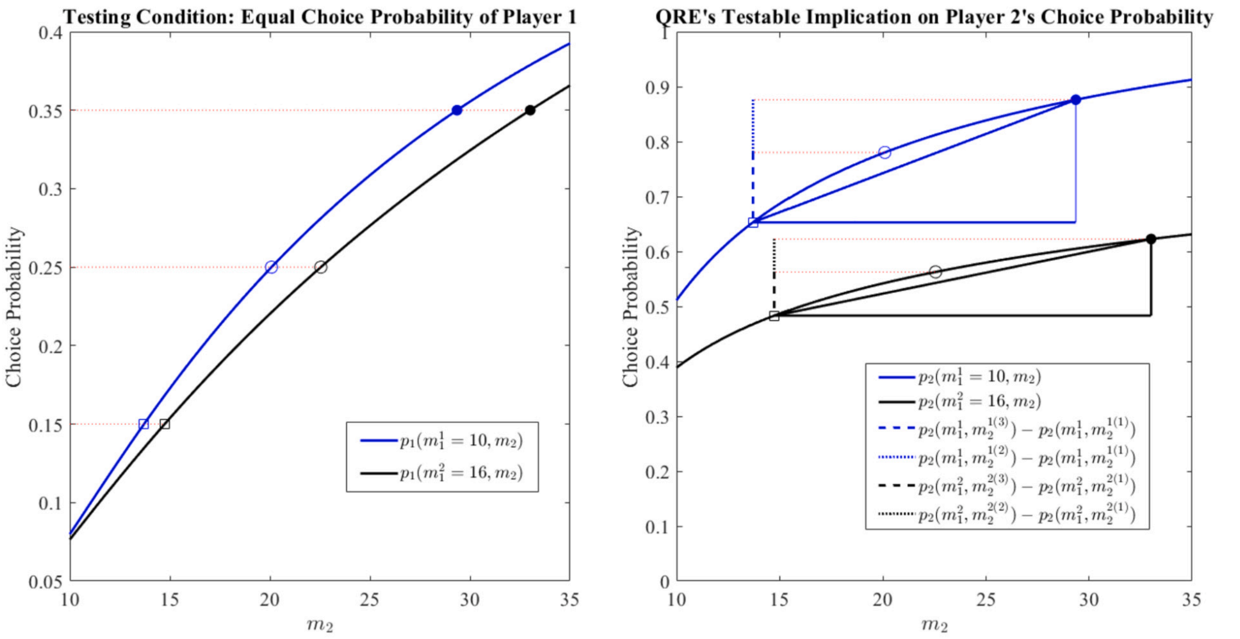

Fig. 2. Violation of the testable implication under non-QRE behaviors.

(black) curve. It also identifies three pairs of games that satisfy the equal choice probability condition for Player 1. Xie (2022) shows that under QRE, the ratio of the *change* in Player 2's choice probabilities across these pairs must be equal. These changes in $p_2(\cdot)$ are depicted by the colored dashed lines in the right panel. Based on Xie's test, the ratio of two dashed blue lines must equal the ratio of two dashed black lines.

As described in the appendix, our over-identification test implies that the ratio of both the *change* and the *level* of Player 2's choice probabilities across pairs must be equal. Consequently, our test can be visualized by at least two restrictions in the right panel. The first one, as described by Xie (2022) on the colored dashed lines, and the second visualized by a set of two similar triangles positioned on the blue and black curves respectively.

Fig. 2 shows a simulation of behavior that violates QRE. We fix Player 2's choice probabilities at their QRE levels, but we assume that Player 1 always underestimates Player 2's choice probability by 20 percentage points. These behaviors clearly violate QRE but they align with the testable implication proposed in Xie (2022). In contrast, such behaviors do not satisfy our additional testable implications, causing the two dissimilar triangles on the blue and black curves. This example graphically highlights the additional statistical power of our test. Notably, in the appendix, we show that there are many other testable implications of QRE in addition to Xie (2022). These implications are difficult to visualize and are not depicted in Figs. 1 and 2. They are, however, included in our over-identification test.

3.3. *Identification of normalizations and the utility function*

Most empirical applications of QRE assume that two actions will be chosen with equal probability if they share the same expected utility. In our framework, this assumption corresponds to assuming that the median of $\tilde{\varepsilon}_i$ is zero. In this subsection, we show that such an assumption is not necessary for identification. In fact, the analyst can identify the median value of $\tilde{\varepsilon}_i$.

As described in footnote 13, we shall refer to $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ as the median and the standard deviation of $\tilde{\varepsilon}_i$, respectively. To identify these two values, we introduce an innocuous scale normalization on the utility function, as described by the following assumption:

Assumption 6. For each player i, there exists a realization $\mathbf{m}_i = \mathbf{m}_i^1$ such that $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) = 1$.

Assumption 6 normalizes the scale of player i's utility function. It is innocuous given that $u_i(m)$ is strictly increasing.[18] Proposition 3 shows that once the analyst determines the scale of the utility function, the standard deviation of the error distribution is identified (i.e., the value of $F_i^{-1}(p^2)$).

[18] Since any affine transformation of utility $u_i(m) = c + \beta\hat{u}_i(m)$ for $\beta > 0$ represents the same preferences, the analyst needs to normalize the values of c and β. For any utility function $\hat{u}_i(m)$, Assumption 6 simply transforms $\hat{u}_i(m)$ to its equivalent form by setting $\beta = \frac{1}{\hat{u}_i[m_i^1(a_i=0,a_{-i}=1)]-\hat{u}_i[m_i^1(a_i=1,a_{-i}=1)]}$. Here, the analyst can ensure a positive denominator by re-labeling each player's action, in any game where $m_i(a_i = 0, a_{-i}) \neq m_i(a_i = 1, a_{-i})$ for some a_{-i}. We do not consider an experiment where $m_i(a_i = 0, a_{-i}) = m_i(a_i = 1, a_{-i})$ for all a_{-i}, as this type of games implies that player i has two *identical* actions. For instance, the utilities of $a_i = 0$ and $a_i = 1$ are the same, regardless of player $-i$'s choices.

Proposition 3. *Suppose that Assumptions 1 to 6 and the QRE restrictions hold. If there exists one probability denoted as $p^1 \in \mathcal{P}_i(\mathbf{m}_i^1)$ such that the value of $F_i^{-1}(p^1)$ is known by the analyst, then the quantile function $F_i^{-1}(p)$ is point identified $\forall p \in \mathcal{P}_i(\mathbf{m}_i^1)$.*

Due to Proposition 3, the only value that requires the analyst's prior information is the median of $\tilde{\varepsilon}_i$ (i.e., $F_i^{-1}(p^1)$). In the appendix, we show that in any experiment, as long as there is exogenous variation of the monetary payoffs for *at least two* action profiles for each player, this median value is identified.

When there is variation of monetary reward for only one action profile, the identification of $F_i^{-1}(p^1)$ requires special structures of monetary payoffs. Assumption 7 summarizes two such structures, which are satisfied in the matching pennies game presented in Table 1.

Assumption 7. Consider the following properties of monetary payoff matrices for each player i:

(a) There exists one realization of $\mathbf{m}_i$, denoted as $\mathbf{m}_i^1$, such that $m_i^1(a_i, a_{-i}) = m_i^1(1-a_i, 1-a_{-i})\ \forall a_i, a_{-i}$.

(b) There exist two realizations of $\mathbf{m}_i$, denoted as $\mathbf{m}_i^1$ and $\mathbf{m}_i^2$, such that $m_i^1(a_i, a_{-i}) = m_i^2(1-a_i, a'_{-i})\ \forall a_i$ and for some a_{-i} and a'_{-i}. Note that a_{-i} and a'_{-i} could be either distinct or identical actions of player $-i$.

Assumption 7(a) considers a design where player i's payoffs for both actions are reversed across the other player's choices within a game. It holds in Table 1 when $m_i = 16$. Assumption 7(b) follows similarly except it reverses payoffs across games with either varying or fixing actions of the other player. It holds in Table 1 when the analyst considers two values $m_i^1 = 16$ and $m_i^2 \neq 16$. Notably, this condition is also satisfied in the coordination game in Table 2 when the two values are $m_i^1 = 0$ and $m_i^2 = 15$.

When either condition in Assumption 7 holds, the median value of $\tilde{\varepsilon}_i$ can be identified, as established by the following proposition:

Proposition 4. *Suppose that Assumptions 1 to 6 and the QRE restrictions hold. Furthermore, suppose there exist two values $\mathbf{m}_i = \mathbf{m}_i^1, \mathbf{m}_i^2$ such that $\mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$ includes an interval. Moreover, either $\mathbf{m}_i^1$ satisfies Assumption 7(a) or $\mathbf{m}_i^1$ and $\mathbf{m}_i^2$ satisfy Assumption 7(b), then the quantile function $F_i^{-1}(p)$ is point identified $\forall p \in \mathcal{P}_i(\mathbf{m}_i^1) \cup \mathcal{P}_i(\mathbf{m}_i^2)$, without assuming that the value of $F_i^{-1}(p^1)$ is known ex ante by the analyst.*

Proof. See the appendix. □

The identification of the median value of $\tilde{\varepsilon}_i$, as established in Proposition 4, can be used to test the common assumption that the errors are i.i.d. across player i's actions. In particular, this i.i.d. restriction on $\varepsilon_i(a_i)$ implies that the difference of errors $\tilde{\varepsilon}_i$ is symmetrically distributed and has a median of zero.

Given that $F_i(\tilde{\varepsilon}_i)$ has been identified, we borrow the results from the existing literature to identify non-parametrically the utility function, as summarized by the following lemma.

