Non-Parametric Identification and Testing of Quantal Response Equilibrium[†]

Johannes Hoelzemann[Φ], Ryan Webb[Ψ], and Erhao Xie[Ω]

[Φ]University of Vienna

[Ψ]University of Toronto

[Ω]Bank of Canada

February 28, 2023

Abstract

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 variation of players' choices across a series of games, we derive a testable implication of QRE. We then show that both the utility function and the distribution of errors are non-parametrically over-identified. This over-identification result suggests a straightforward testing procedure of QRE which achieves the desired type-1 error and maintains a small type-2 error. Finally, we apply the method to an experimental study of the matching pennies game. Our results show that the Quantal Response hypothesis cannot be rejected for the majority of participants when the specification of utility and error distribution is sufficiently flexible and heterogeneous. However, if strong assumptions — such as risk neutrality, logistically distributed errors, and homogeneity — are imposed, QRE is highly rejected.

JEL Classification: C14, C57, C92.

Keywords: Quantal Response Equilibrium (QRE), Falsifiability, Non-Parametric Identification, Experiment, Matching Pennies.

[†] Financial support by the University of Toronto and the University of Vienna is gratefully acknowledged. The experimental part of this study was conducted according to "standard procedures for experimental research," the official guidelines of the University of Vienna's VCEE (Vienna Center for Experimental Economics). The views in this paper do not necessarily represent the views of the Bank of Canada.

1 Introduction

In many strategic settings, choice behaviour systematically deviates from the canonical solution concept of Nash Equilibrium (NE). These deviations have been documented using both experimental data from 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 the failures of NE, the *Quantal Response Equilibrium* (QRE) of McKelvey and Palfrey (1995) has been proposed as an alternative equilibrium concept. QRE extends the random utility framework to strategic settings, where the expected utility of each action is randomly perturbed. This "error" can be interpreted as either a noisy decision process or private information on the part of the player. QRE is defined as a fixed point in the space of the choice probabilities implied by this error. By incorporating decision errors into strategic settings, yet preserving the equilibrium feature, QRE makes predictions about behavior in games that reduce to NE as noise vanishes.

While QRE has successfully explained many deviations from Nash Equilibrium[1] and has become an important benchmark in game theory (Goeree et al., 2020), the empirical falsifiability of QRE requires assumptions by the analyst on the distribution of utility. As long as utility errors are not i.i.d. across players' actions, QRE can rationalize any observed choice behaviour *within* a game (Haile et al., 2008). To ensure falsifiability, the literature has developed two approaches. The first approach has imposed additional restrictions on the distribution of random utility across actions and players, yielding testable implications of QRE within repetitions of the same game. In particular, it assumes that each players' utility function is known, identical, and given by monetary payoffs (henceforth, the *known utility* assumption).[2] However, a formal statistical test

[1]For recent work on endogenizing QRE, see Friedman (2020), as well as on an axiomatic variant of QRE, see 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).

[2]Two important examples of this approach are the regular QRE by Goeree et al. (2005) and the rank-dependent choice equilibrium by Goeree et al. (2019). Each considers a restriction that is weaker than i.i.d. errors.

for QRE under these assumptions has not been derived.

In contrast, this paper follows a second approach that exploits the variation of players' choices across games (or more accurately, versions of a game with variation in payoffs). The idea of exploiting cross-game variation was first suggested by Haile et al. (2008) and subsequently explored by Melo et al. (2019) who non-parametrically specify the error distribution but still impose the *known utility* assumption. In comparison, Aguirregabiria and Xie (2021) consider a non-parametric utility function but impose that the errors follow a known distribution (henceforth, the *distributional* assumption). In our framework, each player's utility and distribution of the random error are both unknown non-parametric functions. Under the assumption that the error distribution remains unchanged across games (henceforth, the *invariance* assumption), we derive a testable necessary condition for QRE in any normal-form game with a finite number of players and actions. We then show that both the utility and the distribution of the random perturbation are non-parametrically (over)-identified under the hypothesis of QRE.[3] This implies a straight-forward statistical test for rejecting QRE which does not require repeated choices from identical games. This test achieves ideal type-1 error rates and therefore guards against over-rejection of QRE in sample sizes typical of economic experiments.

To apply QRE in an empirical study, it is important to relax both the known utility and the distributional assumptions. If either is mis-specified, we show via a Monte-Carlo exercise that tests of QRE can substantially over-reject in typical sample sizes from laboratory studies. Mis-specification of the known-utility assumption – which essentially restricts all participants to have a homogeneous risk-neutral utility function – is particularly problematic with type-1 error rates near 50% or higher on a purported 5% test. In laboratory settings where QRE has been traditionally tested, heterogeneous small-stakes risk aversion (as identified by Harrison and Cox (2008) among many others) thus presents a serious empirical hurdle. An important contribution of this paper is to make

[3] Aguirregabiria and Xie (2021) only derive a semi-parametric identification result. Moreover, Goeree et al. (2003) is an early attempt to relax the known utility assumption. However, they assume a parametric utility function and impose the distributional restriction.

laboratory tests of QRE robust to heterogeneous risk preferences and error distributions. Commonly used distributional assumptions such as Logit and Probit impose strong shape restrictions on the random perturbation, and are considered mainly due to their statistical convenience. These distributional restrictions could be mis-specified, especially when the analyst fits aggregate data that consists of heterogeneous participants.[4] Again, in our Monte-Carlo study we observe substantial over-rejection of QRE when the utility error distribution is mis-specified. The known utility and distributional assumptions are even more problematic in field data where the monetary payoff (e.g., profits) is usually unobserved by the analyst and instead must be estimated. By considering a non-parametric unknown utility function and error distribution, this paper unlocks the possibility of applying QRE to field data.[5]

Our proposed test of QRE follows from conditions on a player's behaviour across similar games that only vary in the magnitude of player-specific payoffs. For some intuition, we describe our procedure in a common example of experimental economics: the matching pennies game. Suppose that the analyst designs a series of 2×2 games where the magnitude of player-specific payoffs varies. Under some regularity conditions, there are multiple pairs of games such that the choice probabilities for player i must be the same. Given any three such pairs, we derive an equality condition that QRE imposes on the relative change in player $-i$'s choice probabilities across these games. This condition is a testable implication of QRE that does not depend on any further assumptions on the utility function or the error distribution. In games with more players and/or more actions, the testable implication of QRE is a rank restriction on a matrix that solely depends on player $-i$'s choice probabilities.

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

[5] To give more detail, QRE has an identical mathematical structure as the Bayesian Nash Equilibrium of an incomplete information game where private information is independent across players. The latter framework is commonly estimated using field data, mainly in empirical industrial organization. Identification results based on the latter framework have been obtained by Bajari et al. (2010), Liu et al. (2017), and Xie (2022).

Our identification results provide a means to implement this test in practice. In particular, we show that both the utility function and the distribution of the random perturbation are non-parametrically over-identified. Consequently, the testable implication of QRE can be viewed as an over-identification test for an additive reduced form *bias term* into each player's expected utility function. This bias term is interpreted as the departure from QRE and is equal to zero if and only if QRE is true. Therefore, testing QRE can be easily conducted via a null hypothesis test using existing econometric tools, such as the likelihood ratio test. In small samples, additional power can be achieved by imposing a semi-parametric functional form for utility, like CRRA, that allows risk-aversion to differ at the participant-level. Our Monte-Carlo simulation results suggest that our test is well-powered and achieves the desired rejection rate in sample sizes typical of economic experiments.

Our identification results also have an important implication for the estimation of QRE. Specifically, our non-parametric specification can be viewed as a population level fit of QRE which allows heterogeneous error distributions across subjects (Golman, 2011). This approach has not been exploited empirically because – even under the known utility assumption – there is no previous identification result for a non-parametric error distribution, only a test of QRE (Melo et al., 2019). The identification results in this paper provide a means to fit heterogeneous QRE at the population level.

Finally, we conduct a laboratory experiment ($N = 100$) of the matching pennies game to test QRE using our procedure. We demonstrate a substantial reduction in rejections of QRE once the known utility and known distribution assumptions are eliminated, with rejection rates dropping from 90% to 40% of participants. We also find that QRE with an unknown error distribution fits the data substantially better than previous methods, both in-sample and out-of-sample. These results highlight the importance of relaxing both the known utility assumption and the distributional assumption in applied settings.

The rest of the paper proceeds as follows. Section 2 reviews QRE in 2×2 games, and Section 3 presents the testable implication and identification results. 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 also in 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 by a_i, from their action set $A_i = \{0,1\}$. Moreover, let $\mathbf{a} = (a_i, a_{-i}) \in A = A_i \times A_{-i}$ be an action profile of this game. Player i's utility of $\mathbf{a} = (a_i, a_{-i})$ is represented by $\pi_i(\mathbf{m}_i, a_i, a_{-i})$.

The expression of the utility function $\pi_i(\mathbf{m}_i, \mathbf{a})$ includes control variables $\mathbf{m}_i$. In field data, such as the entry game estimated in empirical industrial organization, the control variables $\mathbf{m}_i$ could include market conditions and firm i's characteristics. In experimental data, $\mathbf{m}_i$ could be a vector that consists of player i's monetary reward for each action profile. For instance, $\mathbf{m}_i = \big(m_i(a_i = 0, a_{-i} = 0), m_i(a_i = 1, a_{-i} = 0), m_i(a_i = 0, a_{-i} = 1), m_i(a_i = 1, a_{-i} = 1)\big)'$, where $m_i(\mathbf{a})$ represents player i's monetary payoff of the profile $\mathbf{a}$. Our general specification $\pi_i(\cdot)$ allows the utility of the profile $\mathbf{a}$ to depend on the entire vector $\mathbf{m}_i$. For instance, the utility depends not only on the monetary reward of the chosen profile $\mathbf{a}$, but also on the payoffs of other un-chosen profiles (i.e., reference-dependent preference). We derive our key results under this general specification. The experimental literature commonly restricts utility to depend only on received payoffs; for instance, $\pi_i(\mathbf{m}_i, \mathbf{a}) = u[m_i(\mathbf{a})]$, where $u(\cdot)$ represents utility as a function of money. This additional restriction eases estimation and we will exploit this restriction in our Monte-Carlo exercise and experimental analysis. Finally, we assume that players know their utility function $\pi_i(\cdot)$ but the analyst does not.

Our objective is to derive a robust test of QRE without prior information about $\pi_i(\cdot)$, and to identify $\pi_i(\cdot)$ when QRE holds in the data. To achieve these objectives, we consider a *non-parametric* specification of $\pi_i(\cdot)$ under the following assumption.

Assumption 1. $\mathbf{m}_i$ *includes continuous variables. Moreover,* $\pi_i(\mathbf{m}_i, a_i, a_{-i})$ *is bounded and continuous in its continuous arguments.*

In the above framework, exogenous variation of $(\mathbf{m}_i, \mathbf{m}_{-i})$ generates a series of different games. For instance, consider the matching pennies game illustrated in Table 1 that we use throughout the paper.[6]

Table 1: Monetary Payoff Matrix of Matching Pennies ($m_1 > 8$, $m_2 > 8$)

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

The numbers in each cell represent the monetary reward for the corresponding action profile. The variables (m_1, m_2) are controlled by the analyst and vary across a series of trials, generating the required variation to test QRE and to identify the model primitives.

