

Hoelzemann, Johannes; Webb, Ryan; Xie, Erhao

Working Paper

Non-parametric identification and testing of quantal response equilibrium

Bank of Canada Staff Working Paper, No. 2024-24

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Hoelzemann, Johannes; Webb, Ryan; Xie, Erhao (2024) : Non-parametric identification and testing of quantal response equilibrium, Bank of Canada Staff Working Paper, No. 2024-24, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2024-24>

This Version is available at:

<https://hdl.handle.net/10419/299360>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Non-Parametric Identification and Testing of Quantal Response Equilibrium

by Johannes Hoelzemann,¹ Ryan Webb² and Erhao Xie³

¹University of Vienna

²University of Toronto

³Banking and Payments Department, Bank of Canada
erhaoxie@bankofcanada.ca



Acknowledgements

We thank Victor Aguirregabiria, Jose Apesteguia, Ben Greiner, Yoram Halevy, Sota Ichiba, 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 experimental part of 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.

Abstract

We study 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 variations 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 result further suggests a straightforward testing procedure for QRE that 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 risk neutrality, logistically distributed errors and homogeneity lead to substantially higher rejection rates.

Topics: Econometric and statistical methods; Economic models

JEL codes: C14, C57, C92

Résumé

Dans la présente étude, nous nous penchons sur la réfutabilité et l'identification de l'équilibre de réponse quantale (tout ou rien) lorsque la fonction d'utilité et la distribution des erreurs de chaque joueur sont assouplies pour devenir des fonctions non paramétriques inconnues. À partir de variations dans les choix que font les joueurs dans un éventail de jeux, nous montrons d'abord que tant la fonction d'utilité que la distribution des erreurs sont non paramétriquement suridentifiées. Ce résultat sous-entend la possibilité de tester l'équilibre de réponse quantale au moyen d'une méthode simple permettant d'atteindre le taux souhaité d'erreurs de type 1 et de maintenir un faible taux d'erreurs de type 2. Pour appliquer cette méthode, nous menons une étude expérimentale du jeu de l'appariement des sous. Nos estimations non paramétriques rejettent fermement le modèle logit habituel pour les probabilités de choix. De plus, lorsque la fonction d'utilité et la distribution des erreurs sont suffisamment souples et hétérogènes, l'hypothèse relative à la réponse quantale ne peut être rejetée pour 70 % des participants. Toutefois, les hypothèses solides, comme la neutralité face au risque, la distribution logistique des erreurs et l'homogénéité donnent lieu à des taux de rejet nettement plus élevés.

Sujets : Méthodes économétriques et statistiques ; Modèles économiques

Codes JEL : C14, C57, C92

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 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, *Quantal Response Equilibrium* (QRE; 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 the private information of each player. QRE is then defined as a fixed 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. It has successfully explained many deviations from NE and has become an important benchmark in game theory (Goeree et al., 2020).¹

In a 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 small-stakes risk aversion can be prevalent even in laboratory settings, as identified by Goeree et al. (2003) and Harrison and Cox (2008), among many others. Further, most applications also assume that each player’s random errors follow a common distribu-

¹For 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).

tion and the functional form of this distribution is known by the analyst, e.g., the Logit distribution. 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 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 has full control and 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

²In 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.

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

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 sub-

⁴This idea of exploiting cross-game variation was first conjectured by Haile et al. (2008). Another approach to address the non-falsifiability of QRE is to impose additional restrictions on the distribution of random errors across actions and derive the testable implications of QRE *within* repetitions of the same game. Two important contributions using this approach are the regular QRE by Goeree et al. (2005) and the rank-dependent choice equilibrium by Goeree et al. (2019). However, the testable implication is derived under the known utility assumption, and a formal statistical test has not been derived.

stantially 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 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 risk aversion 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

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.⁶ Finally, Xie’s testable implication is a restriction on

⁵Goeree 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.

⁶In 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

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 the existence of other-regarding preference, departures from expected-utility, 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

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.

experimental currency units.

We define $u_i(m) : \mathbb{R} \rightarrow \mathbb{R}$ 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 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 bounded. Moreover, it is strictly increasing and continuous in m .*

In this 2×2 game, let \mathbf{m}_i be a 4×1 vector that equals $[m_i(a_i = 0, a_{-i} = 0), m_i(a_i = 0, a_{-i} = 1), m_i(a_i = 1, a_{-i} = 0), m_i(a_i = 1, a_{-i} = 1)]'$. 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.⁷ 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.

This paper presents results for the analysis of an experiment that comprises a series of games with varying monetary payoffs. Specifically, let $\mathcal{M}_i(\mathbf{a}) \subset \mathbb{R}$ denote the support of $m_i(\mathbf{a})$ (i.e., the set of all possible values that $m_i(\mathbf{a})$ could take). The set $\mathcal{M}_i(\mathbf{a})$ could be either an interval or even a singleton. Moreover, $\mathcal{M}_i \subset \times_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ represents the support of player i 's own monetary rewards \mathbf{m}_i . Assumption 2 states the exogenous condition imposed on this support. It simply requires that for each player i , the monetary payoff of at least one action profile has exogenous variation conditional on \mathbf{m}_{-i} .

⁷We 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.

Assumption 2. For each player i , the outcome variables $\mathbf{m}_i \in \mathcal{M}_i$ have exogenous variations over their support conditional on $\forall \mathbf{m}_{-i} \in \mathcal{M}_{-i}$. Moreover, $\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ is an interval that could be either open or closed.

The structure of $\mathcal{M} \equiv \mathcal{M}_i \times \mathcal{M}_{-i}$ determines the type of game. Other than the exogeneity condition in Assumption 2, we do not impose additional restrictions on this structure.⁸ Therefore, our framework is applicable to a general class of games. For instance, Table 1 represents a matching pennies game where the analyst 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, so that the support $\mathcal{M}_i(\mathbf{a})$ is a singleton for these profiles. 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 1: Monetary Payoff Matrix of the Matching Pennies ($m_1 > 8, m_2 > 8$) Game

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

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 $a_{-i} = 0$, given the control variables $\mathbf{m} = (\mathbf{m}_i, \mathbf{m}_{-i})$. In strategic settings,

⁸This general structure of \mathcal{M} includes experiments that vary every action profile's payoffs across games as well as experiments that only vary some profiles' payoffs while holding the payoffs of other profiles constant. We only require $\mathcal{M}_i(\mathbf{a})$ to be an interval for at least one $\mathbf{a} \in \mathcal{A}$ so that the data contains variations of monetary payoffs. In addition, across games, player i 's monetary rewards of any two action profiles could be either independent or exhibit arbitrary correlations.

Table 2: Monetary Payoff Matrix of the Coordination ($0 \leq m_i \leq 15$) Game

		Player 2	
		0	1
Player 1	0	m_2	0
	1	m_2	15

a player's decision depends on both players' monetary rewards. Specifically, \mathbf{m}_{-i} indirectly affects player i 's choice by directly affecting the decision of player $-i$. Intuitively, when only \mathbf{m}_{-i} varies, it exogenously shifts the indirect impact while holding player i 's utility of each action profile constant. This variation is the key for our identification results. Given $p_{-i}(\mathbf{m})$, the expected utility of action a_i for player i is as follows:

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

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 $p_{-i}(\mathbf{m})$. QRE places an 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:

$$EU_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 0) \geq EU_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})] + \varepsilon_i(a_i = 1)$$

$$\Leftrightarrow \underbrace{\varepsilon_i(a_i = 1) - \varepsilon_i(a_i = 0)}_{=\tilde{\varepsilon}_i} \leq \underbrace{EU_i[\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})] - EU_i[\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m})]}_{=E\tilde{U}_i[\mathbf{m}_i, p_{-i}(\mathbf{m})]}. \quad (2)$$

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) $F_i(\tilde{\epsilon}_i)$ is continuous and strictly increasing over the real line, $\forall i$.

(b) $F_i(\tilde{\epsilon}_i)$ is independent of $(\mathbf{m}_i, \mathbf{m}_{-i})$, $\forall i$.

Assumption 3(a) is the standard regularity condition. In particular, strict monotonicity implies that the choice probability of a_i is strictly increasing in expected utility when a_i is played, known as “responsiveness” (Goeree et al., 2005). The full support condition of $\tilde{\epsilon}_i$ ensures that each action is chosen with strictly positive probability and avoids the zero-likelihood problem (Goeree et al., 2020). 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:

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

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

$$p_i(\mathbf{m}) = F_i \left[\underbrace{EU_i(\mathbf{m}_i, a_i = 0, p_{-i}(\mathbf{m})) - EU_i(\mathbf{m}_i, a_i = 1, p_{-i}(\mathbf{m}))}_{= \widetilde{EU}_i(\mathbf{m}_i, p_{-i}(\mathbf{m}))} \right], \forall i \text{ and } \mathbf{m} \in \mathcal{M}. \quad (4)$$

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

equilibria).

By Brouwer’s fixed-point theorem, any game has at least one QRE. Moreover, many games could have multiple QRE. When multiple QRE exist, Definition 1 states that there is a *deterministic* equilibrium selection mechanism that chooses one of these equilibria. 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.