Lemma 2. (Bajari et al. 2010) *Under Assumptions 1 to 7 and QRE restrictions, $F_i(\tilde{\varepsilon}_i)$ is point identified. Therefore, the empirical model reduces to the semi-parametric specification by Bajari et al. (2010), where the error distribution is known by the analyst and the utility function is non-parametric. Therefore the difference in utility $\tilde{\pi}_i(\mathbf{m}_i, a_{-i}) = u_i[m_i(a_i = 0, a_{-i})] - u_i[m_i(a_i = 1, a_{-i})]$ is point identified $\forall \mathbf{m}_i \in \mathcal{M}_i$, i and a_{-i}.*

Consider the first case where $\mathcal{M}_i = \times_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ and $\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ is an interval (e.g., the matching pennies game in Table 1). Given the standard location normalization such as $u_i(0) = 0$ or $u_i[\min\{\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})\}] = 0$, the identification of the difference in utilities from Lemma 2 directly identifies the utility function $u_i(m)$ non-parametrically $\forall m \in \cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$. Moreover, under Assumption 2, there is a second scenario where $\mathcal{M}_i \subset \times_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ and / or $\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ consists of finite number of disjoint intervals / singletons. In this case, whether the utility function $u_i(m)$ can be identified hinges on the structure of $\mathcal{M}_i$. Nonetheless, the difference of utilities remains identifiable according to Lemma 2. Crucially, the combination of the identified utility difference and the identified distribution function is sufficient to determine the QRE choice probabilities.

4. Monte Carlo experiment

We now describe our estimation and testing procedures and examine their finite sample performance in a Monte Carlo exercise using the matching pennies game depicted in Table 1. We evaluate the test in two different scenarios: one where data is generated in a QRE, and another in which QRE is not satisfied. Moreover, we design the exercise to closely align with the actual experiment that will be discussed in Section 5. As such, the Monte Carlo results can be used to evaluate the reliability of the empirical findings from our experiment.

4.1. Design of the Monte Carlo experiment

In our experiment, each participant makes a choice in 200 rounds. For our Monte Carlo exercise, in each of $S = 1000$ simulations we generate a dataset with T trials where $T \in \{200, 2000\}$. Therefore, $T = 200$ and $T = 2000$ can be viewed as representing situations where the analyst recruits one or ten participants per player role, respectively.[19] In addition, Assumption 2 requires $m_i(\mathbf{a})$ to be continuously distributed for at least one $\mathbf{a}$. This continuity condition is necessary to identify the continuous error distribution. In practice, the experimental design can approximate the continuity condition by a discrete distribution with a sufficiently small step size. In this exercise, we independently and uniformly draw m_1 and m_2 from the discrete set $\mathcal{M} = \{10, 12, \dots, 46, 48\}$ with step size of 2.[20]

In this Monte Carlo exercise, we consider three data generating processes. The first process assumes that data is generated consistently with QRE. In this scenario, the rejection rate of our proposed test should match the pre-specified significance level. Moreover, our procedure should obtain the estimates of utility and error distributions that are close to their true values.

The second and third processes generate data that is inconsistent with QRE. These scenarios illustrate whether our test has the power to reject a false hypothesis and achieve a small type-2 error. We therefore consider a modification of the Level-k model to generate non-QRE behavior (Nagel, 1995; Stahl and Wilson, 1994, 1995; Halevy et al., 2023). Specifically, the level-0 type randomly selects each action with equal probability. For any $k > 0$, the level-k type believes that their opponent is the level-$(k-1)$ type and *quantal responds* to such belief (i.e., a random error perturbed to the expected utility).

In the second process, we consider the symmetric level-k case where both players are of the same type and generate data for $k \in \{1, 2, 3\}$. Therefore, the quantal response function in Equation (3) does not hold for either player. The third process studies an asymmetric level-k setting where Player 1 reasons one level beyond Player 2. Therefore, the quantal response function in Equation (3) holds for Player 1 but not for Player 2. This scenario illustrates the performance of our test at either the player or the participant level. The test should frequently reject quantal response behavior for Player 2. By contrast, the rejection rate of the same hypothesis for Player 1 should be low and close to the desired type-1 error rate. We consider two different levels ($k = 2, 3$) for Player 1.

4.1.1. *Utility function and error distributions*

For convenience and comparability, we normalize the utility of the lowest possible ($m = 8$) and highest possible ($m = 48$) monetary rewards to zero and one, respectively, via the transformation $\tilde{m} = \frac{m-8}{48-8}$ so that $\tilde{m} \in [0, 1]$. In line with most experimental studies, we consider a CRRA utility function over the transformed monetary payoff $\tilde{m}$:

$$u_i(\tilde{m}) = \tilde{m}^{\nu}. \tag{9}$$

The utility curvature parameter ν is set to 0.6 to be consistent with the estimate obtained by Goeree et al. (2003) in their matching pennies game using the QRE framework.

We consider two candidates for the error distributions, both of which differ from the common specification of Logit or Probit. This allows us to investigate the consequences of imposing common distributional assumptions that are potentially mis-specified. In particular, we consider both a symmetric and an asymmetric error distribution:

$$\begin{aligned}
&\textit{Symmetric:}\ \ \tilde{\varepsilon}_i \sim 0.5 \cdot \mathit{Logistic}(0, 7.5 \cdot 3) + 0.5 \cdot \mathit{Logistic}\left(0, 7.5 \cdot \sqrt{\frac{9}{35}}\right), \\
&\textit{Asymmetric:}\ \ \tilde{\varepsilon}_i \sim 0.5 \cdot \mathit{Logistic}\left(-0.2, 7.5 \cdot \sqrt{\frac{4}{15}}\right) + 0.5 \cdot \mathit{Logistic}(0.2, 7.5 \cdot 2),
\end{aligned} \tag{10}$$

where $\mathit{Logistic}(\mu, \lambda) = \frac{\exp[\lambda(\varepsilon_i - \mu)]}{1+\exp[\lambda(\varepsilon_i - \mu)]}$ is the C.D.F. of the logistic distribution, with a mean of μ and a sensitivity parameter λ.

Equation (10) draws $\tilde{\varepsilon}_i$ from a mixture of logistic distributions, in line with the heterogeneity results described by Golman (2011). The symmetric case represents a population of two types of individuals, where one type makes smaller errors and consequently has a higher sensitivity parameter ($\lambda = 7.5 \cdot 3$) than the other $\left(\lambda = 7.5 \cdot \sqrt{\frac{9}{35}}\right)$. In the asymmetric distribution, alongside the heterogeneity in λ, individuals also make systematic errors. One type systematically under-values the expected utility of $a_i = 0$ by 0.2 while the second type over-values it by the same amount. Since the over-valuing type also has a higher λ, the population level $\tilde{\varepsilon}_i$ is asymmetrically distributed with a higher density in the positive region. Fig. 3 plots the P.D.F. of the symmetric and asymmetric distributions, alongside a comparison with the Logit specification.

We set the scale of $\tilde{\varepsilon}_i$ based on the empirical estimates in Section 5. In particular, the variances of $\tilde{\varepsilon}_i$ for both symmetric and asymmetric cases are set to $Var(\tilde{\varepsilon}_i) = \frac{\pi^2}{3 \times (7.5^2)} \approx 0.0585$. This value of variance corresponds to $\lambda = \sqrt{\frac{\pi^2}{3Var(\tilde{\varepsilon}_i)}} = 7.5$ if $\tilde{\varepsilon}_i$ were logistically distributed. Notably, this closely matches the empirical estimate of our actual experiment in Section 5 (i.e., $\hat{\lambda} = 7.505$) and is consistent with the estimate reported in Goeree et al. (2003), that is, $\hat{\lambda} = 6.67$, for a different matching pennies game but under the same normalization of the utility function.

[19] In the actual experiment, we recruited 50 participants per role, thus corresponding to $T = 10,000$. It is computationally challenging to run a Monte Carlo with this sample size as it requires the estimation to be repeated for 1,000 simulations in total. Moreover, our estimator and test achieve the desired performance when $T = 2000$.

[20] This approximation by discretization is akin to the method that solves a dynamic problem. Ideally, the step size should decrease as the sample size increases so that $\mathcal{M}$ is dense in the continuous interval $[10, 48]$ as $T \to \infty$. However, this simulation maintains a fixed step size for two reasons. First, the step size of 2 aligns with our experiment in Section 5 and the Monte Carlo exercise aims to evaluate the performance of both our estimator and our test using this step size. Second, shrinking the step size in larger samples substantially increases the computational burden.

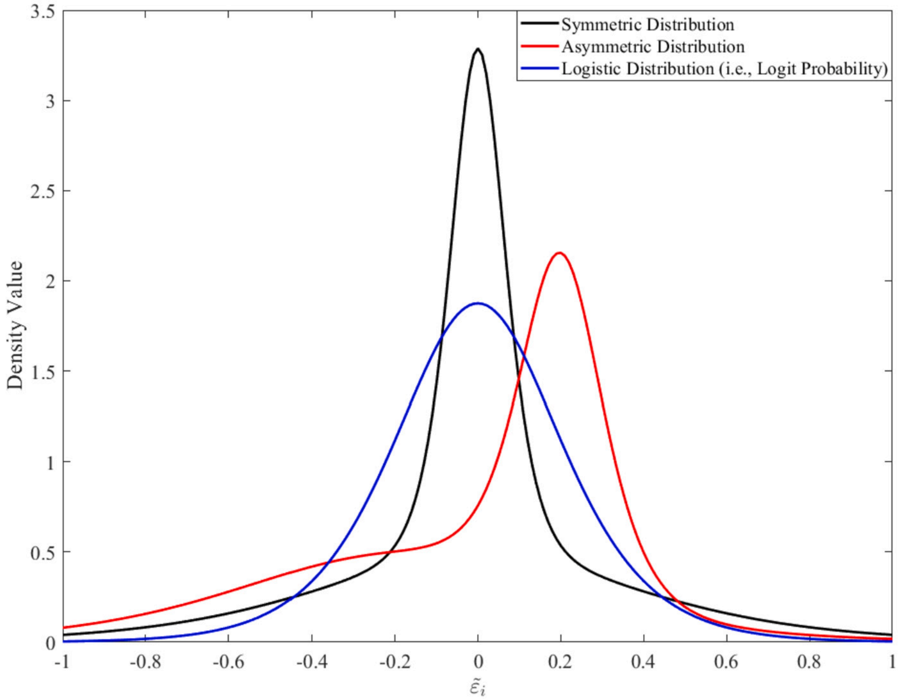

Fig. 3. Probability density functions of the random perturbation $\tilde{\varepsilon}_i$.