Let $p_{-i}(\mathbf{m})$ denote player $-i$'s choice probability of action $a_{-i} = 1$, in the game with control variables $\mathbf{m} = (\mathbf{m}_i, \mathbf{m}_{-i})$. Note that in strategic settings, a player's choice probability depends on all players' control variables $\mathbf{m} = (\mathbf{m}_i, \mathbf{m}_{-i})$. Given the above choice probability, the expected utility of player i's action a_i is:

$$E\pi_i[\mathbf{m}_i, a_i, p_{-i}(\mathbf{m})] = \pi_i(\mathbf{m}_i, a_i, a_{-i} = 0) \cdot [1 - p_{-i}(\mathbf{m})] + \pi_i(\mathbf{m}_i, a_i, a_{-i} = 1) \cdot p_{-i}(\mathbf{m}).$$

QRE places an error distribution on this expected utility. Specifically, let $\varepsilon_i(a_i)$ denote the error on player i expected utility of action a_i. Consequently, player i will choose

[6]For simplicity, we only vary *one* action profile's payoff for each player in this example. Our results hold in general settings where the payoff of *every* action profile varies.

$a_i = 1$ if and only if the following condition holds:

$$
\begin{aligned}
& E\pi_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 1) \geq E\pi_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 0) \\
\Leftrightarrow\ & \varepsilon_i(a_i = 0) - \varepsilon_i(a_i = 1) \leq E\pi_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})] - E\pi_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})]. \quad (1)
\end{aligned}
$$

To derive choice probabilities, let $F_i(\cdot)$ be a continuous and strictly increasing cumulative distribution function (C.D.F.) of $\tilde{\varepsilon}_i = [\varepsilon_i(a_i = 0) - \varepsilon_i(a_i = 1)]$. Importantly, $F_i(\cdot)$ is a *non-parametric* function that is unknown to the analyst, but with the following assumption:

Assumption 2. *$F_i(\cdot)$ is independent of $(\mathbf{m}_i, \mathbf{m}_{-i})$.*

Assumption 2 is known as the *invariance assumption* and simply states that the distribution does not change over the variation in games. It is commonly maintained in empirical applications of QRE, including formal tests of QRE (Goeree et al., 2020; Melo et al., 2019; Aguirregabiria and Xie, 2021). In an extension offered in the appendix, we relax this 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 (1) and $F_i(\cdot)$, player i's choice probability of action $a_i = 1$ takes the following form:

$$
p_i(\mathbf{m}) = F_i[E\pi_i(\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})) - E\pi_i(\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m}))]. \quad (2)
$$

QRE implicitly assumes that each player forms correct beliefs about other players' choice probabilities. Consequently, the choice probabilities satisfy a fixed-point condition under QRE. We refer to this as the *Global QRE condition*:

Definition 1. *The vector $\left(p_i(\mathbf{m}), p_{-i}(\mathbf{m})\right)'$ denotes the QRE choice probabilities if and only if the following condition holds:*

$$
p_i(\mathbf{m}) = F_i[E\pi_i(\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})) - E\pi_i(\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m}))], \forall i \textit{ and } \mathbf{m}. \quad (3)
$$

By Brouwer's fixed point theorem, any game has at least one QRE. Moreover, in many games, there can exist multiple QREs.

3 Testable Implication of QRE and Identification Results

We now demonstrate that the QRE restrictions in Equation (3) can be tested by an analyst who observes players' choices. Moreover, we will demonstrate that the utility function $\pi_i(\cdot)$ and the distribution function $F_i(\cdot)$ are non-parametrically identified when QRE is satisfied in the data. To derive the testable implication, we will work directly with each player's choice probability since these can be consistently estimated from the choice data (i.e., we assume that $p_i(\mathbf{m})$ and $p_{-i}(\mathbf{m})$ are observed by the analyst). In practice, we will show that this simplifying assumption can be relaxed to incorporate datasets for which only a single choice is observed from each $(\mathbf{m}_1, \mathbf{m}_2)$ pair (as in our experiment). 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 Testable Implication of QRE

We begin by noting that since $F_i(\cdot)$ is strictly increasing, we can invert the QRE conditions (3):

$$F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i})] = E\pi_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i})] - E\pi_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i})], \ \forall i. \tag{4}$$

Equation (4) contains all the model restrictions that are imposed on player i's behaviors. Following Aguirregabiria and Magesan (2020) and Aguirregabiria and Xie (2021), we consider any three realizations of $\mathbf{m}_{-i}$, denoted by $\mathbf{m}_{-i}^1$, $\mathbf{m}_{-i}^2$, and $\mathbf{m}_{-i}^3$. When $p_i(\mathbf{m}_i, \mathbf{m}_{-i}^1) \neq p_i(\mathbf{m}_i, \mathbf{m}_{-i}^2)$, evaluating Equation (4) at the above three realizations implies:

$$\frac{F_i^{-1}[p_i(\mathbf{m}_i,\mathbf{m}_{-i}^3)]-F_i^{-1}[p_i(\mathbf{m}_i,\mathbf{m}_{-i}^1)]}{F_i^{-1}[p_i(\mathbf{m}_i,\mathbf{m}_{-i}^2)]-F_i^{-1}[p_i(\mathbf{m}_i,\mathbf{m}_{-i}^1)]}=\frac{p_{-i}(\mathbf{m}_i,\mathbf{m}_{-i}^3)-p_{-i}(\mathbf{m}_i,\mathbf{m}_{-i}^1)}{p_{-i}(\mathbf{m}_i,\mathbf{m}_{-i}^2)-p_{-i}(\mathbf{m}_i,\mathbf{m}_{-i}^1)},\ \forall \mathbf{m}_i. \tag{5}$$

A complete derivation of Equation (5) is left to the appendix, however its intuition is relatively simple. Variation of $\mathbf{m}_{-i}$ will affect player $-i$'s utility and change their choice probabilities $p_{-i}(\cdot)$. In strategic settings, this change of $p_{-i}(\cdot)$ would then have an impact on player i's choice probability. Suppose we fix $\mathbf{m}_i$ so that player i's utility of each action profile remains unchanged. The only reason that $p_i(\cdot)$ will vary with $\mathbf{m}_{-i}$ is due to $\mathbf{m}_{-i}$'s impact on $p_{-i}(\cdot)$. Put differently, when only $\mathbf{m}_{-i}$ varies, the relative change of $p_i(\cdot)$ reveals information about the relative change of $p_{-i}(\cdot)$. For any three realizations of $\mathbf{m}_{-i}$, it is therefore possible to cancel player i's utility function and obtain the relationship as shown by Equation (5).

Note that Equation (5) depends on the distribution function $F_i(\cdot)$. In previous work, Aguirregabiria and Magesan (2020) and Aguirregabiria and Xie (2021) impose the distributional assumption (e.g., Logit) so that $F_i(\cdot)$ is known to the analyst. Under this assumption, Equation (5) becomes a testable restriction of QRE. In contrast, our objective is to cancel $F_i(\cdot)$ and obtain a testable implication of QRE that is robust to any distribution function.

To see how this is possible, we use an example of the matching pennies game depicted in Table 1. In Figure 1, we solve each player's QRE choice probabilities for a series of different values of (m_1, m_2).[7] In particular, m_1 takes on two values: 10 (blue) and 16 (green) and we plot each player's choice probability as a function of m_2.

Consider a pair of games (realizations of (m_1, m_2)) that respectively lie on the blue and green line. Since assumption 1 implies that $p_i(m_1, m_2)$ is continuous in both arguments, there exist *infinite* pairs of games such that player 1's choice probability remains

[7] For illustrative purposes only, our simulation simply assumes players' utilities equal their monetary payoffs ($\pi_i(\mathbf{m}_i,\mathbf{a}) = m_i(\mathbf{a})$) and that $F(\cdot)$ takes a logistic functional form.

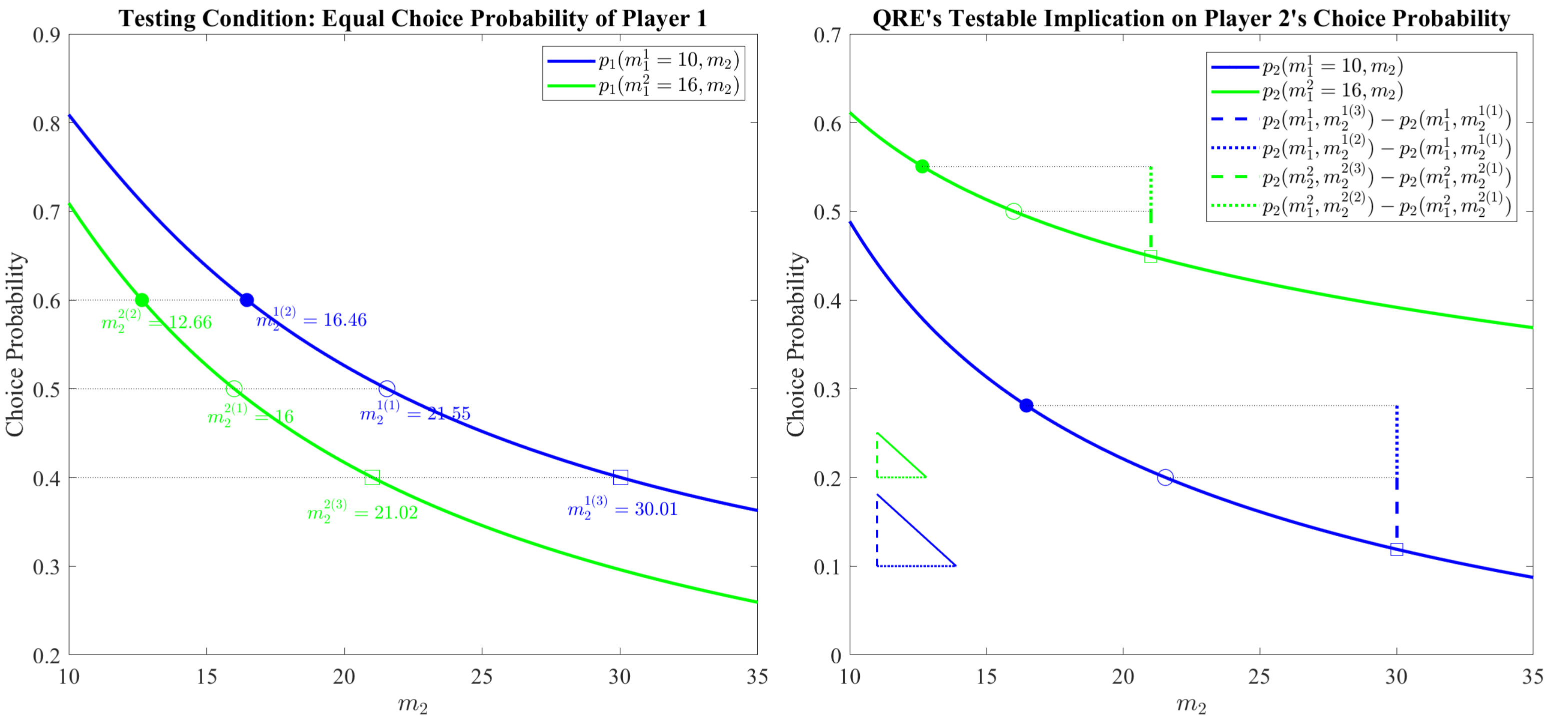

Figure 1: Testable Implication of QRE

constant *within* pairs. Figure 1 (left panel) labels three such pairs that satisfy the condition of *equal choice probability*; for instance, $p_1(m_1^1, m_2^{1(l)}) = p_1(m_1^2, m_2^{2(l)})$ for $l = 1, 2, 3$. These three pairs, together with Equation (5), jointly imply the following:

$$
\begin{array}{ccc}
\dfrac{F_i^{-1}[p_i(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(3)})]-F_i^{-1}[p_i(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(1)})]}{F_i^{-1}[p_i(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(2)})]-F_i^{-1}[p_i(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(1)})]} & = & \dfrac{p_{-i}(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(3)})-p_{-i}(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(2)})-p_{-i}(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(1)})} \\
|| & & || \\
\dfrac{F_i^{-1}[p_i(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(3)})]-F_i^{-1}[p_i(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(1)})]}{F_i^{-1}[p_i(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(2)})]-F_i^{-1}[p_i(\mathbf{m}_i^1,\mathbf{m}_{-i}^{1(1)})]} & = & \dfrac{p_{-i}(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(3)})-p_{-i}(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(2)})-p_{-i}(\mathbf{m}_i^2,\mathbf{m}_{-i}^{2(1)})}
\end{array}. \tag{6}
$$

Due to the equal choice probability condition, the terms on both rows' left-hand sides in Equation (6) are equal, therefore, the terms on the right-hand sides are also equal. Each element in Equation (6) has its corresponding representation in Figure 1. In particular, each row in Equation (6) describes the mapping from the relative change of player i's choice probability (i.e., left panel in Figure 1) to the relative change of $p_{-i}(\cdot)$

(i.e., right panel). Since the relative change of player $-i$'s choice probabilities can be consistently estimated, this serves as a testable implication of QRE. This condition can be visualized by the two similar triangles with sides equal to the change of player $-i$'s choice probabilities.

We formally state the testable implication of QRE in Proposition 1:

Proposition 1. *Under Assumptions 1–2, for any three pairs of realizations* $(\mathbf{m}_i, \mathbf{m}_{-i})$ *that satisfy the following conditions:*

Pair 1: realizations $(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})$ *and* $(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})$ *such that* $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})$.

Pair 2: realizations $(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})$ *and* $(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})$ *such that* $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})$.

Pair 3: realizations $(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)})$ *and* $(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)})$ *such that* $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)})$.

Given these pairs, QRE implies the following testable restriction:

$$\frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} = \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}, \tag{7}$$

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

Proof. A direct implication of Equation (6). □

Proposition 1 describes a necessary condition of QRE. Given three pairs of games that satisfy the equal choice probability condition for player i, QRE requires that the relative change in player $-i$'s choice probabilities across these games must be identical. Figure 2 illustrates how the equality in Proposition 1 no longer holds when choices are generated by Level-2 behavior rather than behavior consistent with QRE. Under the same conditions for player 1's behavior, the relative change in player 2's choice probabilities are not identical, thus QRE can be rejected.

3.2 Identification of Utility and Error Distribution under QRE

Our second key result establishes the non-parametric identification of the utility function and the distribution function. We first define the utility difference of player i's

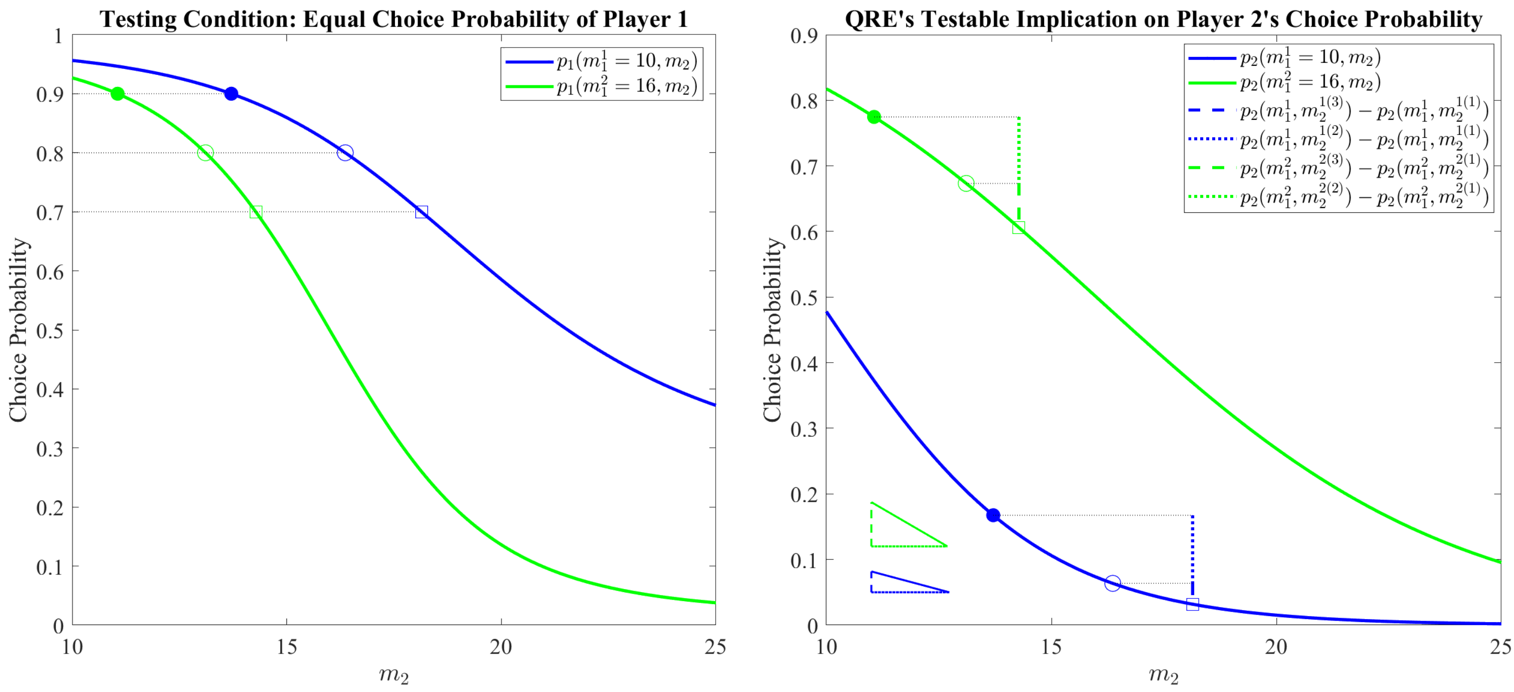

Figure 2: Violation of the Testable Implication under Level-2 Behavior

two actions, given $\mathbf{m}_i$ and the other player's action a_{-i} to be $\tilde{\pi}_i(\mathbf{m}_i, a_{-i}) = \pi_i(\mathbf{m}, a_i = 1, a_{-i}) - \pi_i(\mathbf{m}, a_i = 0, a_{-i})$. Our identification results also require some standard normalization, as summarized by Assumption 3.

Assumption 3. *(a) There exists a realization $\mathbf{m}_i = \mathbf{m}_i^1$, such that* $\left|\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)\right| = \left|\pi_i(\mathbf{m}_i^1, a_i = 1, a_{-i} = 0) - \pi_i(\mathbf{m}_i^1, a_i = 0, a_{-i} = 0)\right| = 1$.

(b) Median$(\tilde{\varepsilon}_i) = 0$, where $\tilde{\varepsilon}_i = \varepsilon_i(a_i = 0) - \varepsilon_i(a_i = 1)$.

In the discrete choice literature, Assumption 3(a) is referred to as a *scale normalization* and Assumption 3(b) as a *location normalization* (Train, 2009). Since any affine transformation of the utility function represents the same preference and predicts the same choice, these normalizations impose no restrictions on players' behaviors and are innocuous (a complete explanation of these normalizations is in the appendix).[8] However we do clarify one important aspect of Assumption 3(b). This normalization is required only for the general specification of the utility function $\pi_i(\mathbf{m}_i, \mathbf{a})$. It is not necessary in ex-

[8]We consider the transformation to apply to both $\pi_i(\cdot)$ and $\varepsilon_i(\cdot)$. If the transformation is only applied to $\pi_i(\cdot)$, it will generally affect a player's choice. See appendix for more details.

perimental settings that define utility over monetary outcomes only, $\pi_i(\mathbf{m}_i, \mathbf{a}) = u_i[m_i(\mathbf{a})]$. This specification imposes the additional restriction that the utility of any two action profiles with identical monetary outcomes must be equal. This restriction identifies the median of $\tilde{\varepsilon}_i$, therefore Assumption 3(b) is redundant (and even testable) in this scenario.

We now describe our identification result. Let $\mathcal{P}_i(\mathbf{m}_i^1)$ denote the *image* of player i's choice probability function $p_i(\mathbf{m}_i = \mathbf{m}_i^1, \mathbf{m}_{-i})$. This choice probability function fixes $\mathbf{m}_i = \mathbf{m}_i^1$ but varies $\mathbf{m}_{-i}$. Given that $p_i(\cdot)$ is a continuous function, we know that $\mathcal{P}_i(\mathbf{m}_i^1) = [\min_{\mathbf{m}_{-i}} p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}), \max_{\mathbf{m}_{-i}} p_i(\mathbf{m}_i^1, \mathbf{m}_{-i})]$. Next, denote $int[\mathcal{P}_i(\mathbf{m}_i^1)]$ as the set of all interior points in $\mathcal{P}_i(\mathbf{m}_i^1)$. With this notation, we establish the non-parametric identification of the inverse of the distribution function:

Proposition 2. *Given Assumptions 1–3, suppose that the QRE restrictions are satisfied whenever $\mathbf{m}_i = \mathbf{m}_i^1$, regardless of the realization of $\mathbf{m}_{-i}$. Suppose further that $1/2 \in int[\mathcal{P}_i(\mathbf{m}_i^1)]$, then $F_i^{-1}(p)$ is point identified $\forall p \in \mathcal{P}_i(\mathbf{m}_i^1)$.*

Since the distribution function $F_i(\cdot)$ is invertible given its strict monotonicity, the identification of its inverse – as in Proposition 2 – implies the identification of $F_i(\cdot)$.[9]

The "local" nature of Proposition 2 is particularly noteworthy. To identify $F_i^{-1}(\cdot)$, Proposition 2 imposes QRE restrictions for *only one* realization of $\mathbf{m}_i = \mathbf{m}_i^1$. For any other realization of $\mathbf{m}_i$, players' behaviors are unrestricted (e.g., they need not satisfy QRE). We refer to this condition as the *local QRE restriction*. It is a substantially weaker condition than the *global QRE restriction* that assumes QRE in the entire space of $(\mathbf{m}_i, \mathbf{m}_{-i})$; that is, Definition 1. Therefore the global QRE condition imposes many more restrictions that substantially over-identify $F_i(\cdot)$.[10]

[9] This identification result requires player i's choice probability function $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i})$ to go through the point $1/2$. This condition is easily satisfied if $\mathbf{m}_{-i}$ has sufficient variation. For instance, consider the matching pennies games presented in Table 1. Figure 1 presents player 1's choice probability functions under QRE restrictions. It is clear that all these functions go through the point $1/2$.