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 (i.e., 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 4. *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 4 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 4 allows players to select an equilibrium in any arbitrary way and only restricts the number of

equilibrium switching points to be finite.⁹ Assumption 4 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).¹⁰ 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). This inversion expresses the difference of expected utilities for player i , which is linear in player $-i$ ’s choice probability:

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

⁹Notably, 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 4 includes this reasonable equilibrium selection mechanism as a special case.

¹⁰In 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).

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 4, these images are either an interval or a union of finite numbers of disjoint intervals. Proposition 1 describes the over-identification of $F_i(\cdot)$.

Proposition 1. *Suppose that Assumptions 1 to 4 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 $F^{-1}(p)$ 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 structures that allow 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).¹¹ 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,

¹¹For 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 mean (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.

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:

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

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 4 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:*

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

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:

Lemma 1. (Xie 2022) *Suppose that Assumptions 1 to 4 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:*

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

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.¹² This condition is difficult to construct in an actual dataset because it requires equating two estimated quantities and has yet to be implemented.¹³ 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

¹²Precisely, 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.

¹³Specifically, 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.

fully non-parametric specification. The benefit of dimension reduction is even more pronounced as the number of actions and / or players increases.¹⁴

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). Figure 1 shows simulated choice probabilities under QRE restrictions for various monetary payoffs.¹⁵ 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 (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.

Figure 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

¹⁴The 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.

¹⁵For 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.

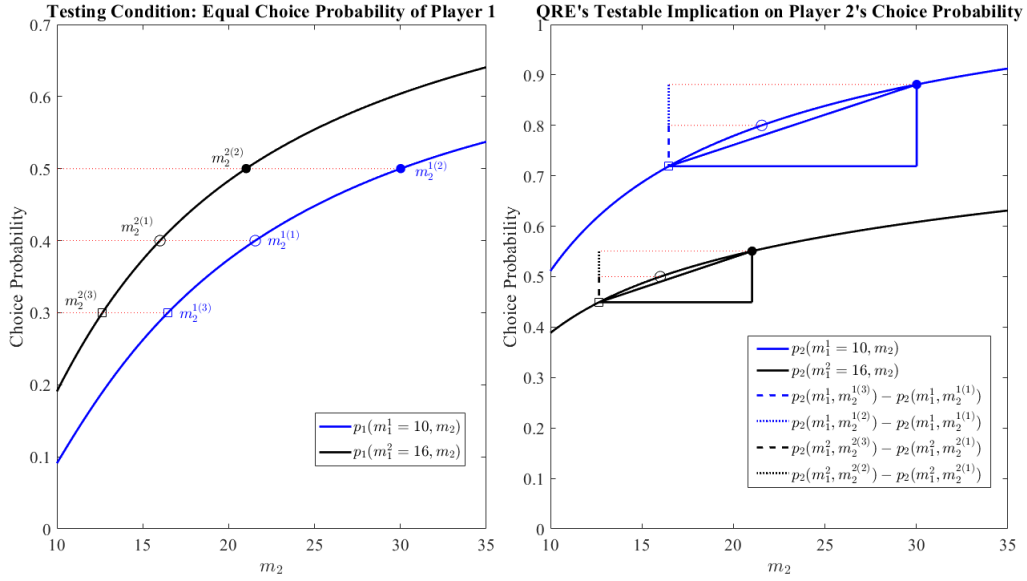


Figure 1: Testable Implication of QRE

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 Figures 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 $\tilde{\epsilon}_i$ 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 $\tilde{\epsilon}_i$.

As described in footnote 11, we shall refer to $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ as the median and the standard deviation of $\tilde{\epsilon}_i$, respectively. To identify these two values, we introduce

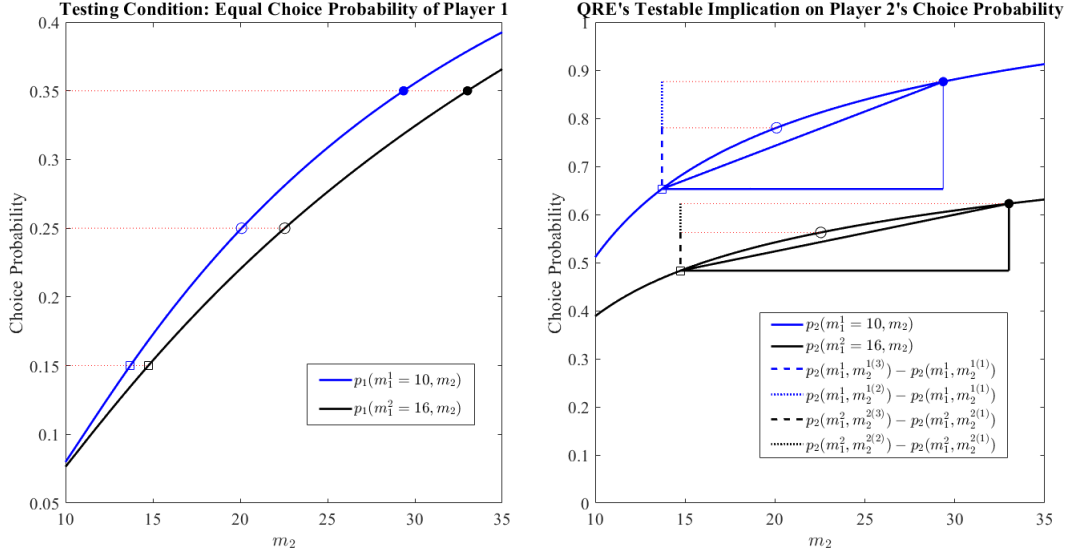


Figure 2: Violation of the Testable Implication under Non-QRE Behaviors

an innocuous scale normalization on the utility function, as described by the following assumption:

Assumption 5. 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 5 normalizes the scale of player i 's utility function. It is innocuous given that $u_i(m)$ is strictly increasing.¹⁶ 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 $F_i^{-1}(p^2)$).

Proposition 3. Suppose that Assumptions 1 to 5 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 $\tilde{\mathcal{E}}_i$ (i.e., $F_i^{-1}(p^1)$). In the appendix, we show that in any experiment, as long

¹⁶Since 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 5 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, of course, the analyst can ensure a positive denominator by relabeling each player's action.

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 6 summarizes two such structures, which are satisfied in the matching pennies game presented in Table 1.

Assumption 6. *Consider the following properties of monetary payoff matrices $\forall i$:*

(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 6(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 6(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 6 holds, the median value of $\tilde{\epsilon}_i$ can be identified, as established by the following proposition:

Proposition 4. *Suppose that Assumptions 1 to 5 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 6(a) or \mathbf{m}_i^1 and \mathbf{m}_i^2 satisfy Assumption 6(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 6 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} .*

By setting either $u_i(0) = 0$ or $u_i[\min\{\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})\}] = 0$, the identification of the difference in utilities from Lemma 2 consequently identifies the utility function $u_i(m)$ non-parametrically $\forall m \in \cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$.

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.¹⁷ 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 2.¹⁸

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

¹⁷In 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 achieves the desired performance when $T = 2000$.

¹⁸The continuity condition of m_i is approximated by a discrete uniform distribution with a small step size. Ideally, the step size should decrease as the sample size increases so that \mathcal{M} is dense in the continuous interval $[10, 48]$ as $T \rightarrow \infty$. 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 our estimator and test it using this specific step size. Second, shrinking the step size in larger samples substantially increases the computational burden.

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} :

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

The risk preference parameter ν is set to 0.6 to model risk aversion, in line 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:

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

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

Equation (10) draws $\tilde{\epsilon}_i$ 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 ($\lambda = 7.5 \cdot \sqrt{\frac{9}{35}}$). 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 $\tilde{\epsilon}_i$ is asymmetrically distributed with a higher density in the positive region. Figure 3 plots the P.D.F. of the symmetric and asymmetric distributions, alongside a comparison with the Logit specification.

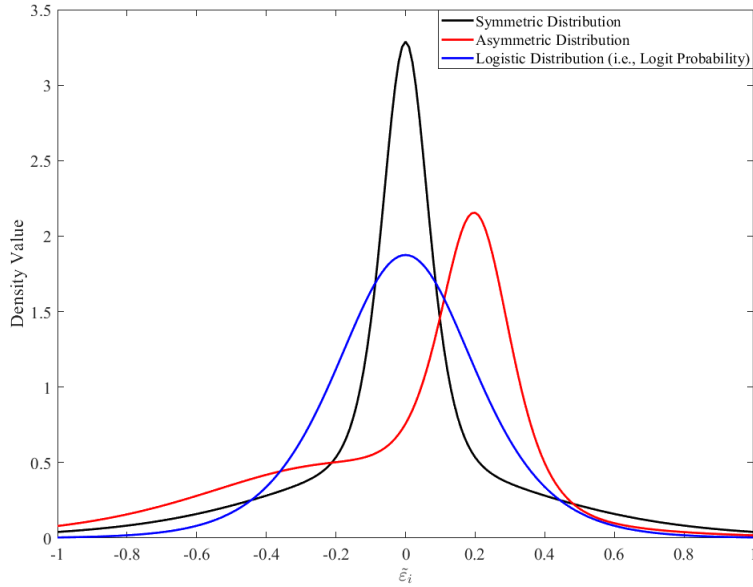


Figure 3: Probability Density Functions of the Random Perturbation $\tilde{\epsilon}_i$

We set the scale of $\tilde{\epsilon}_i$ based on the empirical estimates in Section 5. In particular, the variances of $\tilde{\epsilon}_i$ for both symmetric and asymmetric cases are set to $Var(\tilde{\epsilon}_i) = \frac{\pi^2}{3 \times (7.5^2)} \approx 0.0585$. This value of variance corresponds to $\lambda = \sqrt{\frac{\pi^2}{3Var(\tilde{\epsilon}_i)}} = 7.5$ if $\tilde{\epsilon}_i$ 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.