The scale of $\tilde{\varepsilon}_i$ plays an important role in shaping players' choice probabilities under both QRE and Level-k. In particular, the key convergence properties in Logit QRE, as derived in McKelvey and Palfrey (1995), also hold for the general distribution function $F_i(\tilde{\varepsilon}_i)$ in our framework.[21] In the appendix, we prove these convergence properties and provide a detailed analysis of comparative statics.

4.2. Estimation, testing procedures, and results

4.2.1. Estimation

We exploit the method of sieves to perform a non-parametric maximum likelihood estimation. As reviewed in Chen (2007), this method replaces a non-parametric function by a less complex function with finite-dimensional parameters. The dimension of these parameters increases as the sample size increases so that a less complex function can asymptotically approximate the original non-parametric function arbitrarily well.

The utility function is approximated using a Bernstein polynomial of order L_u,

$$u(\tilde{m}) = \sum_{l=0}^{L_u} \theta_l^u \cdot B_{l,L_u}(\tilde{m}) \tag{11}$$

with the l^{th} basis function denoted as $B_{l,L_u}(\tilde{m}) = \binom{L_u}{l} \cdot \tilde{m}^l \cdot (1-\tilde{m})^{L_u-l}$. As $L_u \to \infty$, the Bernstein polynomial will converge uniformly to the continuous function $u(\tilde{m})$ with $\theta_l^u = u(\frac{l}{L_u})$.[22] We set $L_u = 3$ (4) when $T = 200$ (2000).

Further, we approximate the distribution function using a mixture of normal distributions:

$$F_i(\tilde{\varepsilon}_i) \approx \sum_{l=1}^{L_F} \theta_l^{Pr} \cdot \Phi\Big(\frac{\tilde{\varepsilon}_i - \theta_l^{\mu}}{\theta_l^{\sigma}}\Big), \tag{12}$$

where $\Phi(\cdot)$ is the C.D.F. of the standard normal distribution. Each distribution indexed by l has a mean of θ_l^{μ} and a standard deviation of θ_l^{σ}. As the number of mixing distributions $L_F \to \infty$, the mixture of these distributions can effectively approximate any continuous distribution with high accuracy (Chen, 2007). In our estimation, we find that $L_F = 2$ performs well.

[21] For some intuition, as $Var(\tilde{\varepsilon}_i)$ decreases, player i tends to choose the action with a higher expected utility more frequently. When $Var(\tilde{\varepsilon}_i) \to 0$, player i deterministically selects the action that maximizes the expected utility. Conversely, as $Var(\tilde{\varepsilon}_i) \to \infty$, player i randomizes each action with equal probability.

[22] In a large sample with a sufficiently high order, this property could be exploited to impose regular restrictions on the utility function, such as strict monotonicity and concavity, which can improve the performance of the estimator (Compiani, 2022). We, however, do not utilize this property in our analysis due to a limited sample size and a relatively low order.

Let $\theta = (\theta^{u\prime}, \theta^{Pr\prime}, \theta^{\mu\prime}, \theta^{\sigma\prime})'$ denote the vector of unknown parameters in both the utility function and the distribution function. We estimate these parameters using the *equilibrium correspondence approach* (Goeree et al., 2020). In particular, given θ, we obtain the approximated utility and distribution functions and then solve for each player's QRE choice probability, denoted as $p_i^{QRE}(\tilde{\mathbf{m}}|\theta)$, using Equation (4). In the matching pennies game represented in Table 1, there exists a unique QRE for all $\tilde{\mathbf{m}}$. In applications with multiple QRE, the analyst has to select an equilibrium selection mechanism. Given $p_i^{QRE}(\tilde{\mathbf{m}}|\theta)$, the unknown parameters θ are estimated by maximizing the following log-likelihood function:

$$LL^{QRE} = \max_{\theta} \sum_{i=1}^{2} \sum_{t=1}^{T} \left\{ \mathbb{1}(a_{i,t} = 0) \cdot \log[p_i^{QRE}(\tilde{\mathbf{m}}_t|\theta)] + \mathbb{1}(a_{i,t} = 1) \cdot \log[1 - p_i^{QRE}(\tilde{\mathbf{m}}_t|\theta)] \right\}. \tag{13}$$

4.2.2. Testing

We exploit the over-identification result in Proposition 2 to test QRE. First, given $p_i^{QRE}(\tilde{\mathbf{m}}|\theta)$, we obtain the difference in expected utilities for player i under the QRE restriction, which we denote as $\widetilde{EU}_i(\tilde{\mathbf{m}}_i, p_{-i}^{QRE}(\tilde{\mathbf{m}}_i|\theta))$. To test QRE, we consider a general model that explicitly allows each player i to exhibit non-QRE behavior with the following choice probability:

$$p_i^{Non-QRE}(\tilde{\mathbf{m}}|\theta, \gamma_i) = \begin{cases} F_i[\widetilde{EU}_i(\tilde{\mathbf{m}}_i, p_{-i}^{QRE}(\tilde{\mathbf{m}}|\theta))] = p_i^{QRE}(\tilde{\mathbf{m}}|\theta) & \text{if } \tilde{\mathbf{m}} \notin \tilde{\mathcal{M}} \\ F_i[\widetilde{EU}_i(\tilde{\mathbf{m}}_i, p_{-i}^{QRE}(\tilde{\mathbf{m}}|\theta)) + \gamma_i(\tilde{\mathbf{m}})] & \text{if } \tilde{\mathbf{m}} \in \tilde{\mathcal{M}} \end{cases}. \tag{14}$$

$\tilde{\mathcal{M}}$ is a subset of the support for $\tilde{\mathbf{m}} = (\tilde{m}_1, \tilde{m}_2)'$. When $\tilde{\mathbf{m}} \notin \tilde{\mathcal{M}}$, Equation (14) states that player i behaves according to QRE. When $\tilde{\mathbf{m}} \in \tilde{\mathcal{M}}$, Equation (14) permits non-QRE behavior. The *bias term* $\gamma_i(\tilde{\mathbf{m}})$ captures the degree of player i's departure from QRE in the metric of expected utility.

The specification of Equation (14) follows directly from Proposition 2. In particular, $\gamma_i(\tilde{\mathbf{m}})$ can be alternatively interpreted as the difference between the quantile function identified by observations in $\tilde{\mathcal{M}}$ as opposed to the one obtained by observations not in $\tilde{\mathcal{M}}$. As in Proposition 2, QRE implies that $\gamma_i(\tilde{\mathbf{m}}) = 0$ and we test this restriction.

To ease the estimation burden, we consider a linear specification of $\gamma_i(\tilde{\mathbf{m}})$.

$$\gamma_i(\tilde{\mathbf{m}}) = \gamma_{i,0} + \gamma_{i,1}\tilde{m}_i + \gamma_{i,2}\tilde{m}_{-i}. \tag{15}$$

We set $\tilde{\mathcal{M}} = [0.2, 0.85]^2$, ensuring that approximately 50% of the observations fall within $\tilde{\mathcal{M}}$ and the rest fall outside of $\tilde{\mathcal{M}}$. The model and bias parameters are estimated by MLE:

$$LL^{Non-QRE} = \max_{\theta,\gamma} \sum_{i=1}^{2} \sum_{t=1}^{T} \left\{ \mathbb{1}(a_{i,t} = 0) \cdot \log[p_i^{Non-QRE}(\tilde{\mathbf{m}}_t|\theta, \gamma_i)] + \mathbb{1}(a_{i,t} = 1) \cdot \log[1 - p_i^{Non-QRE}(\tilde{\mathbf{m}}_t|\theta, \gamma_i)] \right\}. \tag{16}$$

The test of QRE is equivalent to testing whether $\gamma = (\gamma_1', \gamma_2')' = 0$. The latter can simply be performed by the standard likelihood ratio test with the test statistic $2(LL^{Non-QRE} - LL^{QRE})$. Under the QRE hypothesis, this statistic follows an asymptotic Chi-squared distribution with the degree of freedom given by the dimension of γ (i.e., 6).

Test of quantal response behavior for each player Proposition 2 can also be exploited to test the hypothesis that player i quantal responds to player $-i$'s choice probability. This test focuses on the choice of each individual player rather than their joint behavior. To perform such a test, we consider the *empirical payoff approach* (Goeree et al., 2020) and first non-parametrically estimate each player's choice probability in a reduced form:

$$\hat{p}_i(\tilde{\mathbf{m}}) = \frac{\exp[\sum_{l=0}^{L}(\sum_{h=0}^{l} \hat{\rho}_{i,l,h} \cdot \tilde{m}_1^h \cdot \tilde{m}_2^{l-h})]}{1 + \exp[\sum_{l=0}^{L}(\sum_{h=0}^{l} \hat{\rho}_{i,l,h} \cdot \tilde{m}_1^h \cdot \tilde{m}_2^{l-h})]}. \tag{17}$$

Equation (17) considers a Logit probability with a non-parametric specification for the index value. This non-parametric index is approximated by a high order polynomial. In the estimation, we find that an order of $L = 3$ performs well.