[10] Proposition 2 identifies $F_i^{-1}(p)$ in the region where $p \in \mathcal{P}_i(\mathbf{m}_i^1)$. When p lies outside this region, the analyst has to impose QRE restrictions for more realizations of $\mathbf{m}_i$ to identify $F_i^{-1}(\cdot)$. Specifically, consider H realizations of $\mathbf{m}_i$, denoted by $\mathbf{m}_i^1$ to $\mathbf{m}_i^H$. With an appropriate choice of these H realizations, $\cup_{h=1}^H \mathcal{P}_i(\mathbf{m}_i^h)$ could well approximate the image of player i's probability function $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$ where variations of both $\mathbf{m}_i$ and $\mathbf{m}_{-i}$ are considered. Therefore, if the analyst imposes the local QRE restriction for

Given that $F_i(\cdot)$ has been identified, we now establish the non-parametric identification of the difference of utility function $\tilde{\pi}_i(\cdot)$. It is well known in the discrete choice literature that the analyst can, at most, identify such utility differences (Train, 2009). Furthermore, in experimental settings that specify $\pi_i(\mathbf{m}_i, \mathbf{a}) = u_i[m_i(\mathbf{a})]$, it is standard to normalize the utility of \$0 or the minimum payoff to 0. With either of these two location normalizations, the function $u_i(\cdot)$ is non-parametrically identified.

Proposition 3. *Suppose that the conditions met in Proposition 2 hold so that $F_i(\cdot)$ is identified. Moreover, consider two realizations of $\mathbf{m}_{-i}$, say $\mathbf{m}^1_{-i}$ and $\mathbf{m}^2_{-i}$. Suppose that QRE restrictions are satisfied whenever $\mathbf{m}_{-i} = \mathbf{m}^1_{-i}$ or $\mathbf{m}_{-i} = \mathbf{m}^2_{-i}$, regardless of the realization of $\mathbf{m}_i$. These conditions imply that the difference of utility function $\tilde{\pi}_i(\mathbf{m}_i, a_{-i}) = \pi_i(\mathbf{m}_i, a_i = 1, a_{-i}) - \pi_i(\mathbf{m}_i, a_i = 0, a_{-i})$ is identified $\forall \mathbf{m}_i, a_{-i}$.*

Proposition 3 also requires a local QRE restriction, but it is slightly different than the one required by Proposition 2. While Proposition 2 imposes the restriction on one realization of the *own* control variables $\mathbf{m}_i$, Proposition 3 considers the local QRE condition on two realizations of the *other player*'s $\mathbf{m}_{-i}$. Using the matching pennies game in Table 1 as a motivating example, Figure 3 illustrates the two local QRE restrictions in the space of (m_i, m_{-i}).

The black line represents all combinations of (m_i, m_{-i}) that satisfy $m_i = m_i^1$. If the local QRE restriction is imposed on all points on this line, then the distribution function $F_i(\cdot)$ is identified as stated in Proposition 2. Moreover, if the local QRE restriction is imposed on the two blue lines (i.e., the combinations of (m_i, m_{-i}) with the property that m_{-i} equals either m^1_{-i} or m^2_{-i}); then the difference of utility function $\tilde{\pi}_i(\cdot)$ is also identified as shown in Proposition 3. Importantly, since m_i and m_{-i} are continuous variables, the region required for identification by the local QRE restriction has a measure of zero (one black line and two blue lines as opposed to the entire space). Therefore the global

each of these H realizations, function $F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i})]$ could be point identified for (almost) the entire image of the function $p_i(\mathbf{m}_i, \mathbf{m}_{-i})$. Importantly, since $\mathbf{m}_i$ includes continuous variables, the region with the local QRE restriction (i.e., H realizations of $\mathbf{m}_i$) still has a measure of zero.

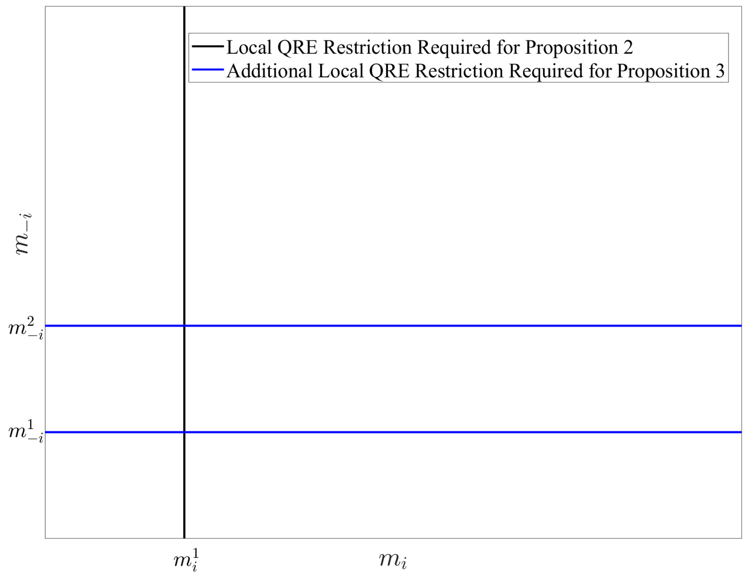

Figure 3: Illustration of the Local QRE Restrictions

QRE condition (Definition 1) substantially over-identifies the model primitives.

The over-identification results from Propositions 2 and 3 have an important implication on the test of QRE. Recall the testable implication of QRE described by Proposition 1. It is challenging to construct a statistical test based directly on Equation (7) because it requires the analyst to choose three pairs of $(\mathbf{m}_i, \mathbf{m}_{-i})$ such that the equal choice probability condition holds. To determine these three pairs in practice, one has to estimate player i's choice probability with some estimation error which enters Equation (7) in a cumbersome manner. Therefore it is difficult to derive the limiting distribution of this test.[11]

Alternatively, our identification results suggest that Proposition 1 can be also viewed as an over-identification test which substantially simplifies the testing procedure. To see how, consider the choice probability function plotted in the left panel of Figure 1. Proposition 2 states that $F_i(\cdot)$ can be identified if QRE is imposed on either the blue or green line. Therefore, if QRE holds globally, the $F_i(\cdot)$ identified from each line must be iden-

[11]A test based on Proposition 1 takes the form: $\mathbb{1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})] \cdot \mathbb{1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})] \cdot \mathbb{1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)})] \cdot [\frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} - \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}]$. The first-step estimation error enters into a non-linear and discontinuous indicator function $\mathbb{1}(\cdot)$.

tical. This is a standard over-identification test. Next, recall Equation (6). The equality of $F_i(\cdot)$ between the blue and green lines (i.e., the left hand side column) directly leads to the testable implication in Proposition 1. However, if QRE is not satisfied, then the $F_i(\cdot)$ identified from the blue line will be different from the one identified from the green line. This difference suggests that the two terms on the left hand side of Equation (6) are no longer identical and further implies the violation of the testable implication by Equation (7). Importantly, because Proposition 1 can be interpreted as an over-identification test, we can construct a likelihood ratio test of QRE that is simple to implement in practice (described in Section 4). Our testing procedure also similarly exploits the over-identification restrictions for the utility function, as established in Proposition 3.

4 Over-Identification Test and Monte Carlo Experiment

This section details our testing procedure in our empirical setting. We begin by conducting a Monte Carlo exercise based on the matching pennies game in Table 1. We then use this exercise to illustrate our testing procedure of QRE and study its finite sample properties. We evaluate the test under two scenarios: one where data are generated by QRE behavior, and another in which QRE is not satisfied.

4.1 Design of the Monte Carlo Experiment

The Monte Carlo experiment focuses on the matching pennies game discussed above. In each simulation, we generate a dataset with T trials where $T \in \{200, 1000, 8000\}$. Our empirical framework assumes that m_i is a continuous variable. To mimic this continuity in a computationally feasible manner, we independently draw (m_1, m_2) from a uniform $M \times M$ grid with resolution that depends on T. For $T = 200$, $M = \{10,\ 12, \ldots,\ 46,\ 48\}$. The asymptotic property of our test depends on the number of trials $T \to \infty$. As T

increases, the grid becomes finer and approaches zero.[12]

In line with most experimental studies, we simulate players with a CRRA utility function over monetary payoffs of the form:

$$u_i(m) = m^{\sigma}. \tag{8}$$

The risk preference parameter σ is set to 0.6 to model risk aversion.

For the distribution function $F_i(\tilde{\varepsilon}_i)$, we consider two cases:

$$\begin{aligned} \textit{Symmetric Distribution:} \quad & \tilde{\varepsilon}_i \sim 0.5N(-1.5,1) + 0.5N(1.5,1), \\ \textit{Asymmetric Distribution:} \quad & \tilde{\varepsilon}_i \sim 0.5N(-1.5,0.5) + 0.5N(1.5,2). \end{aligned} \tag{9}$$

In the symmetric case, $\tilde{\varepsilon}_i$ is drawn from a mixture of two normal distributions with equal weight. These two distributions have a standard deviation of 1 and opposite means. Therefore, $\tilde{\varepsilon}_i$ is distributed symmetrically around 0 and follows a bi-modal distribution. The asymmetric case is almost identical except that the two mixing distributions have standard deviations of 0.5 and 2, respectively. Such an adjustment preserves the mean of $\tilde{\varepsilon}_i$ at zero. However, the distribution of $\tilde{\varepsilon}_i$ turns out to be asymmetric and has more density in the negative region compared to the positive region. The probability density functions for both symmetric and asymmetric cases are shown in Figure 4, with a comparison to

[12] For T, the values of m_1 and m_2 are independently and uniformly drawn from a discrete set M, defined below:

$$M = \{10, 10+\tau, 10+2\tau, 10+3\tau, \cdots, 48\}, \text{where } \tau = \frac{2}{\lfloor \sqrt{T/200} \rfloor}.$$

Note that $\lfloor \cdot \rfloor$ is the floor function; for instance $\lfloor x \rfloor = a$ if and only if $a \leq x < a+1$. We introduce this floor function so that set M could be properly defined. Therefore, the simulation asymptotically approximates the scenario that m_i is uniformly drawn from a continuous distribution. We choose this discrete approximation due to two reasons. First, any experimental dataset has a finite discrete set M and our approximation matches this experimental environment. Second, it substantially reduces the computational cost. To see this point. Let S denote the number of Monte Carlo samples or datasets. When m_i is drawn from a continuous distribution, QRE has to be solved for each observation and the computational cost increases in the order of $T \times S$. In contrast, under discrete approximation, QRE only needs to be solved in the space of $M \times M$. The computational cost increases in the order of T. Our simulation chooses a high value of $S = 1000$; therefore, the discrete approximation is computationally efficient.