The scale of $\tilde{\epsilon}_i$ 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 $F_i(\tilde{\epsilon}_i)$ in our framework.¹⁹ 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 ,

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

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 $L_u \rightarrow \infty$, the Bernstein polynomial will converge uniformly to the continuous function $u(\tilde{m})$ with $\theta_l = u(\frac{l}{L_u})$.²⁰ We set $L_u = 3$ (4) when $T = 200$ (2000).

Further, we approximate the distribution function using a mixture of normal distribu-

¹⁹For some intuition, as $Var(\tilde{\epsilon}_i)$ decreases, player i tends to choose the action with a higher expected utility more frequently. When $Var(\tilde{\epsilon}_i) \rightarrow 0$, player i deterministically selects the action that maximizes the expected utility. Conversely, as $Var(\tilde{\epsilon}_i) \rightarrow \infty$, player i randomizes each action with equal probability.

²⁰In a large sample with 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.

tions:

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

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 $L_F \rightarrow \infty$, 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.

Let $\theta = (\theta^u, \theta^{Pr}, \theta^\mu, \theta^\sigma)'$ 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:

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

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:

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

$\tilde{\mathcal{M}}$ 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 $\gamma_i(\tilde{\mathbf{m}})$.

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

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:

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

The test of QRE is equivalent to testing whether $\boldsymbol{\gamma} = (\gamma_1, \gamma_2) = \mathbf{0}$. 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 $\boldsymbol{\gamma}$ (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 prob-

ability. 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:

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

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 :

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

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

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

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 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.²¹ 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.

Table 3: Rejection Rates of the Over-Identification Test of QRE (QRE Data)

Significance Level	Symmetric Distribution			Asymmetric Distribution		
	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
<i>T</i> = 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%
<i>T</i> = 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.

²¹To 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, which tests model primitives at the participant level.

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.²²

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.

Figure 4 plots the averages of the estimated utility functions and the distribution

²²The 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.

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.

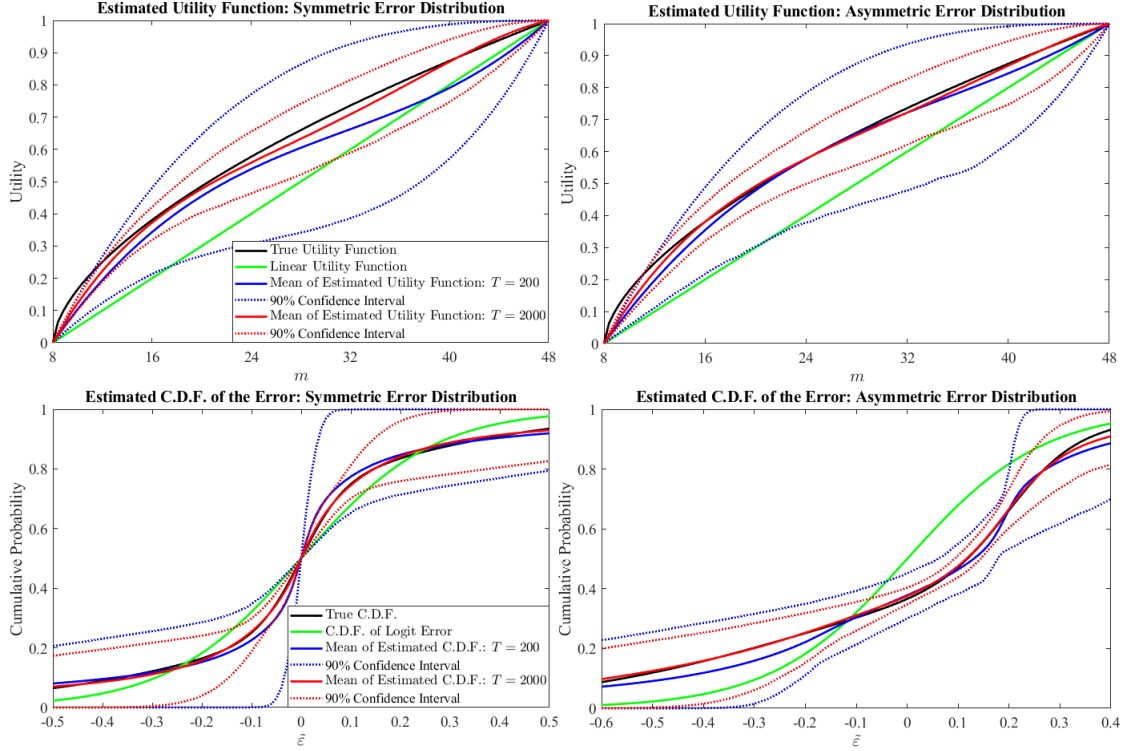


Figure 4: Estimates of the Utility Function and the Distribution Function

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 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

tion 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 \bar{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 \bar{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.

Table 4: Rejection Rates of the Over-Identification Test of QRE (Symmetric Level- k Data)

Panel A: $T = 200$						
Significance Level	Symmetric Distribution			Asymmetric Distribution		
	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
<i>Level-1 Reasoning Behavior</i>						
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%
<i>Level-2 Reasoning Behavior</i>						
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%
<i>Level-3 Reasoning Behavior</i>						
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: $T = 2000$						
Significance Level	Symmetric Distribution			Asymmetric Distribution		
	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
<i>Level-1 Reasoning Behavior</i>						
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%
<i>Level-2 Reasoning Behavior</i>						
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%
<i>Level-3 Reasoning Behavior</i>						
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%

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.

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	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
<i>Player 1 is Level-2 Reasoner & Player 2 is Level-1 Reasoner</i>						
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%
<i>Player 1 is Level-3 Reasoner & Player 2 is Level-2 Reasoner</i>						
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	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
<i>Player 1 is Level-2 Reasoner & Player 2 is Level-1 Reasoner</i>						
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%
<i>Player 1 is Level-3 Reasoner & Player 2 is Level-2 Reasoner</i>						
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%

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 QRE at the population level for the row player using the data from Goeree and Holt (2001). 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 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, $\mathcal{M} = \{10, 12, 14, \dots, 48\}$. 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 rematching 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.

Figure 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.

Your Choice

You are in Round 1 of 200

You are randomly matched with **another** participant.

Please make your choice by clicking on one of the two buttons on the left in "Your Earnings" table.

		Opponent's action	
		Top	Bottom
Your action	Top	22	8
	Bottom	8	16

		Your action	
		Top	Bottom
Opponent's action	Top	8	18
	Bottom	16	8

Figure 5: Matching Pennies Game – Experimental Implementation

To improve participants' experience and to assist in selecting an action, we implemented a highlighting tool that uses a yellow color. When a participant moves their

mouse over a row in their matrix (“Your Earnings”), the action is highlighted in yellow 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 unhighlight 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.²³

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

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

²³The interactive experimental interface can be accessed anytime upon request. Example screenshots can be found in the appendix.

the 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.

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.

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) 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.

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	$\chi^2 = 93.81$ $p < 0.0001$

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:

$$\text{Normalized Log-Likelihood} = \frac{LL^{Model} - LL^{Random}}{LL^{Sample} - LL^{Random}}, \quad (20)$$

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.

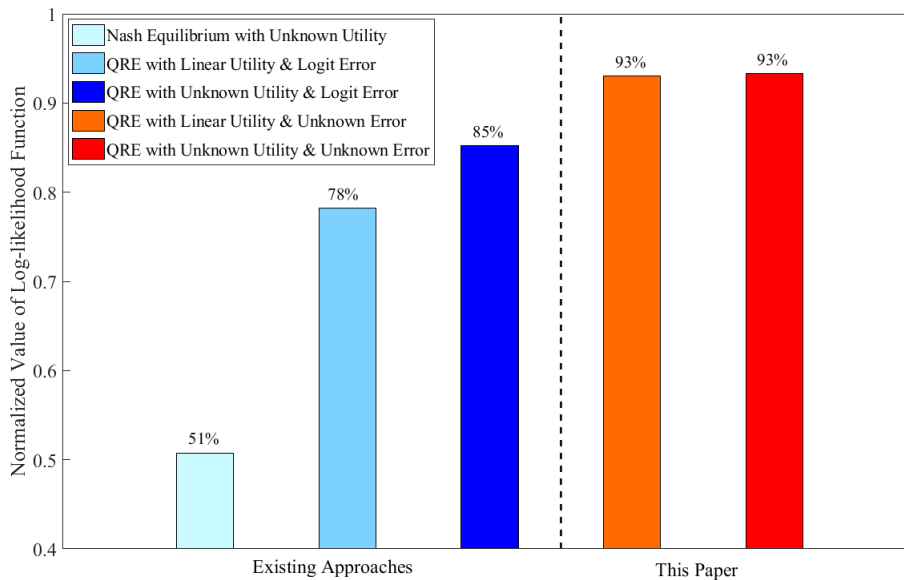


Figure 6: Model Fitness

Figure 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 Figure 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, Figure 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.