We reconsider Equation (14) but now we replace $p_{-i}^{QRE}(\tilde{\mathbf{m}}|\theta)$ with the other player's empirical choice probability $\hat{p}_{-i}(\tilde{\mathbf{m}})$. It leads to the following choice probability for player i:

$$p_i^{Non-QR}(\tilde{\mathbf{m}}|\theta_i, \gamma_i) = \begin{cases} F_i[\widetilde{EU}_i(\tilde{\mathbf{m}}_i, \hat{p}_{-i}(\tilde{\mathbf{m}}))] & \text{if } \tilde{\mathbf{m}} \notin \tilde{\mathcal{M}} \\ F_i[\widetilde{EU}_i(\tilde{\mathbf{m}}_i, \hat{p}_{-i}(\tilde{\mathbf{m}})) + \gamma_i(\tilde{\mathbf{m}})] & \text{if } \tilde{\mathbf{m}} \in \tilde{\mathcal{M}} \end{cases}. \tag{18}$$

Equation (18) explicitly permits non-quantal response behavior and the bias parameters γ_i can be consistently estimated by MLE:

$$LL_i^{Non-QR} = \max_{\theta_i,\gamma_i} \sum_{t=1}^{T} \left\{ \mathbb{1}(a_{i,t} = 0) \cdot \log[p_i^{Non-QR}(\tilde{\mathbf{m}}_t|\theta_i, \gamma_i)] + \mathbb{1}(a_{i,t} = 1) \cdot \log[1 - p_i^{Non-QR}(\tilde{\mathbf{m}}_t|\theta_i, \gamma_i)] \right\}. \tag{19}$$

Therefore, the null hypothesis of player i's quantal response behavior can be assessed by testing whether $\gamma_i = 0$, which is performed by the standard likelihood ratio test.

Similar to the above test of quantal response behavior at the participant level, the empirical payoff approach can also be exploited to estimate model primitives and test the QRE hypothesis. This approach is computationally efficient since it first estimates the

Table 3
Rejection rates of the over-identification test of QRE (QRE data).

	Symmetric Distribution			Asymmetric Distribution		
Significance Level	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
$T = 200$						
Known Utility & Known Error	12.4%	6.4%	1.0%	13.0%	7.4%	2.1%
Unknown Utility & Unknown Error	**15.8%**	**8.7%**	**1.8%**	**23.9%**	**16.1%**	**6.1%**
Linear Utility & Known Error	48.8%	35.7%	16.8%	28.6%	17.9%	7.6%
Known Utility & Logit Error	14.2%	7.5%	1.5%	65.4%	52.4%	29.3%
$T = 2000$						
Known Utility & Known Error	11.1%	4.3%	0.3%	11.8%	6.1%	1.2%
Unknown Utility & Unknown Error	**12.6%**	**6.7%**	**1.4%**	**13.2%**	**7.7%**	**1.7%**
Linear Utility & Known Error	100.0%	100.0%	99.9%	97.2%	94.4%	84.9%
Known Utility & Logit Error	46.5%	33.2%	13.5%	100.0%	100.0%	100.0%

Notes: Rejection rates are calculated based on 1,000 Monte Carlo samples.

equilibrium choice probability and avoids having to solve QRE for each iteration of θ. Moreover, in applications with multiple QRE, it does not require the analyst to impose an equilibrium selection mechanism but instead estimates the actual equilibrium observed in data. However, the empirical payoff approach produces a first-step estimation error which leads to inefficient estimates of model primitives.[23] Moreover, since it requires the estimation of the multivariate choice probability function, the benefit due to dimension reduction—as described in Subsection 3.2—vanishes. For these reasons, we focus on the equilibrium correspondence approach to estimate model primitives and test QRE in this study.

4.2.3. Monte Carlo results

Data generated by QRE Table 3 presents the rejection rates of our test when the data is generated by QRE behavior, consequently representing the type-1 error. The rejection rates are calculated based on 1,000 Monte Carlo datasets. We compare the results for four specifications. The first assumes that the analyst knows the true utility and distribution functions (labeled as "Known Utility & Known Error"). It inserts these true functions into the estimation procedure and tests whether the vector $\gamma = 0$. Obviously, this model is infeasible in an actual dataset, but it serves as a natural benchmark for the comparison of other specifications.

Our framework that non-parametrically specifies both the utility and the distribution functions and tests QRE as in Subsection 4.2.2 is labeled as "**Unknown Utility & Unknown Error**." As shown in Table 3, when the sample size is moderate (i.e., $T = 2000$ or 10 participants per player role), the rejection rates align with the pre-specified significance levels, for both symmetric and asymmetric distributions. Consequently, our test achieves the desired type-1 error rate. When the sample size is small (i.e., $T = 200$ or 1 participant per player role), our test tends to over-reject QRE due to small sample bias, especially when the error distribution is asymmetric.[24]

The remaining two specifications illustrate the consequences when either the utility function or the distribution function is mis-specified. The third specification assumes that the analyst knows the true distribution function but mis-specifies the utility function as $u_i(\tilde{m}) = \tilde{m}$ (labeled as "Linear Utility & Known Error"). The fourth specification assumes that the analyst knows the utility function but mistakenly considers the Logit choice probability (labeled as "Known Utility & Logit Error"). As shown in Table 3, when either the utility or the error distribution is mis-specified, the QRE hypothesis is substantially over-rejected with a moderate sample size (i.e., $T = 2000$). In most scenarios, the rejection rates are close to 100%. This over-rejection issue is less of a concern in small samples (i.e., $T = 200$). However, the mis-specification of either model primitives still leads to a substantially higher rejection rate than our proposed method (i.e., "Unknown Utility & Unknown Error"), except for the case of "Known Utility & Logit Error" under the symmetric distribution. In this scenario, even when the Logit formula is mis-specified, it correctly imposes the symmetry condition. In a small sample, this correct shape restriction leads to a rejection ratio that is slightly lower than our test.

Fig. 4 plots the averages of the estimated utility functions and the distribution functions across 1,000 Monte Carlo samples with their corresponding 90% confidence intervals. It shows that with a moderate sample size, the model primitives can be reliably and non-parametrically estimated.

Data generated by non-QRE behavior: symmetric iterative reasoning We next generate data according to the standard Level-k model to assess the power of our test to reject an incorrect hypothesis; that is, the type-2 error. Table 4 presents the rejection rates when each player has the same level of sophistication in their reasoning (i.e., symmetric Level-k). Specifically, when $T = 2000$, the test obtains a rejection rate of almost 100% for any error distribution and any level of iterative reasoning. This suggests that the proposed testing

[23] To deal with the first-step estimation error in the test of quantal response behavior, we input the true value of $p_i(\mathbf{m})$ in the Monte Carlo exercise. In the actual experiment, we assume away the first-step error. This is because the first-step estimates choice probabilities at the population level, which are based on a substantially larger sample size than the second step that tests model primitives at the participant level.

[24] The over-rejection of QRE in small samples is akin to the well-known problem of over-fitting. In particular, the general choice probability structure in Equation (14) includes the bias parameter γ. When QRE holds, these parameters are unnecessary to explain players' behavior. However, if the sample size is small, these parameters would fit idiosyncratic sample noise. This over-fitting problem then translates to the over-rejection of QRE. Note that the benchmark specification ("Known Utility & Known Error") also exhibits a comparable over-rejection problem in small samples.