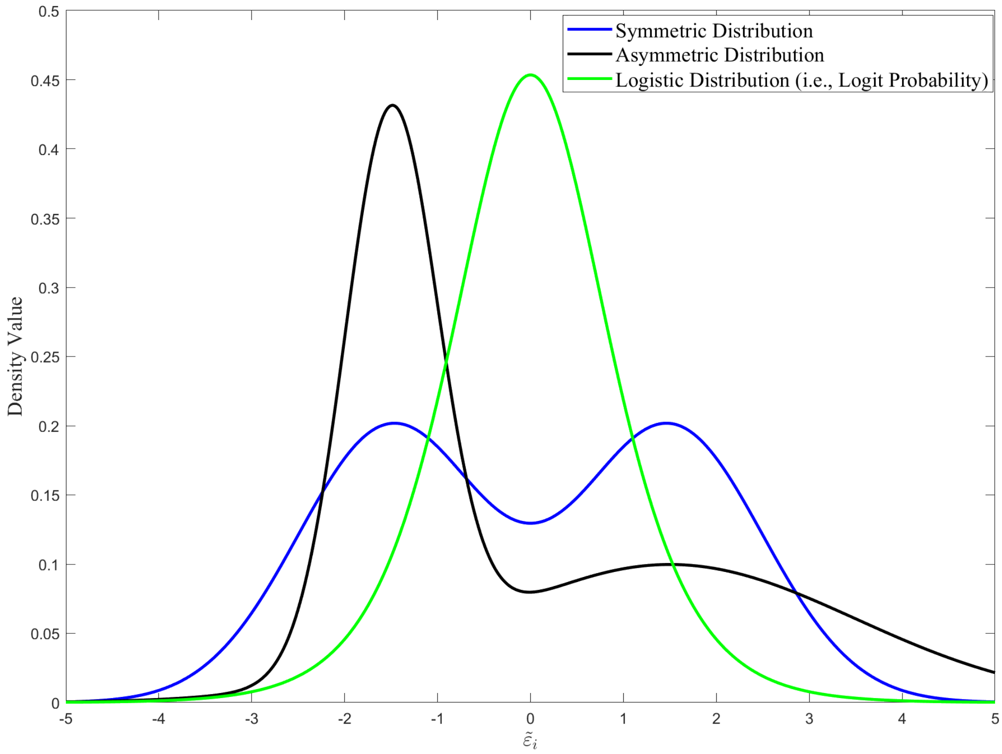

Figure 4: Probability Density Functions of the Random Perturbation $\tilde{\varepsilon}_i$

the Logit specification – that is, logistically distributed $\tilde{\varepsilon}_i$.

The Monte Carlo exercise examines these two distributions instead of other common specifications such as, for example, Logit and Probit. This allows us to investigate the consequences when an analyst uses commonly assumed distribution functions that are, in fact, mis-specified.

Finally, we consider two data generating processes. The first process assumes that data are generated consistently with QRE. In this scenario, the rejection rate of our proposed test should match the pre-specified significance level. Put differently, our test should obtain the desired type-1 error and not over-reject the true hypothesis.

The second process generates data that are inconsistent with QRE. This scenario illustrates whether our test has the power to reject a false hypothesis and achieve a small type-2 error. We 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. However, instead of best responding to such belief, this type plays a Quantal response. In particular, their expected utility of each action is perturbed by the calculation error ε_i. The data are then generated by the behavior of level-k types, where $k \in \{1,2,3\}$.

4.2 Testing Procedure and Results

4.2.1 Estimation

Our estimation procedure borrows elements from recent developments in semi-nonparametric estimation (Chen, 2007). First, note that each player i's choice probability function $p_i(\mathbf{m})$ can be consistently estimated by the following Logit specification with high-order polynomials:

$$\hat{p}_i(\mathbf{m}) = \frac{\exp\left[\sum_{l=0}^{L}(\sum_{h=0}^{l}\hat{\beta}_{l,h}m_1^h m_2^{l-h})\right]}{1+\exp\left[\sum_{l=0}^{L}(\sum_{h=0}^{l}\hat{\beta}_{l,h}m_1^h m_2^{l-h})\right]}, \tag{10}$$

where L represents the order of polynomials. In all models that we simulate in this Monte-Carlo exercise, we find that an order of $L = 3$ approximates the choice probability function well.[13]

As described above, $p_i(\cdot)$ and $p_{-i}(\cdot)$ can be consistently estimated; therefore, this exercise treats them as observables to the analyst.[14] With these choice probabilities, Equation (2) suggests that the only unknown in player i's expected utility $E\pi_i(\cdot)$ is their utility function. If the utility function $u_i(\cdot)$ is non-parametrically specified, it can be approximated by high-order polynomials or other sieve methods. In a typical experimental dataset where the number of observations per participant is not large, a parametric utility function can be imposed. This specification can simply be plugged in to Equation (2).

To non-parametrically approximate the distribution function $F_i(\tilde{\varepsilon}_i)$, we exploit a Logit

[13] Alternatively, it can be estimated by the Kernel or Nadaraya-Watson estimator: $\hat{p}_i(\mathbf{m}) = \frac{\sum_{t=1}^{T} K_h(\mathbf{m}-\mathbf{m}_t)\cdot\mathbb{1}(a_{i,t}=1)}{\sum_{t=1}^{T} K_h(\mathbf{m}-\mathbf{m}_t)}$, where $K_h(\cdot)$ is a kernel with a bandwidth h.

[14] Admittedly, this treatment assumes away the first-step estimation error. In practice, this error could be dealt with using the correction in Chen (2007) or via a bootstrap procedure.

transformation of the non-parametric error distribution, $g_i(\tilde{\varepsilon}_i) = \log[\frac{F_i(\tilde{\varepsilon}_i)}{1-F_i(\tilde{\varepsilon}_i)}]$. Inverting this relationship implies that $F_i(\tilde{\varepsilon}_i) = \frac{\exp[g_i(\tilde{\varepsilon}_i)]}{1+\exp[g_i(\tilde{\varepsilon}_i)]}$, which is the standard Logit formula applied to $g_i(\cdot)$. We then approximate $F_i(\cdot)$ with high-order polynomials applied on $g_i(\cdot)$; for instance, $g_i(\tilde{\varepsilon}_i) = \sum_{l=0}^{L} \beta_{i,l} \cdot \tilde{\varepsilon}_i^l$ where L increases with the sample size T. In practice, we find that an order of $L = 4$ approximates the distribution function in Equation (9) precisely and use it throughout.

Finally, define the difference of player i's expected utility as $E\tilde{\pi}_i[m_i, p_{-i}(\mathbf{m})] = E\pi_i[m_i, a_i = 1, p_{-i}(\mathbf{m})] - E\pi_i[m_i, a_i = 0, p_{-i}(\mathbf{m})]$. Then player i's choice probability under QRE restrictions can be expressed as:

$$p_i^{QRE}(\mathbf{m}) = \frac{\exp\left\{\sum_{l=0}^{L} \beta_{i,l} \cdot E\tilde{\pi}_i[m_i, p_{-i}(\mathbf{m})]^l\right\}}{1+\exp\left\{\sum_{l=0}^{L} \beta_{i,l} \cdot E\tilde{\pi}_i[m_i, p_{-i}(\mathbf{m})]^l\right\}}. \tag{11}$$

By Propositions 2 and 3, the utility function and the distribution function are non-parametrically identified. Therefore, these model primitives can be consistently estimated by maximizing the following log-likelihood function:

$$LL_i^{QRE} = \max_{\sigma_i, \beta_{i,l}} \sum_{t=1}^{T} \left\{ \mathbb{1}(a_{i,t} = 1) \log[p_i^{QRE}(\mathbf{m}_t)] + \mathbb{1}(a_{i,t} = 0) \log[1 - p_i^{QRE}(\mathbf{m}_t)] \right\}. \tag{12}$$

4.2.2 Testing

As described in Section 3, when player i's behavior is inconsistent with QRE, the difference of player i's expected utility can be specified as the difference in expected utility under QRE restrictions plus a bias term: $E\tilde{\pi}_i[m_i, p_{-i}(\mathbf{m})] + \gamma_i(\mathbf{m})$. The bias term, $\gamma_i(\mathbf{m})$, measures the departure of player i's behavior from QRE in terms of utility.[15] We established that QRE restrictions hold if and only if $\gamma_i(\mathbf{m})$ equals zero. In our test, we consider

[15]Consider any choice probability function $p_i(\mathbf{m})$ not necessarily consistent with QRE. This function is equivalent to an expected utility difference $F_i^{-1}[p_i(\mathbf{m})]$ due to the fact that $p_i = F_i[F_i^{-1}(p_i)]$. Consequently, the bias term is precisely $\gamma_i(\mathbf{m}) = F_i^{-1}[p_i(\mathbf{m})] - E\tilde{\pi}_i[m_i, p_{-i}(\mathbf{m})]$.

a linear bias term function,

$$\gamma_i(\mathbf{m}) = \gamma_{i,0} + \gamma_{i,1}m_i + \gamma_{i,2}m_{-i}. \tag{13}$$

which is identified by variation as described by the application of Propositions 2 and 3.[16]

Player i's choice probability without QRE restrictions is then:

$$p_i^{Non-QRE}(\mathbf{m}) = \frac{\exp\left\{\sum_{l=1}^{L}\beta_{i,l}\cdot\left(E\tilde{\pi}_i[m_i,p_{-i}(\mathbf{m})]+\gamma_i(\mathbf{m})\right)^l\right\}}{1+\exp\left\{\sum_{l=1}^{L}\beta_{i,l}\cdot\left(E\tilde{\pi}_i[m_i,p_{-i}(\mathbf{m})]+\gamma_i(\mathbf{m})\right)^l\right\}}. \tag{14}$$

The model primitives can also be consistently estimated by MLE:

$$\begin{aligned} & LL_i^{Non-QRE} \\ &= \max_{\sigma_i,\beta_i,\gamma_i}\sum_{t=1}^{T}\left\{\mathbb{1}(a_{i,t}=1)\log[p_i^{Non-QRE}(\mathbf{m}_t)]+\mathbb{1}(a_{i,t}=0)\log[1-p_i^{Non-QRE}(\mathbf{m}_t)]\right\}. \end{aligned} \tag{15}$$

The model described by Equation (14) nests QRE restrictions (i.e., Equation (11)) as a special case. Therefore, we can test QRE by the following likelihood ratio test on the hypothesis that the vector $\gamma_i = 0$:

$$\lambda_i = 2(LL_i^{Non-QRE} - LL_i^{QRE}). \tag{16}$$

The likelihood ratio statistic, λ_i, naturally lends itself to test the QRE restriction for player i. It follows an asymptotic Chi-squared distribution with degrees of freedom given by the number of restrictions on $\gamma_{i,0}, \gamma_{i,1}$, and $\gamma_{i,2}$. To test whether QRE holds for both

[16]Recall that – as shown in Figure 3 – imposing QRE restrictions *only* on the black and blue lines suffices to identify the utility and distribution functions. Naturally, any point or realization of (m_i, m_{-i}) that lies outside of these lines provides information about $\gamma_i(\cdot)$. Given the linear specification by Equation (13), this information identifies the coefficient on m_{-i} (i.e., $\gamma_{i,2}$). Further, the specification of $u_i(m)$ that defines utility over the space of monetary reward only will identify the coefficient on m_i (i.e., γ_i). Finally, under a symmetric distribution, the median restriction – $Median(\tilde{\varepsilon}_i) = 0$ – identifies the constant $\gamma_{i,0}$. In contrast, this constant cannot be separately identified from the median of $\tilde{\varepsilon}_i$ when it is asymmetrically distributed. Therefore, under asymmetric case, we test whether $(\gamma_{i,1}, \gamma_{i,2})$ jointly equal zero; instead whether $(\gamma_{i,0}, \gamma_{i,1}, \gamma_{i,2})$ are all zeros – as in the symmetric case.

players, one can simply consider the statistic $\lambda = \lambda_i + \lambda_{-i}$ and double the degrees of freedom.