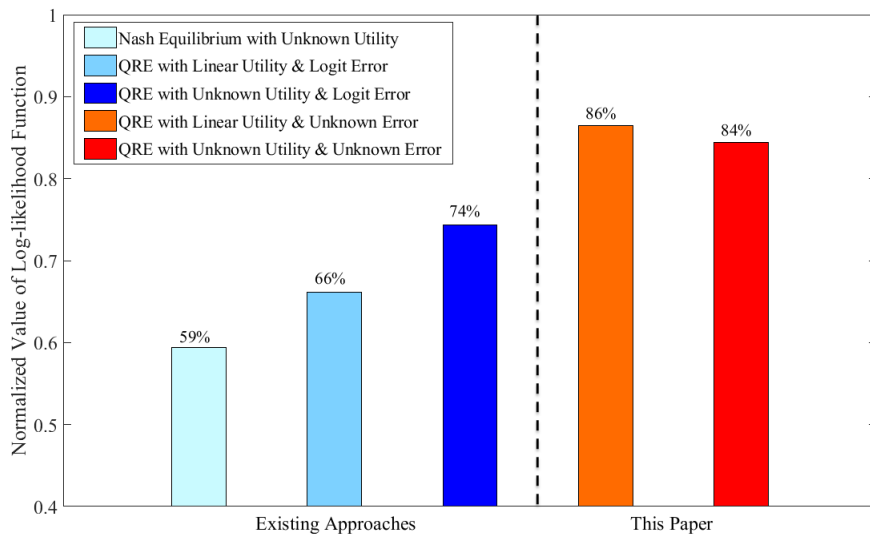


Figure 7: Out of Sample Fitness

In Figure 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 \approx \$43) do we find risk aversion to be significant at the 10% level.

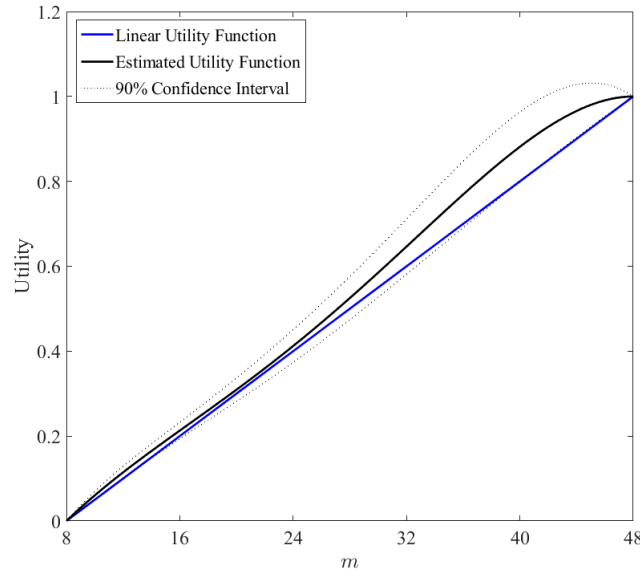


Figure 8: Estimated Utility Function

Figure 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 $\tilde{\epsilon}_i$ appears symmetric, it is not symmetric around 0. In particular, the estimate of $mean(\tilde{\epsilon}_i)$ 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 $mean(\tilde{\epsilon}_i)$ suggests that participants tend to over-estimate the reward of the action presented at the top of the screen by around €3 \approx \$3.20, which is approximately 6% of the maximum reward. Our non-parametric estimate is able to recover this sizable position effect, which

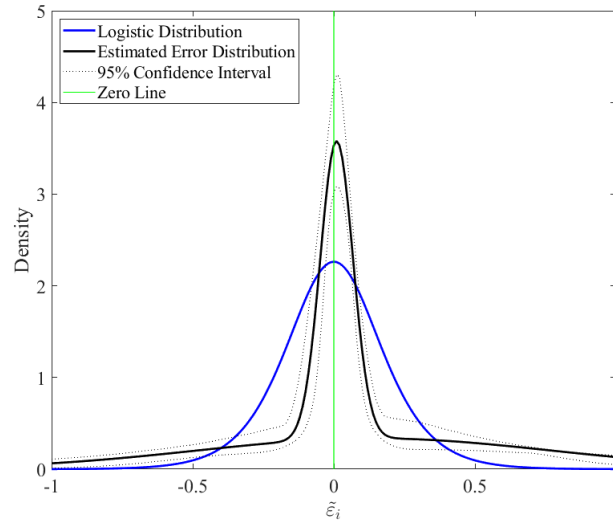


Figure 9: Estimated P.D.F. for $\tilde{\epsilon}_i$

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.

Figure 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

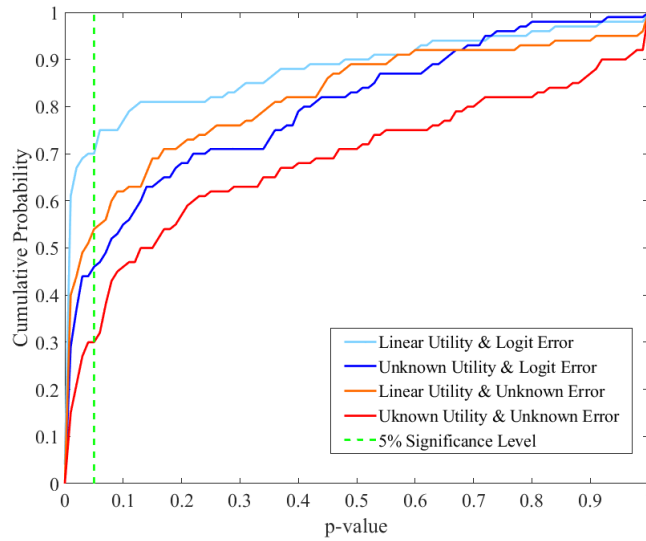


Figure 10: Empirical C.D.F. of the p -Value of the Test of Quantal Response Behaviors

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.

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 a 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.²⁴ 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.

²⁴If the analyst assumes that the marginal error distribution is invariant across games with different number of actions, then the predicted choice probability is a singleton. Without this invariance assumption on the marginal error distribution, the predicted choice probabilities would form a set, and this set is sufficiently narrow (Goeree et al., 2020).

References

- Aguirregabiria, Victor 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.
- **and Pedro Mira**, “Identification of Games of Incomplete Information with Multiple Equilibria and Unobserved Heterogeneity,” *Quantitative Economics*, 2019, 10 (4), 1659–1701.
- 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.
- Camerer, Colin F.**, *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, 2003.
- Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong**, “A Cognitive Hierarchy Model of Games,” *Quarterly Journal of Economics*, 2004, 119 (3), 861–898.
- 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.
- Compiani, Giovanni**, “Market Counterfactuals and the Specification of Multiproduct Demand: A Nonparametric Approach,” *Quantitative Economics*, 2022, 13 (2), 545–591.
- Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberry**, “Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications,” *Journal of Economic Literature*, 2013, 51 (1), 5–62.
- Delgado, Miguel A and Javier Hidalgo**, “Nonparametric Inference on Structural Breaks,” *Journal of Econometrics*, 2000, 96 (1), 113–144.
- 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 Matching 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.
- Klein, Roger W and Richard H Spady**, “An Efficient Semiparametric Estimator for Binary Response Models,” *Econometrica*, 1993, 61 (2), 387–421.
- Lewbel, Arthur**, “Semiparametric Qualitative Response Model Estimation with Unknown Heteroscedasticity or Instrumental Variables,” *Journal of Econometrics*, 2000, 97 (1), 145–177.

- Liu, Nianqing, Quang Vuong, and Haiqing Xu**, “Rationalization and Identification of Binary Games with Correlated Types,” *Journal of Econometrics*, 2017, 201 (2), 249–268.
- 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.
- Müller, Hans-Georg**, “Change-Points in Nonparametric Regression Analysis,” *The Annals of Statistics*, 1992, 20 (2), 737–761.
- 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.
- Rogers, Brian W., Thomas R. Palfrey, and Colin F. Camerer**, “Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies,” *Journal of Economic Theory*, 2009, 144 (4), 1440–1467.
- Sørensen, Jesper R.-V. and Mogens Fosgerau**, “How McFadden Met Rockafellar and Learned To Do More with Less,” *Journal of Mathematical Economics*, 2022, 100.
- 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**, “Inference in Games without Nash Equilibrium: An Application to Restaurants’ Competition in Opening Hours,” Technical Report, Bank of Canada 2018.
- , “Nonparametric Identification of Incomplete Information Discrete Games with Non-Equilibrium Behaviors,” Technical Report, Bank of Canada 2022.