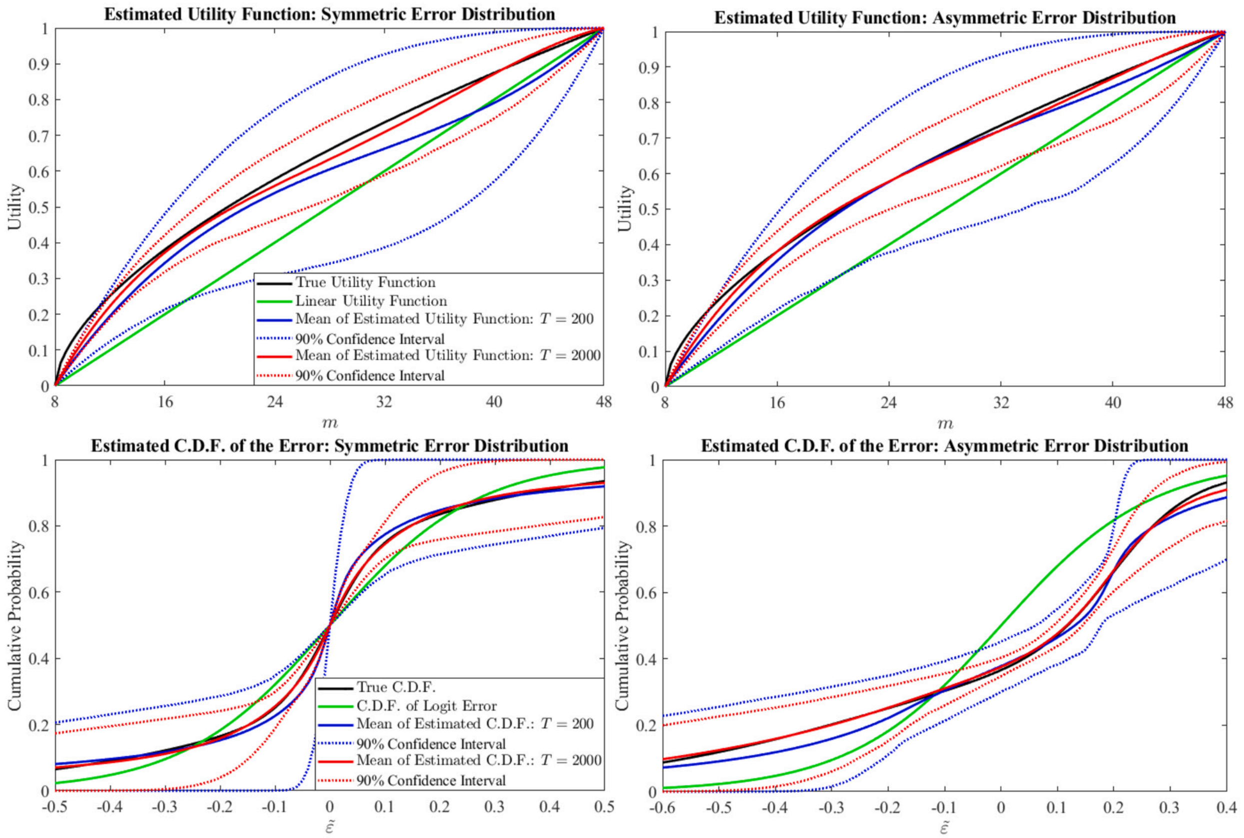

Fig. 4. Estimates of the utility function and the distribution function.

procedure possesses the power to reject an incorrect null hypothesis with a moderate sample size. When the sample size is small (i.e., $T = 200$), the test's performance crucially depends on the level of iterative reasoning. In cases where players are not sufficiently sophisticated (i.e., level-2 or below), our test exhibits lower rejection rates compared to the benchmark "Known Utility & Known Error," and these rates fall below 50%. Intuitively, the choice probabilities of low-level players exhibit weak dependence on $\tilde{\mathbf{m}}$. For instance, a level-1 player's decision is independent of the other player's monetary reward. In small samples, this limited dependence on $\tilde{\mathbf{m}}$ could lead to imprecise estimates and reduce the power of our test. In contrast, when players are more sophisticated (i.e., level-3), our test rejects the incorrect null hypothesis of QRE almost certainly, regardless of the shape of the error distribution.

Data generated by non-QRE behavior: asymmetric iterative reasoning Our final exercise considers players with heterogeneous levels of sophistication in their reasoning process (i.e., asymmetric Level-k) and studies the performance of the test for each individual player's quantal response behavior. Table 5 presents the test results.

In this exercise, Player 1 has the ability to perform an additional step of iterative reasoning compared to Player 2. Therefore, the quantal response function in Equation (3) holds for Player 1 and the rejection ratio for the test of this player's quantal response behavior should be close to the pre-specified significance level. As shown in Table 5, the rejection rates align with the desired type-1 error rate with a moderate sample size (i.e., $T = 2000$). In a small sample (i.e., $T = 200$), the test tends to over-reject the hypothesis of Player 1's quantal response behavior due to small sample bias.

Player 2, on the other hand, is characterized by a lower level in their iterative reasoning and does not quantal respond to Player 1's choice probability. Furthermore, players' joint behaviors violate QRE. Therefore, our test should frequently reject two incorrect null hypotheses: (i) quantal response behavior for Player 2, and (ii) QRE. As shown in Table 5, the rejection rates for these two tests are consistently close to 100% across different sample sizes and error distributions. These results demonstrate the high statistical power of our test.

5. Empirical application: an experimental study

Our empirical application focuses on the matching pennies game as presented in Table 1 and maintains the same structure as Goeree and Holt (2001). In a previous study of this game, Aguirregabiria and Xie (2021) do not reject quantal response behavior for the row player using the data from Goeree and Holt (2001). Moreover, in a generalized 3×3 matching pennies game, Melo et al.

Table 4
Rejection rates of the over-identification test of QRE (symmetric level-k data).

	Panel A: $T = 200$					
	Symmetric Distribution			Asymmetric Distribution		
Significance Level	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Level-1 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	89.2%	82.5%	62.1%
Unknown Utility & Unknown Error	62.0%	49.3%	25.9%	36.8%	26.0%	10.4%
Level-2 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	100.0%	99.9%	99.3%
Unknown Utility & Unknown Error	47.0%	35.5%	17.6%	80.5%	70.2%	47.4%
Level-3 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	100.0%	99.7%	98.7%
Unknown Utility & Unknown Error	100.0%	100.0%	100.0%	99.9%	99.6%	96.8%
	Panel B: $T = 2000$					
	Symmetric Distribution			Asymmetric Distribution		
Significance Level	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Level-1 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Unknown Utility & Unknown Error	100.0%	100.0%	99.9%	99.3%	98.3%	94.1%
Level-2 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Unknown Utility & Unknown Error	99.3%	98.6%	94.9%	100.0%	100.0%	100.0%
Level-3 Reasoning Behavior						
Known Utility & Known Error	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Unknown Utility & Unknown Error	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

Table 5
Rejection rates of the over-identification test of QRE (asymmetric level-k data).

	Panel A: $T = 200$					
	Symmetric Distribution			Asymmetric Distribution		
Significance Level	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Player 1 is Level-2 Reasoner & Player 2 is Level-1 Reasoner						
Test of QRE	97.8%	96.2%	88.5%	96.8%	91.9%	80.4%
Test of Quantal Response for Player 1	15.7%	8.8%	1.6%	16.6%	8.1%	1.6%
Test of Quantal Response for Player 2	99.9%	99.9%	99.9%	100.0%	99.9%	98.8%
Player 1 is Level-3 Reasoner & Player 2 is Level-2 Reasoner						
Test of QRE	100.0%	100.0%	100.0%	98.7%	96.8%	90.2%
Test of Quantal Response for Player 1	19.4%	11.1%	2.4%	19.9%	12.5%	3.0%
Test of Quantal Response for Player 2	100.0%	100.0%	100.0%	96.1%	93.9%	82.4%
	Panel B: $T = 2000$					
	Symmetric Distribution			Asymmetric Distribution		
Significance Level	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
Player 1 is Level-2 Reasoner & Player 2 is Level-1 Reasoner						
Test of QRE	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Test of Quantal Response for Player 1	12.5%	6.8%	1.2%	13.3%	6.9%	1.4%
Test of Quantal Response for Player 2	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Player 1 is Level-3 Reasoner & Player 2 is Level-2 Reasoner						
Test of QRE	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Test of Quantal Response for Player 1	12.6%	7.0%	2.2%	14.1%	8.6%	2.1%
Test of Quantal Response for Player 2	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

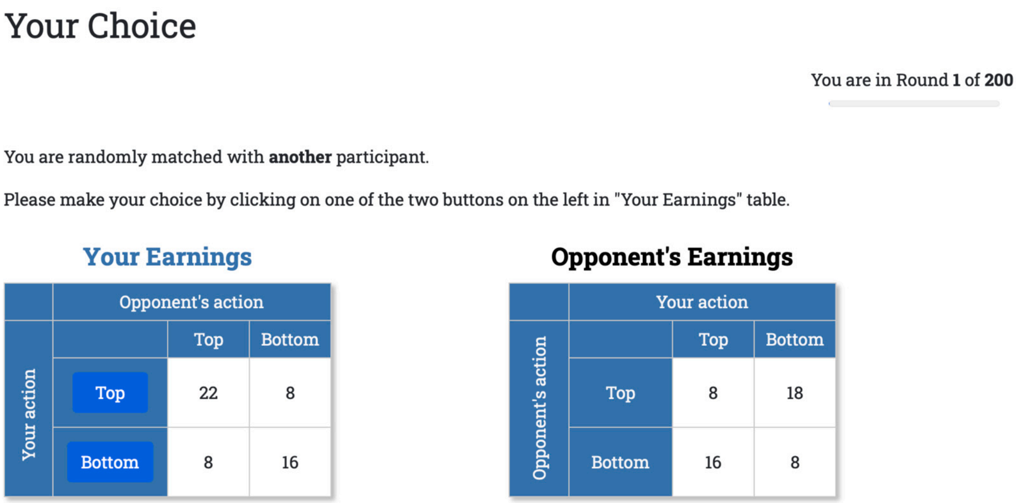

Fig. 5. Matching pennies game – experimental implementation.

(2019) do reject QRE at the population level, but cannot reject quantal response behavior at the participant level for more than 50% of participants.

5.1. *Experimental design*

Our design closely follows the Monte Carlo exercise in Section 4. In particular, we exogenously varied two variables, m_1 and m_2, that directly enter the utility function, one for each player. These variables were unique combinations drawn uniformly from a discrete set of 20 values, $\mathcal{M} = \{10, 12, 14, ..., 48\}$. We randomized the order of these combinations for a given experiment session. Each session was comprised of 20 participants who were allocated to two separate matching groups and player roles were assigned. Throughout the experiment, each participant maintained their player role and remained in their group. To ensure efficient data collection, each group played, in total, 200 matching pennies games with varying monetary payoffs and with random re-matching to mute potential order effects. Thus, in a given experiment session with 20 participants we collected data using $|\mathcal{M}|^2 = 20^2 = 400$ unique monetary payoff combinations.

Fig. 5 visualizes the experimental implementation of the bimatrix matching pennies game, where the variables m_1 and m_2 were exogenously varied and changed in each round (in this example, $m_1 = 22$ and $m_2 = 18$). To create a more natural and intuitive interface, we displayed one 2×2 matrix for each player separately as in Halevy et al. (2023). The first matrix represents player 1's monetary payoffs and the second matrix represents player 2's monetary rewards, respectively.