4.2.3 Monte-Carlo Results

Our Monte-Carlo exercise studies *four* specifications. The first assumes that the analyst knows the true utility and distribution functions (labelled as "Known Utility & Error"). It inserts these true functions into the estimation procedure and tests whether the vector $\gamma = 0$. Obviously, this model is not feasible in an actual dataset, but it serves as a natural benchmark for the comparison of other specifications. The second model permits the analyst to be unaware of both the utility function and the distribution function (labelled as "Unknown Utility & Error"). It tests QRE according to the procedure described in Equations (11) to (16).

The remaining two specifications study the consequences when either the utility function or the distribution function is mis-specified. The third assumes that the analyst knows the true distribution function but mis-specifies the utility function to take the form $u_i(m) = m$ (labelled as "Linear Utility & Known Error"). In contrast, the fourth specification assumes that the analyst knows the utility function but mistakenly imposes the Logit formula for each player's choice probability (labelled as "Known Utility & Logit Error).

Data Generated By QRE Behavior Table 2 presents the rejection rates of our test when the data are generated by QRE behavior, and consequently represents the type-1 error. The rejection rates are calculated based on 1,000 Monte Carlo datasets.

The testing procedure suggested in this paper ("Unknown Utility & Error") achieves a rejection rate that is extremely close to the pre-specified significance level, for both symmetric and asymmetric distributions and any sample size. However, if the analyst imposes risk neurality on a risk averse participant ("Linear Utility & Known Error"), the true hypothesis of QRE is substantially over-rejected irrespective of the distribution

Table 2: Rejection Rates of the Test of QRE when Data Are Generated by QRE Behavior

	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 & Error	10.2%	5.7%	1.6%	9.9%	4.5%	1.2%
Unknown Utility & Error	9.1%	4.8%	1.3%	14.2%	8.5%	1.8%
Linear Utility & Known Error	61.4%	48.9%	24.2%	78.8%	68.6%	45.2%
Known Utility & Logit Error	12.2%	7.2%	2.1%	65.6%	52.3%	27.8%
$T = 1000$						
Known Utility & Error	11.2%	5.6%	1.8%	11.2%	6%	0.8%
Unknown Utility & Error	9.3%	4.7%	0.5%	9.1%	5.6%	1%
Linear Utility & Known Error	100%	99.7%	98.8%	100%	100%	99.9%
Known Utility & Logit Error	15.1%	7.9%	1.9%	100%	99.9%	99.9%
$T = 8000$						
Known Utility & Error	9.7%	6.1%	1.3%	9.4%	4.3%	0.4%
Unknown Utility & Error	12.7%	5.7%	1.1%	13.5%	6.7%	1.6%
Linear Utility & Known Error	100%	100%	100%	100%	100%	100%
Known Utility & Logit Error	38.2%	27.3%	11.5%	100%	100%	100%

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

and regardless of the sample size. In sample sizes typical of economics experiments ($T = 200$), rejection rates are near 50%. In larger samples, we observe extremely high rejection rates that are close to 100%.

If the analyst imposes the true utility function but incorrectly specifies the restriction of the Logit formula ("Known Utility & Logit Error"), the type-1 error rate depends on the symmetry of $\tilde{\varepsilon}_i$'s distribution. If $\tilde{\varepsilon}_i$ is symmetrically distributed, the Logit formula is mis-specified but imposes the correct symmetry condition. With a moderate sample size $T \in \{200, 1000\}$, this correct shape restriction yields a rejection rate somewhat higher than the pre-specified significance level. The rejection rate is substantially higher only in an unusually large dataset ($T = 8000$). In contrast, when the distribution of $\tilde{\varepsilon}_i$ is asymmetric, the Logit formula imposes an incorrect shape restriction and leads to considerable over-rejection of QRE regardless of sample size.

Data Generated By Non-QRE Behavior: Iterative Reasoning Table 3 presents rejection rates when the data are generated by Level-k model instead of by QRE behavior.

It consequently studies whether our procedure has the power to reject incorrect hypothesis, i.e., type-2 error.

Table 3: Rejection Rates of the Test of QRE when Data Are Generated by Level-k

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 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	86%	81%	67.9%	54.4%	49.3%	43.8%
Level-2 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	99.8%	99.7%	99.2%	82.8%	75.9%	58.5%
Level-3 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	99.1%	99%	98.9%	99.4%	99.2%	98.6%
Panel B: $T = 1000$						
	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 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	99.5%	99.4%	99.4%	84.9%	80.8%	72.8%
Level-2 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	100%	100%	100%	100%	99.9%	99.7%
Level-3 Behavior						
Known Utility & Error	100%	100%	100%	100%	100%	100%
Unknown Utility & Error	100%	100%	100%	99.9%	99.9%	99.9%

When player types are sophisticated (i.e., level-2 and above), the test obtains a rejection rate of almost 100% for any error distribution and any sample size. This suggests that the proposed testing procedure possesses the power to reject an incorrect null hypothesis in typical sample sizes. However when players are the level-1 type, the rejection rates tend to be lower than in the "Known Utility & Error" benchmark, dropping to around 50% with an asymmetric error distribution and a small sample size $T = 200$. This arises because the "level-1" type player i believes with certainty that player $-i$ is the level-0 type who randomizes uniformly over the set of actions. According to this belief, player $-i$'s choice probability remains constant across all realizations of (m_i, m_{-i}). In turn, the Quantal response to such belief implies that player i reacts only to their own control variable m_i and does not respond to the other player's m_{-i}. Equivalently, player i's choice probability function reduces from $p_i(m_i, m_{-i})$ to $p_i(m_i)$. Therefore, the data lose one important dimension of variation provided by m_{-i}. As shown in Propositions 1 to 3, it is indeed the variation of both m_i and m_{-i} that identifies the model and tests QRE. Naturally, due to the disappearance of m_{-i} in $p_i(\cdot)$, our test obtains a lower rejection rate when data are generated by the level-1 type.[17]

5 Empirical Application: An Experimental Study

For our empirical application, we focus on the matching pennies game as presented in Table 1. Our experiment maintains the same structure as Goeree and Holt (2001). Using the data from Goeree and Holt (2001), Aguirregabiria and Xie (2021) do not reject QRE at the population level for the row player. Moreover, in a generalized 3×3 matching pennies game, Melo et al. (2019) do reject QRE at the population level, but cannot reject QRE at the individual level for more than 50% of participants.

5.1 Experimental Design

Our design requires sufficient variation in the monetary payoff magnitudes while holding the error distribution constant. As described in our Monte-Carlo exercise, we exogenously varied two parameters, m_1 and m_2, that directly enter the utility function, one for each player. These parameters were unique combinations drawn from a discrete set of 20 values, $M = \{10, 12, 14, \cdots, 48\}$. 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

[17] In principle, one can still distinguish between data that are consistent with level-1 or QRE. In particular, as described above, level-1 type behavior restricts $p_i(\cdot)$ to be independent of m_{-i} while QRE requires $p_i(\cdot)$ to depend simultaneously on both m_i and m_{-i}.

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 $|m_1| \times |m_2| = 20 \times 20 = 400$ unique parameter combinations. This approach enables the analyst to concentrate on the relevant range of monetary payoff combinations in order to estimate the choice probability function precisely.

Figure 5 visualizes the experimental implementation of the bimatrix matching pennies game, where the parameters 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.

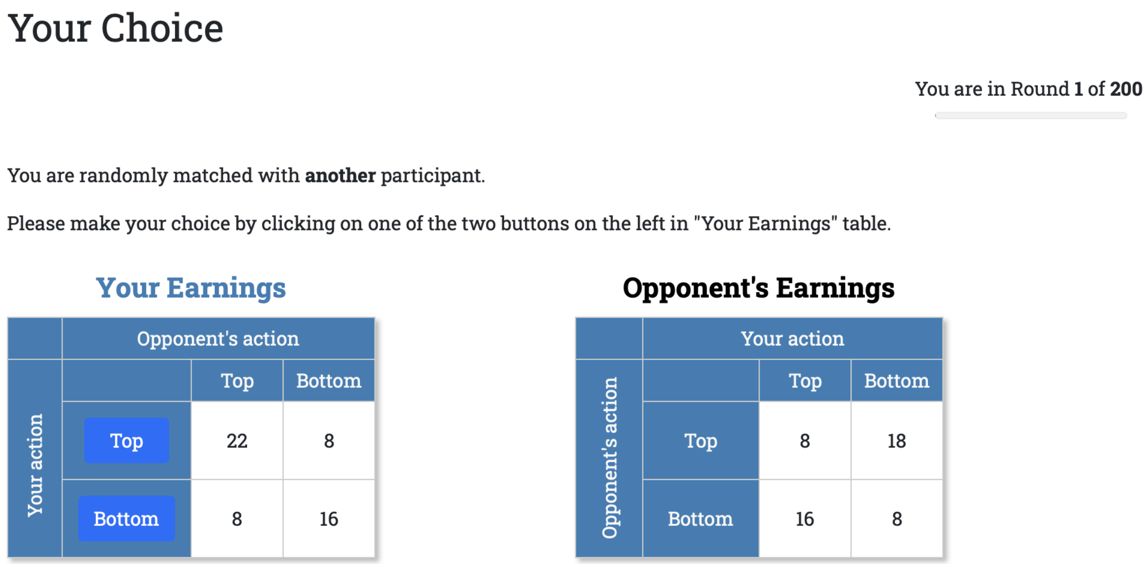

Figure 5: Matching Pennies – Experimental Implementation

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.[18]

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 $\approx$ \$20.50, including a show-up payment of €5, in a session that typically lasted around 70 minutes.