Non-Parametric Identification and Testing of Quantal Response Equilibrium

Online Appendix: Proofs

Johannes Hoelzemann Ryan Webb Erhao Xie

April 22, 2024

Omitted Proofs

Proof of Proposition 1: Since $p^1, p^2 \in \mathcal{P}_i(\mathbf{m}_i^1)$, there must exist two values of \mathbf{m}_{-i} —denoted as \mathbf{m}_{-i}^1 and \mathbf{m}_{-i}^2 —such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) = p^1$ and $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^2) = p^2$. Evaluating Equation (5) at these two values leads to the following equations:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) = p^1] &= \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1), \\ F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^2) = p^2] &= \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^2). \end{aligned} \quad (21)$$

Given that $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$ are known by the analyst, the above system is a linear system with two equations and two unknowns (i.e., $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)$ and $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)$). Moreover, the fact that $p^1 \neq p^2$ implies that $F_i^{-1}(p^1) \neq F_i^{-1}(p^2)$ and therefore $p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) \neq p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^2)$. Consequently, the rank condition of the system by Equation (21) is satisfied and both $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)$ and $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)$ are point identified.

Fix \mathbf{m}_i at \mathbf{m}_i^1 and only consider the variations of \mathbf{m}_{-i} . Equation (5) then becomes:

$$F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i})] = \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}). \quad (22)$$

Since $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)$ and $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)$ have been identified and $p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i})$ is known by the analyst, Equation (22) directly identifies $F_i^{-1}(p) \forall p \in \mathcal{P}_i(\mathbf{m}_i^1)$ with the variations provided by \mathbf{m}_{-i} . This completes the proof. \square

Proof of Proposition 3: Similar to the argument in the proof of Proposition 1, there exists one value $\mathbf{m}_{-i} = \mathbf{m}_{-i}^1$ such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) = p^1$ given that $p^1 \in \mathcal{P}_i(\mathbf{m}_i^1)$. Evaluating Equation (5) at this realization $(\mathbf{m}_i^1, \mathbf{m}_{-i}^1)$ would imply the following relationship:

$$F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) = p^1] = \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1). \quad (23)$$

Since $F_i^{-1}(p^1)$ and $p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1)$ are known to the analyst and $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)$ is normalized to one, Equation (23) contains only one unknown $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)$. Consequently, this utility difference is identified. Given the identification of the utility differences, Equation (22) then identifies $F_i^{-1}(p) \forall p \in \mathcal{P}_i(\mathbf{m}_i^1)$ due to the exogenous variation of \mathbf{m}_{-i} . This completes the proof. \square

Proof of Proposition 4: To prove this proposition, it is suffice to prove that $F_i^{-1}(p^1)$ is identified at only one value p^1 . The identification of $F_i^{-1}(p) \forall p \neq p^1$ simply follows Proposition 3.

First consider Assumption 6(a) so that $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) = -\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)$. Plugging this relationship into Equation (23), one could obtain the following equation:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1) = p^1] &= [1 - 2p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1)] \cdot \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) \\ \Rightarrow F_i^{-1}(p^1) &= 1 - 2p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1). \end{aligned} \quad (24)$$

The second line identifies the value of $F_i^{-1}(p^1)$ and is the result of the normalization by Assumption 5 such that $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) = 1$.

Next, suppose instead that Assumption 6(b) holds. We prove the case that $m_i^1(a_i, a_{-i}) = m_i^2(1 - a_i, a'_{-i}) \forall a_i$ and for $a_{-i} = a'_{-i} = 1$. Therefore, we have $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) = -\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} =$

1). The proofs for the other two cases (i.e., $a_{-i} \neq a'_{-i}$ and $a_{-i} = a'_{-i} = 0$) follow a similar argument and are suppressed.

Let us consider $p^1 \in \mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$. As described above, there must exist two games—denoted as $(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})$ and $(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})$ —such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}) = p^1$. When we evaluate these two games, Equation (5) then turns to:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) = p^1] &= \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) \\ F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}) = p^1] &= \tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}) \\ &= -\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 0) + \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}). \end{aligned} \quad (25)$$

The last line of Equation (25) follows from the result that $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) = -\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1)$.

1). Solving Equation (25), one could identify $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) = \frac{F_i^{-1}(p^1) - 1}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} + 1$ and $\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 0) = \frac{F_i^{-1}(p^1) + 1}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})} - 1$. Next, consider another two games—denoted as $(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})$ and $(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})$ —such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) = p^2 \in \mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$. Since $\mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$ includes an interval, we could always find such $p^2 \neq p^1$.

Evaluating Equation (5) at the above two realizations implies the following equation:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) = p^2] &= \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) \\ F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) = p^2] &= -\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 0) + \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}). \end{aligned} \quad (26)$$

Since the terms on the left-hand side of the above two equations are equal, we could equate them and plug in the identified values of $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0)$ and $\tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 0)$.

This transformation then identifies the value of $F_i^{-1}(p^1)$ as the following:

$$F_i^{-1}(p^1) = \frac{2 - \left[\frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} + \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})} \right]}{\frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})} - \frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}}. \quad (27)$$

It can be shown that the denominator of Equation (27) equals $\frac{F_i^{-1}(p^2)+1}{F_i^{-1}(p^1)+1} - \frac{F_i^{-1}(p^2)-1}{F_i^{-1}(p^1)-1}$.
Therefore, this denominator is different than zero provided that $F_i^{-1}(p^1) \neq F_i^{-1}(p^2)$.
Equation (27) then identifies $F_i^{-1}(p^1)$ and completes the proof. \square

Non-Parametric Identification and Testing of Quantal Response Equilibrium

Online Appendix: Additional Testable Implications

Johannes Hoelzemann Ryan Webb Erhao Xie

April 22, 2024

In Proposition 5, we first show that our over-identification test includes all the testable implications derived by Xie (2022).

Proposition 5. *Suppose that Assumptions 1 to 4 hold. If Equation (7) is satisfied, then the QRE restrictions in Equation (8) hold for any three pairs of games such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = (\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ for $l = 1, 2, 3$ and $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) \neq p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})$.*

Proof. Recall the definition of Equation (6). Consider two distinct realizations of \mathbf{m}_{-i} , denoted as $\mathbf{m}_{-i}^{1(1)}$ and $\mathbf{m}_{-i}^{1(2)}$. When we individually substitute these realizations into Equation (6) and subtract one from the other, we obtain the following equation:

$$\begin{aligned} & \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1] \\ &= [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot [p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})]. \end{aligned} \quad (28)$$

By a similar argument, for realizations $\mathbf{m}_{-i}^{1(1)}$ and $\mathbf{m}_{-i}^{1(3)}$, we could derive the following:

$$\begin{aligned} & \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1] \\ &= [\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)] \cdot [p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})]. \end{aligned} \quad (29)$$

Dividing Equation (29) by Equation (28) would imply the following ratio:

$$\frac{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1]}{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1]} = \frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}. \quad (30)$$

Repeating the above steps for another realization \mathbf{m}_i^2 , one could derive a similar equation:

$$\frac{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)})|\mathbf{m}_i^2] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})|\mathbf{m}_i^2]}{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})|\mathbf{m}_i^2] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})|\mathbf{m}_i^2]} = \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}. \quad (31)$$

Let us consider any three pairs of games that satisfy the condition of equal choice probability; for instance, $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ for $l = 1, 2, 3$. A combination of Equations (30) and (31) would imply the following relationship:

$$\begin{aligned} \frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})} &= \frac{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(3)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1]}{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})|\mathbf{m}_i^1] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})|\mathbf{m}_i^1]} \\ &= \frac{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)})|\mathbf{m}_i^2] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})|\mathbf{m}_i^2]}{\hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})|\mathbf{m}_i^2] - \hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})|\mathbf{m}_i^2]} \\ &= \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(3)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}. \end{aligned} \quad (32)$$

The first and third lines in Equation (32) are direct results of Equations (30) and (31).

The second line follows the equal choice probability condition and Equation (7) so that $\hat{F}_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)})|\mathbf{m}_i^1] = \hat{F}_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})|\mathbf{m}_i^2]$ for $l = 1, 2, 3$. This completes the proof. \square

Following Proposition 5, if Xie (2022)'s testable implication is violated, our over-identification test in Proposition 2 would also reject QRE. Importantly, the reverse is not true since our test includes more restrictions than Xie (2022) and has higher statistical power. Specifically, the structure of monetary payoffs often implies additional restrictions on players' utilities across games or action profiles. As shown in Subsection 3.3, each player's utility function is non-parametrically identified. Therefore, these utility

restrictions become testable implications of QRE in addition to the ones derived by Xie (2022). To better describe these results, we extend Assumption 6 to include another structural property of the matching pennies game presented in Table 1. The property, indexed as Assumption 6(c), preserves player i 's payoffs for one of the other player's actions and varies player i 's payoffs when player $-i$ chooses the other action.

Assumption 6. (c) For each player i , there exist two realizations of \mathbf{m}_i —denoted as \mathbf{m}_i^1 and \mathbf{m}_i^2 —such that $m_i^1(a_i, a_{-i}) = m_i^2(a_i, a_{-i}) \forall a_i$ and for some a_{-i} .