To improve participants' experience and to assist in selecting an action, we implemented a highlighting tool that uses yellow color. When a participant moves their mouse over a row in their matrix ("Your Earnings"), the action is highlighted in yellow color in both matrices: a row in their matrix, and a column in the opponent's matrix ("Opponent's Earnings"). By left clicking the mouse over a row it remains highlighted, and participants can un-highlight it by clicking their mouse again or clicking another row. Similarly, when participants move their mouse over a row that corresponds to an action of the opponent in "Opponent's Earnings," the row is highlighted in yellow and the corresponding column is highlighted in yellow in "Your Earnings." Clicking the mouse over the row keeps it highlighted, and clicking it again (or clicking another action) unhighlights it.[25]

We conducted the experiment with students enrolled at the University of Vienna in December 2022. In total, 100 participants were recruited from Vienna Center for Experimental Economics' (VCEE) pool using ORSEE (Greiner 2015). No participant was allowed to participate in more than one session.

After reading the instructions, participants had to correctly answer three comprehension questions before starting the first task. If participants made a mistake in answering a quiz question, they had to answer it correctly in order to move to the next question. The experiment was programmed in oTree (Chen et al. 2016). For each participant, we randomly selected one of the 200 matching pennies games that they had played, and rewarded them based on the earnings in this selected game. This design mutes potential hedging incentives. The average participant earned € 19.18 ≈ \$20.50, including a show-up payment of € 5, in a session that typically lasted around 70 minutes.

5.2. *Experimental data and results*

Table 6 reports the estimated coefficients from a reduced form Logit regression, where we regress player i's choice probability of action 0 on m_1 and m_2. As would be expected, an increase of m_i strictly increases the expected utility of $a_i = 0$ for player i, holding the

[25] The interactive experimental interface can be accessed anytime upon request. Example screenshots can be found in the appendix.

Table 6
Reduced form logit regression of player i's choice probability function.

	Player 1	Player 2
m_1	0.027*** (0.002)	−0.050*** (0.002)
m_2	0.028*** (0.002)	0.052*** (0.002)
Constant	−1.016*** (0.079)	−0.223*** (0.077)
Log-likelihood	-6339.57	-6184.61
Observations	10,000	

Notes: *, **, and *** represent significant at 10%, 5%, 1% significance levels, respectively.

Table 7
Chi-square statistic and p-value of the test of QRE (population level).

Linear Utility & Logit Error	$\chi^2 = 245.68$ $p < 0.0001$
Unknown Utility & Logit Error	$\chi^2 = 181.53$ $p < 0.0001$
Linear Utility & Unknown Error	$\chi^2 = 94.74$ $p < 0.0001$
Unknown Utility & Unknown Error	$\chi^2 = 93.81$ $p < 0.0001$

other player's choice probability constant. Consequently, the rise of m_i increases $p_i(m_i, m_{-i})$. This effect is known as the *own-payoff effect* and is a common feature in experimental studies of matching pennies games (Ochs, 1995; Goeree et al., 2003). This own-payoff effect is also salient and highly significant in our dataset. Moreover, if a player knows that the other player experiences an own-payoff effect, the structure of the matching pennies game implies that player 1's choice probability of action 0 increases in m_2 (while player 2's probability decreases in m_1). Table 6 shows that such *effect of other-payoff* is also sizable and statistically significant.

Our analysis starts with testing QRE and estimating model primitives under the condition of QRE at the population level. We then test the hypothesis of quantal response behavior for each participant in our experiment. The estimation and testing procedures follow the process described in Subsection 4.2.

Population level analysis of the heterogeneous QRE We allow each participant in the experiment to have a heterogeneous error distribution but assume that they share the same utility function. In this scenario, QRE at the population level can be described by a representative player whose error distribution is non-parametrically specified (Golman, 2011). Due to this interpretation, one can view all participants with the same player role as a single participant or player who makes $T = 50 \times 200 = 10,000$ decisions. Notably, this treatment of heterogeneous QRE that allows different error distributions nests the *heterogeneous Logit QRE* (Rogers et al., 2009; Golman, 2012) as a special case.

Table 7 presents the results of the test of QRE for four different specifications. The first one follows the standard procedure in the literature, assuming the utility is given by the monetary reward (i.e., risk-neutral participants with perfect perception) and the Logit choice probability. Accordingly, this specification is labeled as "Linear Utility & Logit Error." The second and third specifications relax one of the two restrictions and only allow one of the functions to be unknown to the analyst. They are referred to as "Unknown Utility & Logit Error" and "Linear Utility & Unknown Error." The last specification is the one proposed in this study—it allows both functions to be unobserved by the analyst (i.e., "Unknown Utility & Unknown Error"). In each specification, all unknown functions are non-parametrically specified.

As shown in Table 7, the null hypothesis of QRE is rejected in our data under all specifications. However, the results also deliver an important message: when fewer restrictions are imposed on the utility and the distribution functions, QRE becomes more difficult to reject. This is reflected in the decreasing statistic of the likelihood ratio test.

Notably, the test of QRE examines whether QRE *perfectly* matches the actual choice probability $p_i(\cdot)$. Therefore, the population-level rejection of QRE should not be mistakenly interpreted as contradicting the common finding in the literature that QRE generally fits the data well (Camerer, 2003; Crawford et al., 2013). In particular, the literature usually evaluates QRE based on whether its prediction is *sufficiently close* to the true choice probability. Since we non-parametrically identify the model primitives under QRE, we can also evaluate *how close* QRE is to the actual $p_i(\cdot)$ under different specifications. To perform such a comparison, we consider the following measure of normalized log-likelihood:

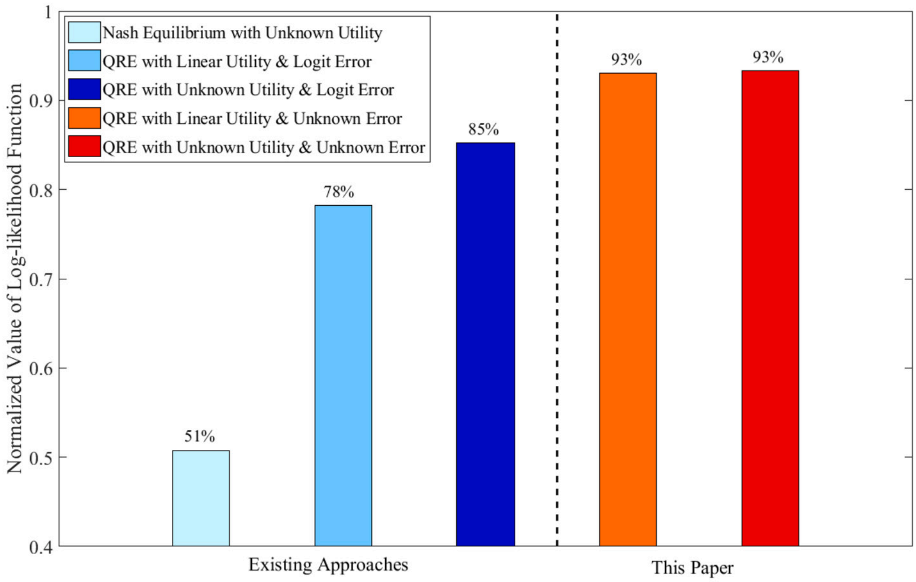

Fig. 6. Model fitness.

$$\textit{Normalized Log-Likelihood} = \frac{LL^{Model} - LL^{Random}}{LL^{Sample} - LL^{Random}}, \tag{20}$$

where LL^{Model} is the log-likelihood value for the corresponding model and LL^{Random} is evaluated when each action is assumed to be chosen with equal probability, representing the lower bound that any model should beat. LL^{Sample} is calculated using the smooth non-parametric reduced form estimates $\hat{p}_i(\mathbf{m})$, as shown in Equation (17). By construction, it represents the maximum value of log-likelihood that any model could reach. As sample size grows, our test will reject QRE as long as it does not achieve a perfect fit (of 100%). In contrast, our estimation procedure allows us to evaluate the closeness of QRE to the perfect fit.

Fig. 6 plots the normalized log-likelihood values for different model specifications. As previously noted, even the simplest QRE (i.e., with Logit error and linear utility) fits the data substantially better than NE *with a non-parametric utility function*. This is because NE only predicts an other-payoff effect, while QRE predicts both own-payoff and other-payoff effects (Table 6). Importantly, when the error distribution is non-parametrically specified, the model fit substantially increases beyond 90% (the right two bars on Fig. 6). These specifications can be viewed as a population level analysis of the heterogeneous QRE (Golman, 2011). As a benchmark, existing approaches which impose the distributional assumption perform worse. Finally, Fig. 7 plots the same measure for an out-of-sample procedure that estimates model primitives for 50% of participants and predicts on the remaining participants. Here again, the two specifications with non-parametric error distributions achieve the best of out-of-sample fit. Consequently, QRE explains much of the variation in participants' behavior in this game.

In Fig. 8, we present the non-parametric estimates of model primitives under QRE. We plot the estimated utility function with a 90% confidence interval (black dotted lines) and compare it with the linear utility assumption (blue line). Our estimates suggest that participants are risk neutral when the monetary payoff is low or moderate. Only when the reward is very high (i.e., above € 40 ≈ $43) do we find utility curvature consistent with risk aversion to be significant at the 10%-level.