5.2 Experimental Data and Results

We first present the summary statistics of our raw data in Table 4. Since our experiment consists of only one symmetric game (i.e., $m_1 = m_2 = 16$) and 399 asymmetric games, we observe asymmetry in choice probabilities across actions.

These choice probabilities do respond to monetary payoffs. Table 5 reports the estimated coefficients of the reduced form Logit regression of player i's choice probability of

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

Table 4: Summary Statistics

	Mean	Std. Dev.	Min	Max
Player 1's Choice of Action 1	0.365	0.481	0	1
Player 2's Choice of Action 1	0.539	0.499	0	1
m_i	29	11.533	10	48
Observations	20,000 = 100 (Participants) $\times$ 200 (Decisions Per Participant)			

action 1 on m_1 and m_2 .[19] As would be expected, an increase of m_i strictly increases the expected utility of action 0 for player *i*, holding the other player's choice probability constant. Consequently, the rise of m_i reduces player *i*'s choice probability of action 1; i.e., $p_i(m_i, m_{-i})$. The above channel is referred to as the *own-utility effect* and is prevalent in experimental studies of matching pennies games (Ochs, 1995; Goeree et al., 2003). Such an effect is also salient and highly significant in our dataset. However, if a player knows that the other player has an own-utility effect, the structure of the matching pennies game then implies that player 1's choice probability of action 1 decreases in m_2 (while player 2's probability increases in m_1). Table 5 shows that such *effect of other-utility* is also sizable and statistically significant.

Table 5: Reduced Form Logit Regression of Player *i*'s Choice Probability Function

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

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

To test QRE, we first obtain a non-parametric estimate of player *i*'s choice probability, using the specification outlined in Equation (10) and MLE. Given the first-step estimate

[19]Note that we label action 1 based on Table 1. This action 1 corresponds to the choice of "Bottom" as presented in Figure 5.

$\hat{p}_i(\cdot)$, we then conduct the test of QRE in two different scenarios.

Imposing Homogeneity We begin with the *homogeneous* case where we assume that all participants in the experiment share the same utility function and error distribution. In this case, one can interpret all participants with the same player role as a single participant or player who makes $T = 50 \times 200 = 10,000$ decisions, rather than 50 different participants making individually $T = 200$ choices. With this in mind, we then conduct the test for each player i following the procedure described in Equations (11) to (16). Similar to the Monte-Carlo experiment discussed in Section 4, the test considers a parametric utility function as in Equation (8) and non-parametrically approximates the error distribution.

Table 6 presents the test results for four different specifications.

Table 6: Chi-Square Statistic and p-Value of the Test of QRE (Population Level)

	Player 1	Player 2
Linear Utility & Logit Error	$\chi^2 = 126.62$ $p < 0.0001$	$\chi^2 = 665.03$ $p < 0.0001$
Unknown Utility & Logit Error	$\chi^2 = 73.79$ $p < 0.0001$	$\chi^2 = 216.01$ $p < 0.0001$
Linear Utility & Unknown Error	$\chi^2 = 68.34$ $p < 0.0001$	$\chi^2 = 12.17$ $p = 0.0023$
Unknown Utility & Error	$\chi^2 = 7.01$ $p = 0.0300$	$\chi^2 = 26.71$ $p < 0.0001$

The first specification assumes that utility is known and is given by monetary reward so that participants are risk neutral. Moreover, the choice probability takes the usual Logit formula. Accordingly, this specification is labelled as "Linear Utility & Logit Error." The second and third specifications only impose one of the two restrictions and specify the other function 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 & Error”).

When each individual has a heterogeneous error distribution but shares the same utility function, QRE can be described by a representative player whose error distribution is non-parametrically specified (Golman, 2011). Consequently, our “Linear Utility & Unknown Error” and “Unknown Utility & Error” specifications can be interpreted as a test of such heterogeneous QRE at the population level and with the corresponding utility specification. Note that the above two specifications nest the *heterogeneous Logit QRE* (Brian et al., 2009; Golman, 2012) as a special case.

Whenever the homogeneity restriction is applied, the null hypothesis of QRE is rejected in our data under any specification. However, Table 6 also delivers an important message: When less restrictions are imposed on the utility and the distribution functions, QRE becomes more difficult to reject. This is reflected in the decreasing test statistic.

Importantly, the rejection of QRE presented in Table 6 should not be mistaken for the common finding in the literature that QRE generally fits data well (Camerer, 2003; Crawford et al., 2013). The literature evaluates QRE based on whether its prediction is *closer* to the true choice probability, as compared with other alternatives. In contrast, our test studies whether QRE *perfectly* matches the actual $p_i(\cdot)$. Consequently, it is possible that QRE is rejected by our test but still has a relatively good fit of the data. Figure 6 plots the normalized value of the log-likelihood which measures how well a model matches the sample data compared to a random guess, where 100% suggests a perfect fit.[20]

Even the simplest QRE (i.e., with Logit error and linear utility) fits the data much better than NE *with non-linear utility*. In particular, model fit substantially increases when the error distribution is non-parametrically specified. This specification can be viewed as an aggregate fit of QRE with heterogeneous error distributions (Golman, 2011). Our non-

[20]The scale is $\frac{LL^{Model} - LL^{Random}}{LL^{Sample} - LL^{Random}}$, where LL^{Random} indicates the log-likelihood value when each action is assumed to be chosen with equal probability. It thus represents the lower bound that any model should beat. LL^{Sample} is calculated using the choice probability $\hat{p}_i(\mathbf{m}_i, \mathbf{m}_{-i})$ estimated from the sample. By construction, it represents the maximum value of log-likelihood that any model could reach.

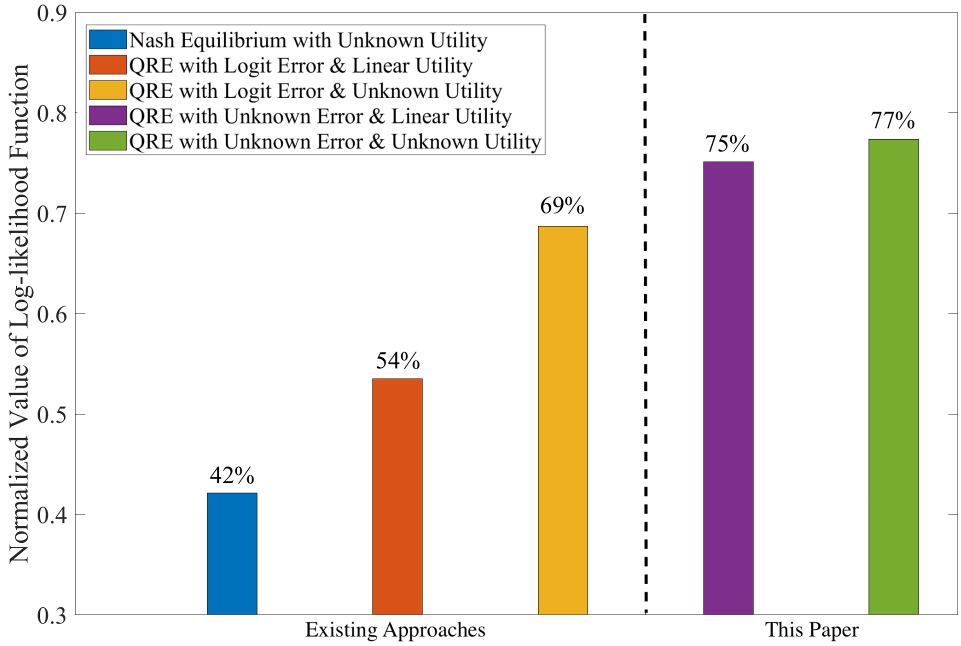

Figure 6: Model Fitness

parametric identification allow us to estimate such heterogeneous QRE at the population level. This is indicated by the two models with non-parametric error distributions, as presented by the last two bars on the right in Figure 6. As a benchmark, existing approaches which impose a distributional assumption perform worse (three bars on left) (Goeree et al., 2003; Aguirregabiria and Xie, 2021). Finally, Figure 7 plots the same measure for an out-of-sample prediction that estimates model parameters for 50% of participants and predict on the remaining participants. The two specifications with non-parametric error distributions do indeed achieve the best of out-of-sample fit.

Permitting Heterogeneity Our second analysis allows the utility function and error distribution of the participants in our experiment to be *heterogeneous*. In particular, we perform the test outlined in Equations (11) to (16) for every participant n. Formally, we allow $u_n(\cdot) \neq u_{n'}(\cdot)$ and $F_n(\cdot) \neq F_{n'}(\cdot)$ for $n \neq n'$. In line with the interpretation in Melo et al. (2019), the above process essentially tests the following null hypothesis: Participant n features Quantal response behavior with respect to other players' actual choice probabilities. If this hypothesis holds for every participant, one interpretation of this re-

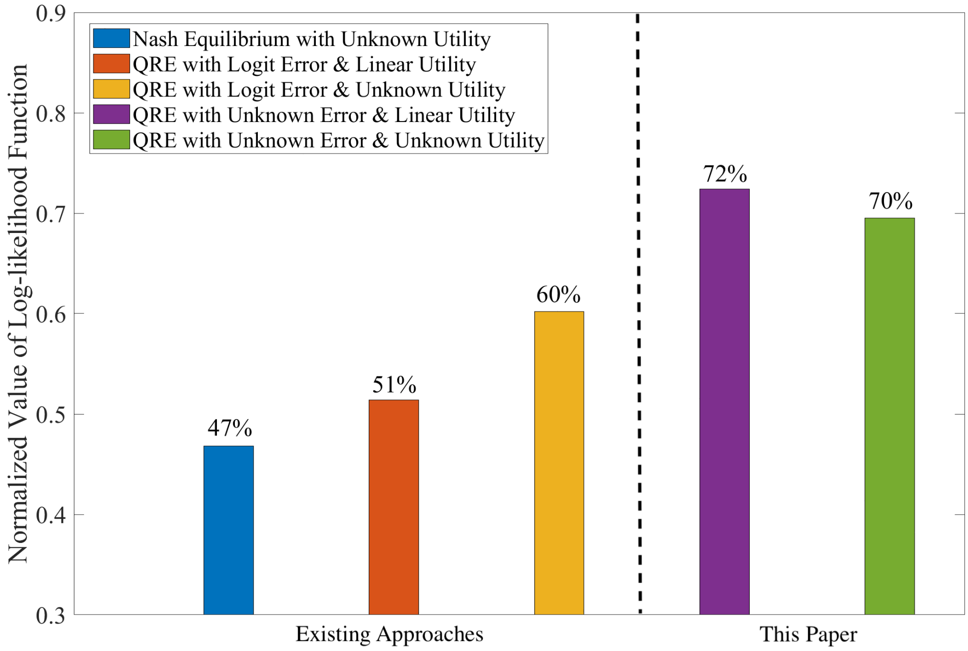

Figure 7: Out of Sample Fitness

sult is that the *heterogeneous QRE* holds in the data. In addition to the representative player interpretation of heterogeneous QRE offered by Golman (2011), this second exercise estimates and tests QRE while allowing heterogeneity in the utility function at the individual level.