The strict monotonicity of the utility function and each condition in Assumption 6 implies different testable implications of QRE. Proposition 6 shows that these implications are included in our over-identification test.

Proposition 6. Suppose that Assumptions 1 to 4 hold, then Equation (7) implies the following testable restrictions of QRE:

(a) $\forall \mathbf{m}_i \in \mathcal{M}_i$:

$$\begin{aligned} & \text{Sign}\left\{ \frac{(-1)^{a-i} p_{-i}(1 - a_{-i} | \mathbf{m}_i, \mathbf{m}_{-i}^2) F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^1)] + (-1)^{1-a-i} p_{-i}(1 - a_{-i} | \mathbf{m}_i, \mathbf{m}_{-i}^1) F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^2)]}{p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^1) - p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^2)} \right\} \\ & = \text{Sign}[m_i(a_i = 0, a_{-i}) - m_i(a_i = 1, a_{-i})], \quad \forall \mathbf{m}_{-i}^1, \mathbf{m}_{-i}^2 \in \mathcal{M}_{-i} \text{ and } \forall a_{-i}. \end{aligned} \quad (33)$$

(b) $\forall \mathbf{m}_i^1$ that satisfies Assumption 6(a):

$$\frac{1 - 2p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^1)}{1 - 2p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^2)} = \frac{F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^1)]}{F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^2)]}, \quad \forall \mathbf{m}_{-i}^1, \mathbf{m}_{-i}^2 \in \mathcal{M}_{-i}. \quad (34)$$

(c) For each pair of \mathbf{m}_i^1 and \mathbf{m}_i^2 that satisfies Assumption 6(b):

$$\begin{aligned} & \frac{(-1)^{a-i} p_{-i}(1 - a_{-i} | \mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)}) F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})] + (-1)^{1-a-i} p_{-i}(1 - a_{-i} | \mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})]}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)}) - p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})} = \\ & \frac{(-1)^{a'-i} p_{-i}(1 - a'_{-i} | \mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)}) F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})] + (-1)^{1-a'-i} p_{-i}(1 - a'_{-i} | \mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}) F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})]}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)}) - p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})}, \\ & \quad \forall \mathbf{m}_{-i}^{1(1)}, \mathbf{m}_{-i}^{1(2)}, \mathbf{m}_{-i}^{2(1)}, \mathbf{m}_{-i}^{2(2)} \in \mathcal{M}_{-i}. \end{aligned} \quad (35)$$

(d) Consider each pair of \mathbf{m}_i^1 and \mathbf{m}_i^2 that satisfies both Assumption 6(c) and the condition that $\mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2)$ includes an interval. Then for any two pairs of games such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ for $l = 1, 2$, we have:

$$\frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})} = \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})}. \quad (36)$$

Proof. For any \mathbf{m}_i , consider two realizations denoted as \mathbf{m}_{-i}^1 and \mathbf{m}_{-i}^2 . Evaluating the definition of $\hat{F}_i^{-1}(p|\mathbf{m}_i)$ by Equation (6) at these two realizations leads to the following system:

$$\begin{aligned} F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^1)] &= \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^1) \\ F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^2)] &= \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1) + [\tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 0) - \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1)] \cdot p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^2). \end{aligned} \quad (37)$$

In the left-hand side of this system, we replace $\hat{F}_i^{-1}(p|\mathbf{m}_i)$ by $F_i^{-1}(p)$. This follows the implication by Equation (7) such that $\hat{F}_i^{-1}(p|\mathbf{m}_i) = F_i^{-1}(p) \forall \mathbf{m}_i$. This linear system by Equation (37) identifies utility difference $\tilde{\pi}_i(\cdot)$ as the following expression:

$$\begin{aligned} &\tilde{\pi}_i(\mathbf{m}_i, a_{-i}) \quad (38) \\ &= \frac{(-1)^{a_{-i}} p_{-i}(1 - a_{-i} | \mathbf{m}_i, \mathbf{m}_{-i}^2) F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^1)] + (-1)^{1-a_{-i}} p_{-i}(1 - a_{-i} | \mathbf{m}_i, \mathbf{m}_{-i}^1) F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^2)]}{p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^1) - p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^2)}. \end{aligned}$$

Note that there always exist \mathbf{m}_{-i}^1 and \mathbf{m}_{-i}^2 such that the denominator in the second line is non-zero. This is because $p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i})$ varies with \mathbf{m}_{-i} .

The property of the utility function and the structure of monetary payoffs impose restrictions on $\tilde{\pi}_i(\cdot)$. It is these restrictions that lead to Proposition 6. Specifically, the strict increasing property of $u_i(m)$ implies that $\tilde{\pi}_i(\mathbf{m}_i, a_{-i})$ and $[m_i(a_i = 0, a_{-i}) - m_i(a_i = 1, a_{-i})]$ have the same sign. It leads to Proposition 6(a). Assumption 6(a) suggests that

$\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}) = -\tilde{\pi}_i(\mathbf{m}_i^1, 1 - a_{-i})$. It leads to Proposition 6(b). Assumption 6(c) restricts $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i}) = -\tilde{\pi}_i(\mathbf{m}_i^2, a'_{-i})$ and implies Proposition 6(c).

To prove Proposition 6(d), assume that Assumption 6(c) holds for the action $a_{-i} = 1$. The proof for the case that $a_{-i} = 0$ follows a similar argument and is suppressed. By transforming the system by Equation (37), we obtain the following relationship:

$$\frac{p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^1)}{p_{-i}(\mathbf{m}_i, \mathbf{m}_{-i}^2)} = \frac{F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^1)] - \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1)}{F_i^{-1}[p_i(\mathbf{m}_i, \mathbf{m}_{-i}^2)] - \tilde{\pi}_i(\mathbf{m}_i, a_{-i} = 1)}. \quad (39)$$

Now consider two pairs of games such that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ for $l = 1, 2$, we then obtain Proposition 6(d) through the following steps:

$$\begin{aligned} \frac{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})}{p_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})} &= \frac{F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(1)})] - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)}{F_i^{-1}[p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(2)})] - \tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1)} \\ &= \frac{F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})] - \tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1)}{F_i^{-1}[p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})] - \tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1)} \\ &= \frac{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(1)})}{p_{-i}(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(2)})}. \end{aligned} \quad (40)$$

The first and third lines follow directly from Equation (39). The second line is due to the equal choice probability condition that $p_i(\mathbf{m}_i^1, \mathbf{m}_{-i}^{1(l)}) = p_i(\mathbf{m}_i^2, \mathbf{m}_{-i}^{2(l)})$ and the implication of Assumption 6(c) such that $\tilde{\pi}_i(\mathbf{m}_i^1, a_{-i} = 1) = \tilde{\pi}_i(\mathbf{m}_i^2, a_{-i} = 1)$. This completes the proof. \square

Given Proposition 1 that identifies $F_i^{-1}(\cdot)$, all equations in Proposition 6 are then restrictions on known functions of players' choice probabilities and are therefore testable.

In Proposition 6, restriction (a) exploits only the strict monotonicity of $u_i(m)$ and applies to any payoff structure and any type of games. In contrast, restrictions (b) to (d) focus on matching pennies games. In addition, restrictions (a) to (c) do not require the equal choice probability condition. Restriction (d) requires this condition but only for two pairs of games as opposed to the three pairs in Xie (2022). Therefore, all restrictions

in Proposition 6 are additional testable implications in our over-identification test, but they are excluded from Xie (2022).

Even though Proposition 6(b) to (d) focuses on matching pennies games, other types of games have their own monetary payoffs structure. These structural properties can be exploited to derive additional testable restrictions of QRE. For instance, consider the coordination game illustrated in Table 2. Assumption 6(b) holds when the analyst considers two values $m_i^1 = 0$ and $m_i^2 = 15$. Therefore, Proposition 6(c) applies. Moreover, in this coordination game, the payoff of $a_i = 0$ does not depend on the other player's action. It implies that $\tilde{\pi}_i(m_i, a_{-i} = 0) - \tilde{\pi}_i(m_i, a_{-i} = 1) = u_i(15) - u_i(0)$, which is independent of m_i . Consequently, the following is a natural testable implication of QRE:

$$\frac{F_i^{-1}[p_i(m_i, m_{-i}^1)] - F_i^{-1}[p_i(m_i, m_{-i}^2)]}{p_{-i}(m_i, m_{-i}^1) - p_{-i}(m_i, m_{-i}^2)} \text{ is independent of } m_i, \forall m_{-i}^1, m_{-i}^2. \quad (41)$$

In another section of this online appendix (i.e., Section “Generalizations and Extensions”), we consider an experiment that varies at least two action profiles' payoffs without further restrictions on the payoff structures. For instance, it does not require Assumption 6. This general structure includes common types of games as special cases. In that section, we demonstrate that there are additional testable restrictions of QRE.

Even though experimental data typically provides additional structure to test QRE, the concrete form of the restrictions depends on the structure of \mathcal{M} and is therefore application specific. While it can be cumbersome to derive and list all such restrictions for a given application, these additional testable restrictions are included in our over-identification test as shown in Proposition 6, and therefore the analyst only needs to test the simple condition in Equation (7).