Fig. 9 plots the estimated P.D.F. for the error distribution with a 95% confidence interval (black dotted lines) and compares it with the logistic distribution with the same variance (blue line). This illustrates the strong rejection of the Logit choice probability. Compared to the logistic distribution, participants tend to make errors of smaller magnitude. Moreover, the estimated distribution has a heavier tail, suggesting that participants also make larger errors with non-negligible probability. In contrast, there is a smaller probability of moderate mistakes.

While the estimated P.D.F. of $\tilde{\varepsilon}_i$ appears symmetric, it is not symmetric around 0. In particular, the estimate of $mean(\tilde{\varepsilon}_i)$ is 0.080 and is highly significant at the 1%-level, with a standard error of 0.020. Given the estimated utility function, this estimate of $mean(\tilde{\varepsilon}_i)$ suggests that participants tend to over-estimate the reward of the action presented at the top of the screen by around € 3 ≈ $3.20, which is approximately 6% of the maximum reward. Our non-parametric estimate is able to recover this sizable position effect, which is usually assumed to be absent in existing applications of QRE.

Test of quantal response behavior at the participant level Our second analysis allows both the utility function and the error distribution to be heterogeneous across participants. We aim to test whether each participant exhibits quantal response behavior with respect to the other player's actual choice probabilities. If this hypothesis holds for each participant, we could interpret the data as being consistent with QRE featuring heterogeneity in both the utility function and the distribution function.

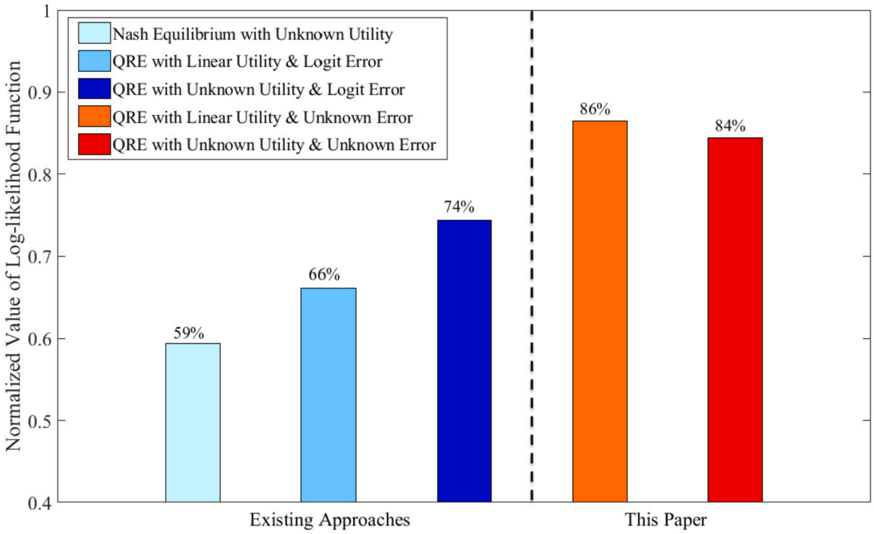

Fig. 7. Out of sample fitness.

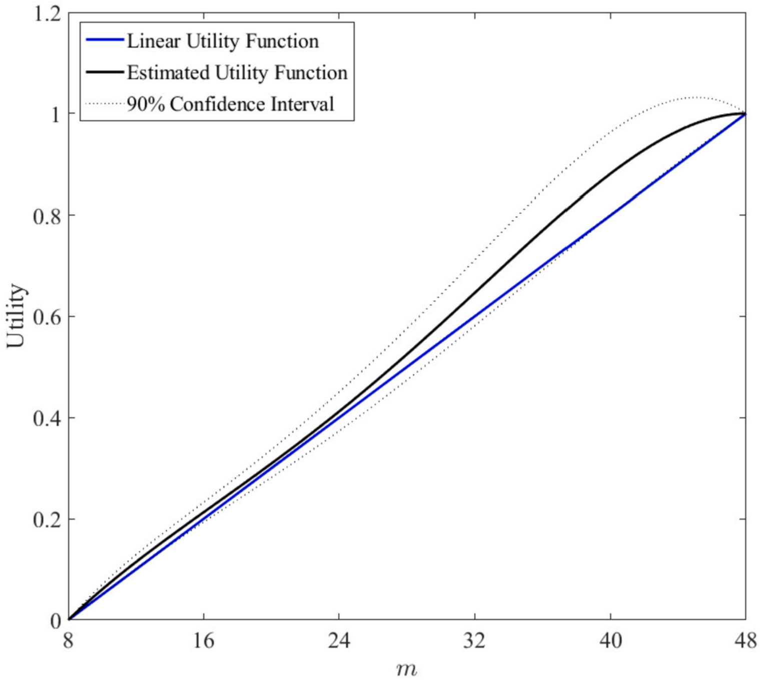

Fig. 8. Estimated utility function.

Fig. 10 presents the empirical C.D.F. for the p-value of the test statistic, with a vertical line that represents statistical significance at the 5%-level. Therefore, the intersection of the empirical C.D.F. and the vertical line shows the fraction of participants for whom quantal response behavior is rejected at the 5%-level.

Similar to the results at the population level, the test highlights a general trend: the fewer restrictions imposed on the utility and the error distribution, the more likely it is that QRE holds in the data. Under the assumption of a risk neutral utility function and a logistically distributed error, quantal response behavior is rejected for 70% of participants. When only one of the two model primitives is restricted, the null hypothesis is rejected for about 50% of participants. In contrast, with unknown and non-parametric specifications of both functions, quantal response behavior is rejected for only 30% of participants. Notably, this test at the participant level is conducted with a sample size of $T = 200$, which may introduce small sample bias as described in our Monte Carlo results (i.e., Table 5). Since this small sample bias tends to over-reject the quantal response hypothesis, rather than reducing the power to reject incorrect hypotheses, these results are more supportive of quantal response behavior than previous methods.

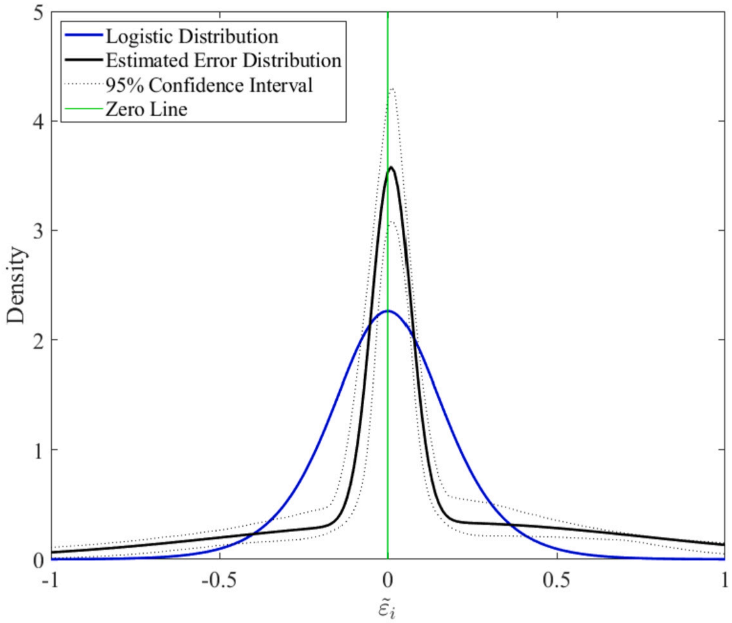

Fig. 9. Estimated P.D.F. for $\tilde{\varepsilon}_i$.

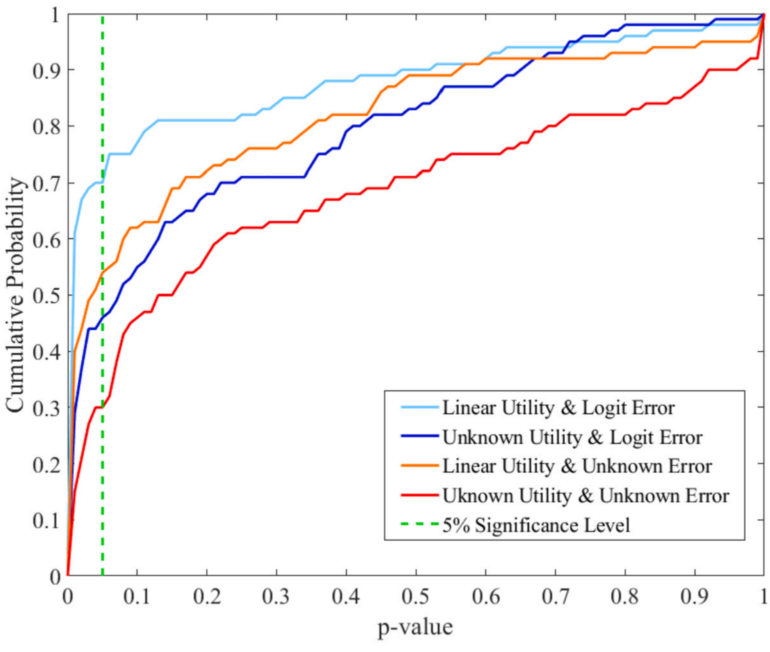

Fig. 10. Empirical C.D.F. of the p-value of the test of quantal response behaviors.

In summary, the quantal response hypothesis has a satisfactory statistical fit when allowing for sufficiently flexible and heterogeneous utility and error distributions. However, when strong assumptions in terms of the functional form or homogeneity are imposed, QRE is strongly rejected. These results emphasize the importance of a flexible and unknown specification of all model primitives. With this specification, the identification results and the testable implication derived in this paper are particularly useful.