Figure 8 plots the empirical cumulative distribution function (C.D.F.) for the p-value of the test statistic. It also adds a vertical line that represents the threshold of the 5% significance level. Therefore, the intersection of the empirical C.D.F. and the vertical line represents the fraction of participants whose Quantal response behavior is rejected at the 5% level. Similar to the results under the homogeneity assumption, the test delivers a general message: The less restrictions imposed on the utility and the error distribution, the more likely it is that QRE holds in the data. In more detail, with a risk neutral utility function and a logistically distributed error, the Quantal response behavior is rejected for almost 90% of participants. When the restriction is imposed for only one of the two model primitives, the null hypothesis is rejected for more than 60% of participants. In contrast, with unknown and flexible specifications of both functions, the Quantal response behavior is rejected for less than 40% of participants.

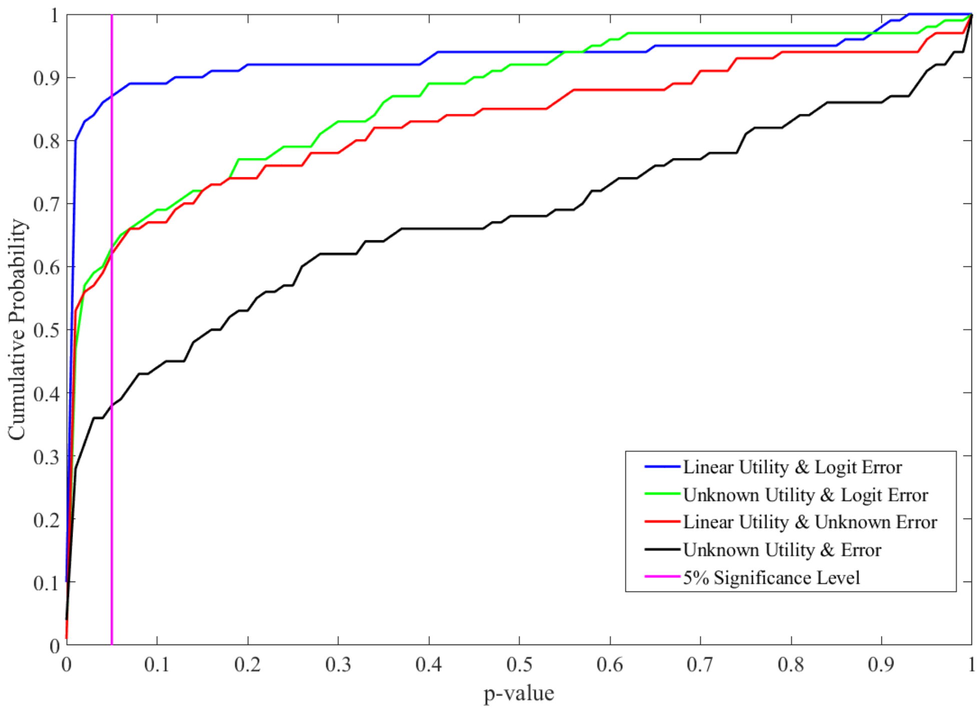

Figure 8: Empirical C.D.F. of the p-Value of the Test of QRE at the Participant Level

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

6 Conclusion

This paper studies the falsifiability and identification of QRE when both utility and the error distribution are non-parametric functions unknown to the analyst. Making use of cross-game variation, we first derive a testable implication of QRE. We then show that both the utility function and the distribution function of the error are non-parametrically over-identified under the hypothesis of QRE. Such an over-identification result directly implies a testing procedure that is straightforward to implement. Our Monte-Carlo experiment illustrates that the test has sufficient power to reject a false hypothesis. Power can be improved by assuming a functional form for utility which allows risk-aversion to vary.

We apply our results to an experimental study of the matching pennies games. With a flexible and heterogeneous specification of the utility and error distribution, the Quantal response hypothesis cannot be rejected for a majority of participants. However, it is highly rejected when strong assumptions on functional form or homogeneity are imposed.

Our identification results have an important implication for estimating QRE and predicting behavior. Our non-parametric error specification can be viewed as a population level fit of QRE with heterogeneous error distributions across participants (Golman, 2011). Previous approaches have not exploited this interpretation because of a lack of identification results. The identification results in this paper provide a means to fit heterogeneous QRE at the population level. We find that QRE with an unknown error distribution fits the data substantially better than previous methods, both in-sample and out-of-sample. This suggests that there is substantial heterogeneity in error distributions in our participant sample.

Our approach that exploits cross-game variation does have an obvious limitation. In a series of games, we can fit and test the null hypothesis that QRE holds jointly in this group of games. However, it is impossible to test whether QRE holds in any single game (i.e., for a particular realization of $(\mathbf{m}_i, \mathbf{m}_{-i})$).

Finally, while QRE has been widely applied in the experimental literature, it is far less exploited in field data. One key reason is that monetary payoffs are usually unobserved. By specifying the utility as an unknown non-parametric function, this paper opens up the possibility to apply QRE to field data. We commend empirical applications of QRE to such data for future research.

References

Aguirregabiria, Victor and Arvind Magesan, "Identification and Estimation of Dynamic Discrete Games when Players' Beliefs are Not in Equilibrium," *The Review of Economic Studies*, 2020, *87*, 582–625.

_ **and Erhao Xie**, "Identification of Non-Equilibrium Beliefs in Games of Incomplete Information Using Experimental Data," *Journal of Econometric Methods*, 2021, *10*, 1–26.

_ **and Jihye Jeon**, "Firms' Beliefs and Learning: Models, Identification, and Empirical Evidence," *Review of Industrial Organization*, 2020, *56*, 203–235.

Allen, Roy and John Rehbeck, "A Generalization of Quantal Response Equilibrium via Perturbed Utility," *Games*, 2021, *12* (1).

Bajari, Patrick, Han Hong, John Krainer, and Denis Nekipelov, "Estimating Static Models of Strategic Interactions," *Journal of Business & Economic Statistics*, 2010, *28*, 469–482.

Brian, Rogers W., Thomas R. Palfrey, and Colin F. Camerer, "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," *Journal of Economic Theory*, 2009, *144* (4), 1440–1467.

Camba-Mendez, Gonzalo and George Kapetanios, "Statistical Tests and Estimators of the Rank of a Matrix and Their Applications in Econometric Modelling," *Econometric Reviews*, 2009, *28* (6), 581–611.

Camerer, Colin F., *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, 2003.

Chen, Daniel, Martin Schonger, and Chris Wickens, "oTree-An Open-Source Platform for Laboratory, Online, and Field Experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Chen, Xiaohong, "Large Sample Sieve Estimation of Semi-Nonparametric Models," in James J. Heckman and Edward E. Leamer, eds., *Handbook of Econometrics*, Vol. 6, Elsevier, 2007, pp. 5549–5632.

Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 2013, *51* (1), 5–62.

Friedman, Evan, "Endogenous Quantal Response Equilibrium," *Games and Economic Behavior*, 2020, *124*, 620–643.

_ **and Felix Mauersberger**, "Quantal Response Equilibrium with Symmetry: Representation and Applications," Technical Report, Paris School of Economics 2022.

Goeree, Jacob and Charles Holt, "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 2001, *91*, 1402–1422.

_ , _ , and Thomas Palfrey, "Risk Averse Behavior in Generalized Mathcing Pennies Games," *Games and Economic Behavior*, 2003, *45* (1), 97–113.

_ , _ , and _ , "Regular Quantal Response Equilibrium," *Experimental Economics*, 2005, *8*, 347–367.

_ , _ , and _ , "Stochastic Game Theory for Social Science: A Primer on Quantal Response Equilibrium," *Handbook of Experimental Game Theory*, 2020, pp. 8–47.

_ , _ , Philippos Louis, Thomas Palfrey, and Brian Rogers, "Rank-Dependent Choice Equilibrium: A Non-parametric Generalization of QRE," *Handbook of Research Methods and Applications in Experimental Economics*, 2019, pp. 8–47.

Goldfarb, Avi and Mo Xiao, "Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets," *American Economic Review*, 2011, *101* (7), 3130–3161.

Golman, Russell, "Quantal Response Equilibria with Heterogeneous Agents," *Journal of Economic Theory*, 2011, *146* (5), 2013–2028.

_ , "Homogeneity Bias in Models of Discrete Choice with Bounded Rationality," *Journal of Economic Behavior & Organization*, 2012, *82* (1), 1–11.

Greiner, Ben, "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

Haile, Philip, Ali Hortacsu, and Grigory Kosenok, "On the Empirical Content of Quantal Response Equilibrium," *American Economic Review*, 2008, *98*, 180–200.

Halevy, Yoram, Johannes Hoelzemann, and Terri Kneeland, "Magic Mirror on the Wall, Who Is the Smartest One of All?," Technical Report, The University of Toronto 2023.

Harrison, G.W. and James C. Cox, *Risk Aversion in Experiments*, JAI Press Inc., 2008.

Hoelzemann, Johannes and Hongyi Li, "Coordination in the Network Minimum Game," Technical Report, The University of Toronto 2022.

Hotz, V. Joseph and Robert A. Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Kleibergen, Frank and Richard Paap, "Generalized Reduced Rank Tests Using the Singular Value Decomposition," *Journal of Econometrics*, 2006, *133* (1), 97–126.

Liu, Nianqing, Quang Vuong, and Haiqing Xu, "Rationalization and Identification of Binary Games with Correlated Types," *Journal of Econometrics*, 2017, *201* (2), 249–268.

Manski, Charles F., "Maximum Score Estimation of the Stochastic Utility Model of Choice," *Journal of Econometrics*, 1975, *3* (3), 205–228.

McKelvey, Richard and Thomas Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 1995, *10*, 6–38.

Melo, Emerson, Kirill Pogorelskiy, and Matthew Shum, "Testing the Quantal Response Hypothesis," *International Economic Review*, 2019, *60*, 53–74.

Nagel, Rosemarie, "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 1995, *85* (5), 1313–1326.

Norets, Andriy and Satoru Takahashi, "On the Surjectivity of the Mapping Between Utilities and Choice Probabilities," *Quantitative Economics*, 2013, *4* (1), 149–155.

Ochs, Jack, "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 1995, *10* (1), 202–217.

Robin, Jean-Marc and Richard J. Smith, "Tests of Rank," *Econometric Theory*, 2000, *16* (2), 151–175.

Stahl, Dale and Paul Wilson, "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior & Organization*, 1994, *25* (3), 309–327.

_ and _ , "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 1995, *10* (1), 218–254.

Train, Kenneth, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.

Xie, Erhao, "Nonparametric Identification of Incomplete Information Discrete Games with Non-Equilibrium Behaviors," Technical Report, Bank of Canada 2022.