Non-Parametric Identification and Testing of Quantal Response Equilibrium

Online Appendix: Comparative Statics Analysis

Johannes Hoelzemann Ryan Webb Erhao Xie

April 22, 2024

We consider a transformed random error $\hat{\tilde{\epsilon}}_i = \frac{1}{\lambda} \tilde{\epsilon}_i$ that scales up the standard deviation of $\tilde{\epsilon}_i$ by $\frac{1}{\lambda}$, where $0 < \lambda < \infty$. Under this scaled error, player i 's choice probability is characterized as:

$$p_i(\mathbf{m}) = F_i\left\{\lambda \left[EU_i(\mathbf{m}_i, a_i = 0, b_i(\mathbf{m})) - EU_i(\mathbf{m}_i, a_i = 1, b_i(\mathbf{m}))\right]\right\}, \quad (42)$$

where $b_i(\mathbf{m})$ represents player i 's belief about the probability that player $-i$ will choose $a_{-i} = 0$. QRE places the restriction that $b_i(\mathbf{m}) = p_{-i}(\mathbf{m})$ so that Equation (42) turns to the quantal response function by Equation (3) when $\lambda = 1$. Moreover, Equation (42) also includes Level- k behaviors when $b_i(\mathbf{m})$ is the belief of the level- k player.

Equation (42) indicates that as λ increases, or equivalently as $Var(\tilde{\epsilon}_i)$ decreases, player i will choose $a_i = 0$ more (less) frequently if such an action has a higher (lower) expected utility than $a_i = 1$. Furthermore, when $\lambda \rightarrow \infty$, player i will unambiguously choose the action that maximizes the expected utility, provided that $F_i(-\infty) = 0$ and $F_i(\infty) = 1$. Conversely, as $\lambda \rightarrow 0$, player i will choose $a_i = 0$ with probability $F_i(0)$. If the analyst imposes the restriction that $Median(\tilde{\epsilon}_i) = 0$ so that $F_i(0) = 1/2$, then player i simply randomizes each action with equal probability.

Next, consider the matching pennies game in Table 1 and suppose that the analyst imposes the QRE restrictions. Given the normalization that the utility of the lowest

payoff (i.e., $m = 8$) is zero, each player's $p_i(\mathbf{m})$ is determined by the following equation system:

$$\begin{aligned} p_1(m_1, m_2) &= F_1 \{ \lambda [(u_1(m_1) + u_1(16)) \cdot p_2(m_1, m_2) - u_1(16)] \}, \\ p_2(m_1, m_2) &= F_2 \{ \lambda [u_2(m_2) - (u_2(m_2) + u_2(16)) \cdot p_1(m_1, m_2)] \}. \end{aligned} \quad (43)$$

Suppose that both $u_i(m)$ and $F_i(\tilde{\epsilon}_i)$ are continuously differentiable, then taking derivative with respect to (m_1, m_2) on both sides of Equation (43) would imply the following comparative statics under the QRE framework:

$$\begin{aligned} \frac{\partial p_1(\mathbf{m})}{\partial m_1} &= \frac{\lambda \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot u'_1(m_1) \cdot p_2(\mathbf{m})}{1 + \lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot [u_1(m_1) + u_1(16)] \cdot [u_2(m_2) + u_2(16)]} > 0, \\ \frac{\partial p_1(\mathbf{m})}{\partial m_2} &= \frac{\lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot [u_1(m_1) + u_1(16)] \cdot u'_2(m_2) \cdot [1 - p_1(\mathbf{m})]}{1 + \lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot [u_1(m_1) + u_1(16)] \cdot [u_2(m_2) + u_2(16)]} > 0, \\ \frac{\partial p_2(\mathbf{m})}{\partial m_1} &= \frac{-\lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot u'_1(m_1) \cdot [u_2(m_2) + u_2(16)] \cdot p_2(\mathbf{m})}{1 + \lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot [u_1(m_1) + u_1(16)] \cdot [u_2(m_2) + u_2(16)]} < 0, \\ \frac{\partial p_2(\mathbf{m})}{\partial m_2} &= \frac{\lambda \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot u'_2(m_2) \cdot [1 - p_1(\mathbf{m})]}{1 + \lambda^2 \cdot f_1(\lambda \cdot \widetilde{EU}_1) \cdot f_2(\lambda \cdot \widetilde{EU}_2) \cdot [u_1(m_1) + u_1(16)] \cdot [u_2(m_2) + u_2(16)]} > 0, \end{aligned} \quad (44)$$

where $f_i(\cdot)$ is the P.D.F. of $\tilde{\epsilon}_i$ and \widetilde{EU}_i is the difference between the expected utilities of actions 0 and 1. Moreover, both $p_i(\mathbf{m})$ and \widetilde{EU}_i are evaluated at the QRE conditions. The directions of the own-payoff effect and the other-payoff effect, as shown in Equation (44), are intuitive and are consistent with the reduced form results in Table 6. Moreover, Equation (44) also provides insights into the comparative statics of these effects with respect to λ . When $\lambda \rightarrow 0$, both $\frac{\partial p_i(\mathbf{m})}{\partial m_i}$ and $\frac{\partial p_i(\mathbf{m})}{\partial m_{-i}}$ converge to zero. These diminishing own-payoff and other-payoff effects are consistent with the property that each player randomizes each action with equal probability when $\lambda \rightarrow 0$. Conversely, consider the other extreme that $\lambda \rightarrow \infty$. Since the expression of $\frac{\partial p_i(\mathbf{m})}{\partial m_i}$ has the term λ on its nominator and the term λ^2 in the denominator, the effect of own payoff m_i on $p_i(\mathbf{m})$ decreases in the

order of λ . Conversely, the expression of $\frac{\partial p_i(\mathbf{m})}{\partial m_{-i}}$ has the term λ^2 in both its nominator and denominator. Therefore, the effect of the other player's payoff m_{-i} on $p_i(\mathbf{m})$ is order-invariant with respect to λ . As $\lambda \rightarrow \infty$, the own-payoff effect disappears while the other-payoff effect remains, as predicted in Nash Equilibrium.

Equation (45) offers another perspective for interpreting the comparative statics of the other-payoff effect. In QRE, player i anticipates that player $-i$'s payoff m_{-i} has a diminishing impact (in order of λ) on player $-i$'s choice probability $p_{-i}(\mathbf{m})$. This diminishing impact is entirely offset by the effect of $p_{-i}(\mathbf{m})$ on $p_i(\mathbf{m})$, which grows in the order of λ as shown in Equation (45). Consequently, the other-payoff effect, quantified by $\frac{\partial p_i(\mathbf{m})}{\partial m_{-i}}$, is order-invariant with respect to λ .

$$\begin{aligned}\frac{\partial p_1(\mathbf{m})}{\partial p_2(\mathbf{m})} &= \lambda f_1(\lambda \widetilde{EU}_1)[u_1(m_1) + u_1(16)], \\ \frac{\partial p_2(\mathbf{m})}{\partial p_1(\mathbf{m})} &= -\lambda f_2(\lambda \widetilde{EU}_1)[u_2(m_2) + u_2(16)].\end{aligned}\quad (45)$$

The structure of the matching pennies game in Table 1 also implies an interesting feature under Level- k behaviors. Specifically, when $m_i < 16$ ($m_i > 16$), the level-1 player would obtain a strictly lower (higher) expected utility of action 0 than action 1. Therefore, as $\lambda \rightarrow \infty$, the level-1 player will choose $a_i = 0$ with probability 0 (1). Due to the hierarchy of beliefs, players with higher types would also choose one of the actions with certainty, and such a choice is independent of players' risk preference. In summary, under level- k models, the effect of players' risk preference parameter ν on their behaviors vanishes in the limiting case as $\lambda \rightarrow \infty$ or $\lambda \rightarrow 0$.²⁵

Figures 11 to 14 plot $p_i(m_i, m_{-i})$ for both players in our Monte Carlo exercise. These figures aim to illustrate how the value of $Var(\tilde{\epsilon}_i)$ will affect each player's behavior under various models, including QRE and Level- k with $k \in \{1, 2, 3\}$. We consider three scenarios: (1) original value of $Var(\tilde{\epsilon}_i)$ in our Monte Carlo exercise, (2) doubling the value of

²⁵Note that when $\lambda \rightarrow 0$, each player randomizes their actions with equal probability, regardless of their expected utilities.

$Var(\tilde{\epsilon}_i)$, and (3) the limiting case where $Var(\tilde{\epsilon}_i) \rightarrow 0$. Clearly, these figures demonstrate the substantial impact of $Var(\tilde{\epsilon}_i)$ on each player's behavior.