6. Conclusion

This paper studies the falsifiability and identification of QRE when both the utility and the error distribution are non-parametric functions. Making use of cross-game variation, we first show that the error distribution and the utility function are non-parametrically over-identified. This over-identification result implies a straightforward testing procedure for QRE. The Monte Carlo experiment suggests that our test has sufficient power to reject a false hypothesis. Moreover, when QRE holds in the data, our estimation procedure can reliably recover both the utility and the distribution functions non-parametrically.

As shown by Golman (2011), the non-parametric error specification can be viewed as a population level fit of QRE with heterogeneous error distributions across participants. Previous studies have not exploited this interpretation because of a lack of identification results. This paper fills this gap by providing a means to fit heterogeneous QRE at the population level. In an experimental study of the matching pennies game, we find that QRE with a non-parametric error distribution fits the data substantially better than previous methods, both in-sample and out-of-sample. This suggests substantial heterogeneity in error distributions in our sample. Moreover, at the participant level, with a heterogeneous and non-parametric specification of the utility and the error distribution, the quantal response hypothesis cannot be rejected for a majority of participants. However, it is highly rejected with strong assumptions on functional form or homogeneity.

Our framework's weak assumptions on the monetary payoff structures enable an analyst to test QRE in a wide class of games, accommodating their various research objectives. For instance, while this paper focuses on the matching pennies game (Table 1), our method is equally applicable to other types of games such as coordination games (Table 2). An important feature of our approach is that it enables an analyst to test the validity of QRE both within and across game types. For instance, the analyst could design an experiment where the payoff structure $\mathcal{M}$ is a union of Tables 1 and 2, this design allows for testing whether QRE jointly holds in both matching pennies and coordination games.

Our results build on the invariance assumption that each player's error distribution remains unchanged across games. Consequently, we focus on games with fixed number of players and actions. When players have different action sets across games, the joint distribution of errors across actions will vary and our results do not apply. However, with some additional restrictions, it is possible to generalize our results. For instance, consider a series of 2×2 games and another series of 3×3 games, with the common restriction that errors are i.i.d. across actions (Goeree et al., 2020). Based on the results in this paper, the analyst could first non-parametrically estimate the utility function and the marginal error distribution using data from the 2×2 games. Under the i.i.d. restriction, these non-parametric estimates then determine the set of predicted choice probabilities under QRE for the 3×3 games.[26] This result could then be used to test QRE by testing whether the set of predicted probabilities contains the true choice probability. In a semi-parametric specification, Xie (2018) shows that the above variations in the action sets could provide extra information to test BNE, and equivalently QRE.

Declaration of competing interest

None.

Data availability

Data will be made available on request.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2024.07.004.

References

Aguirregabiria, Victor, Xie, Erhao, 2021. Identification of non-equilibrium beliefs in games of incomplete information using experimental data. J. Econ. Methods 10, 1–26.
Aguirregabiria, Victor, Jeon, Jihye, 2020. Firms' Beliefs and Learning: Models, Identification, and Empirical Evidence. Rev. Ind. Organ. 56, 203–235.
Aguirregabiria, Victor, Mira, Pedro, 2019. Identification of games of incomplete information with multiple equilibria and unobserved heterogeneity. Quant. Econ. 10 (4), 1659–1701.
Allen, Roy, Rehbeck, John, 2021. A generalization of quantal response equilibrium via perturbed utility. Games 12 (1).
Bajari, Patrick, Hong, Han, Krainer, John, Nekipelov, Denis, 2010. Estimating static models of strategic interactions. J. Bus. Econ. Stat. 28, 469–482.
Camerer, Colin F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press.
Camerer, Colin F., Ho, Teck-Hua, Chong, Juin-Kuan, 2004. A cognitive hierarchy model of games. Q. J. Econ. 119 (3), 861–898.
Chen, Daniel, Schonger, Martin, Wickens, Chris, 2016. oTree-an open-source platform for laboratory, online, and field experiments. J. Behav. Exp. Finance 9, 88–97.
Chen, Xiaohong, 2007. Large sample sieve estimation of semi-nonparametric models. In: Heckman, James J., Leamer, Edward E. (Eds.), Handbook of Econometrics, vol. 6. Elsevier, pp. 5549–5632.
Compiani, Giovanni, 2022. Market counterfactuals and the specification of multiproduct demand: a nonparametric approach. Quant. Econ. 13 (2), 545–591.
Crawford, Vincent P., Costa-Gomes, Miguel A., Iriberri, Nagore, 2013. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. J. Econ. Lit. 51 (1), 5–62.
Delgado, Miguel A., Hidalgo, Javier, 2000. Nonparametric inference on structural breaks. J. Econom. 96 (1), 113–144.
Friedman, Evan, 2020. Endogenous quantal response equilibrium. Games Econ. Behav. 124, 620–643.
Friedman, Evan, Mauersberger, Felix, 2022. Quantal Response Equilibrium with Symmetry: Representation and Applications. Technical Report. Paris School of Economics.

[26] If the analyst assumes that the marginal error distribution is invariant across games with different numbers of actions, then the set of predicted choice probabilities has a measure of zero generically. Without this invariance assumption on the marginal error distribution, the predicted set has a positive measure but is sufficiently narrow (Goeree et al., 2020).

Goeree, Jacob, Holt, Charles, 2001. Ten little treasures of game theory and ten intuitive contradictions. Am. Econ. Rev. 91, 1402–1422.
Goeree, Jacob, Holt, Charles, Palfrey, Thomas, 2003. Risk averse behavior in generalized matching pennies games. Games Econ. Behav. 45 (1), 97–113.
Goeree, Jacob, Holt, Charles, Palfrey, Thomas, 2005. Regular quantal response equilibrium. Exp. Econ. 8, 347–367.
Goeree, Jacob, Holt, Charles, Palfrey, Thomas, 2020. Stochastic game theory for social science: a primer on quantal response equilibrium. In: Handbook of Experimental Game Theory, pp. 8–47.
Goeree, Jacob, Holt, Charles, Louis, Philippos, Palfrey, Thomas, Rogers, Brian, 2019. Rank-dependent choice equilibrium: a non-parametric generalization of QRE. In: Handbook of Research Methods and Applications in Experimental Economics, pp. 8–47.
Goldfarb, Avi, Xiao, Mo, 2011. Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. Am. Econ. Rev. 101 (7), 3130–3161.
Golman, Russell, 2011. Quantal response equilibria with heterogeneous agents. J. Econ. Theory 146 (5), 2013–2028.
Golman, Russell, 2012. Homogeneity bias in models of discrete choice with bounded rationality. J. Econ. Behav. Organ. 82 (1), 1–11.
Greiner, Ben, 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.
Haile, Philip, Hortacsu, Ali, Kosenok, Grigory, 2008. On the empirical content of quantal response equilibrium. Am. Econ. Rev. 98, 180–200.
Halevy, Yoram, Hoelzemann, Johannes, Kneeland, Terri, 2023. Magic Mirror on the Wall, Who Is the Smartest One of All? Technical Report. The University of Toronto.
Harrison, G.W., Cox, James C., 2008. Risk Aversion in Experiments. JAI Press Inc.
Hoelzemann, Johannes, Li, Hongyi, 2022. Coordination in the Network Minimum Game. Technical Report. The University of Toronto.
Hotz, V. Joseph, Miller, Robert A., 1993. Conditional choice probabilities and the estimation of dynamic models. Rev. Econ. Stud. 60 (3), 497–529.
Klein, Roger W., Spady, Richard H., 1993. An efficient semiparametric estimator for binary response models. Econometrica 61 (2), 387–421.
Lewbel, Arthur, 2000. Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables. J. Econom. 97 (1), 145–177.
Liu, Nianqing, Vuong, Quang, Xu, Haiqing, 2017. Rationalization and identification of binary games with correlated types. J. Econom. 201 (2), 249–268.
McKelvey, Richard, Palfrey, Thomas, 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.
Melo, Emerson, Pogorelskiy, Kirill, Shum, Matthew, 2019. Testing the quantal response hypothesis. Int. Econ. Rev. 60, 53–74.
Müller, Hans-Georg, 1992. Change-points in nonparametric regression analysis. Ann. Stat. 20 (2), 737–761.
Nagel, Rosemarie, 1995. Unraveling in guessing games: an experimental study. Am. Econ. Rev. 85 (5), 1313–1326.
Norets, Andriy, Takahashi, Satoru, 2013. On the surjectivity of the mapping between utilities and choice probabilities. Quant. Econ. 4 (1), 149–155.
Ochs, Jack, 1995. Games with unique, mixed strategy equilibria: an experimental study. Games Econ. Behav. 10 (1), 202–217.
Oprea, Ryan, forthcoming. Decisions under risk are decisions under complexity. Am. Econ. Rev.
Rogers, Brian W., Palfrey, Thomas R., Camerer, Colin F., 2009. Heterogeneous quantal response equilibrium and cognitive hierarchies. J. Econ. Theory 144 (4), 1440–1467.
Sørensen, Jesper R.-V., Fosgerau, Mogens, 2022. How McFadden met Rockafellar and learned to do more with less. J. Math. Econ., 100.
Stahl, Dale, Wilson, Paul, 1994. Experimental evidence on players' models of other players. J. Econ. Behav. Organ. 25 (3), 309–327.
Stahl, Dale, Wilson, Paul, 1995. On players' models of other players: theory and experimental evidence. Games Econ. Behav. 10 (1), 218–254.
Train, Kenneth, 2009. Discrete Choice Methods with Simulation. Cambridge University Press.
Xie, Erhao, 2018. Inference in Games without Nash Equilibrium: An Application to Restaurants' Competition in Opening Hours. Technical Report. Bank of Canada.
Xie, Erhao, 2022. Nonparametric Identification of Incomplete Information Discrete Games with Non-Equilibrium Behaviors. Technical Report. Bank of Canada.