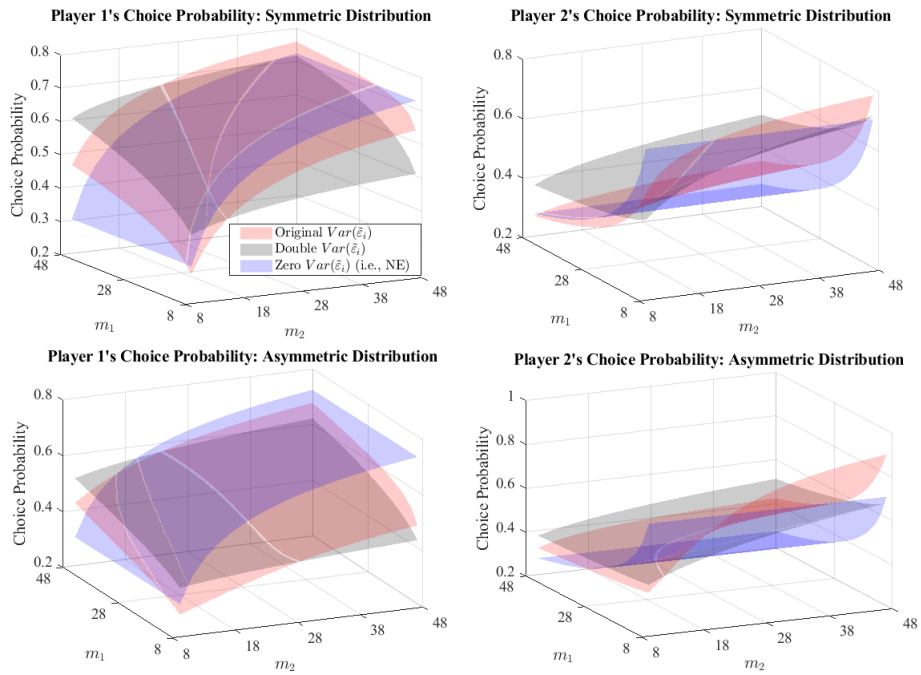


Figure 11: Players' Choice Probabilities: QRE Behavior

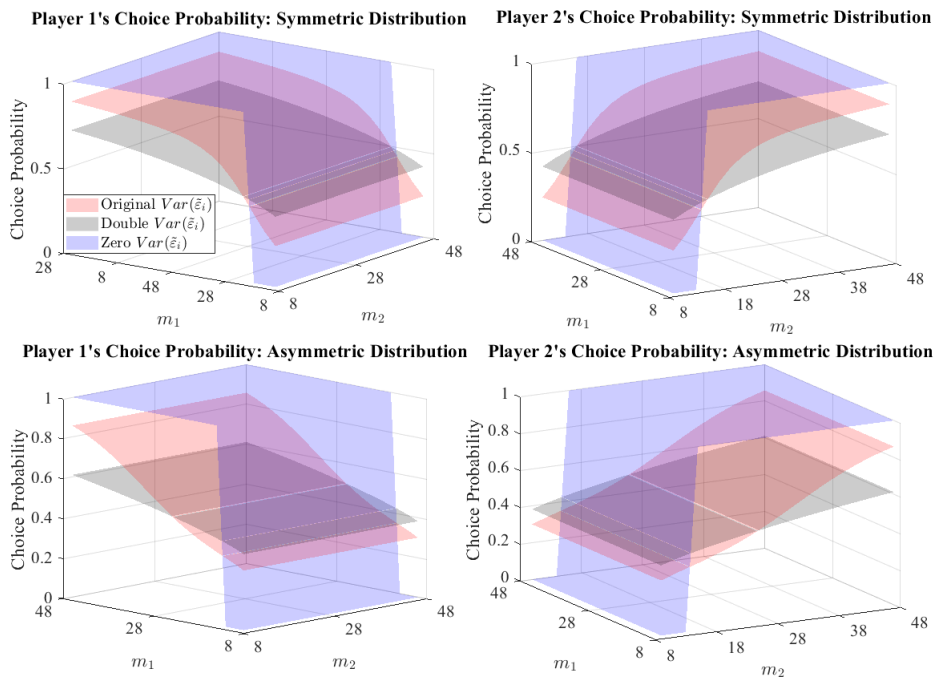
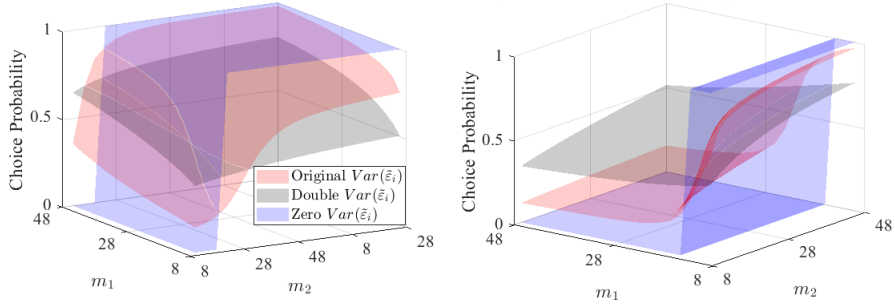


Figure 12: Players' Choice Probabilities: Level-1 Reasoning Behavior

Player 1's Choice Probability: Symmetric Distribution Player 2's Choice Probability: Symmetric Distribution



Player 1's Choice Probability: Asymmetric Distribution Player 2's Choice Probability: Asymmetric Distribution

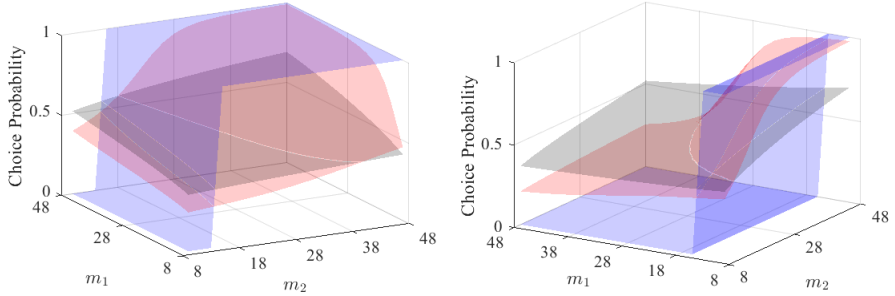
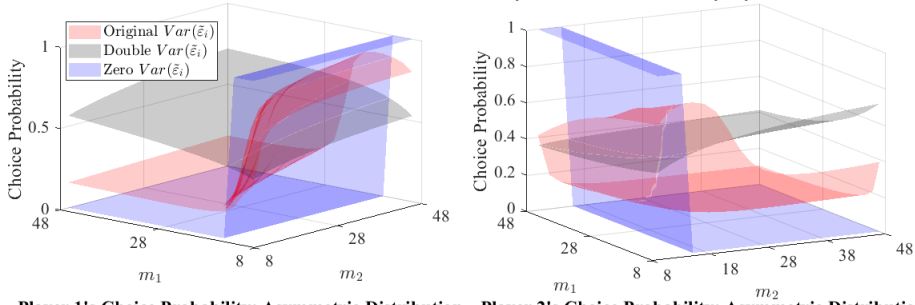


Figure 13: Players' Choice Probabilities: Level-2 Reasoning Behavior

Player 1's Choice Probability: Symmetric Distribution Player 2's Choice Probability: Symmetric Distribution



Player 1's Choice Probability: Asymmetric Distribution Player 2's Choice Probability: Asymmetric Distribution

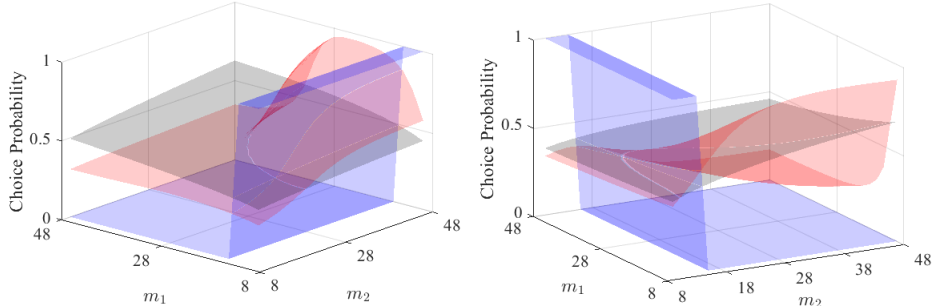


Figure 14: Players' Choice Probabilities: Level-3 Reasoning Behavior

Non-Parametric Identification and Testing of Quantal Response Equilibrium

Online Appendix: Generalizations and Extensions

Johannes Hoelzemann Ryan Webb Erhao Xie

April 22, 2024

Games with More Players and / or More Actions

This section extends the results in the main text to a *general multinomial choice game* with $N \geq 2$ players. We use letters i and j to denote a single player. Letter $-i$ represents all players other than i . Each player i simultaneously chooses an action, denoted as a_i , from their action set $\mathcal{A}_i = \{0, 1, \dots, K_i\}$. The number of actions (i.e., $K_i + 1 \geq 2$) is unrestricted and could be heterogeneous across players. Moreover, let $\mathbf{a} = (a_i, \mathbf{a}_{-i}) \in \mathcal{A} = \times_{j=1}^N \mathcal{A}_j$ be an action profile of this game, where $\mathbf{a}_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_N)$ is the decision profile made by all players other than i . In an experiment, player i will receive a monetary payoff $m_i(\mathbf{a})$, in the unit of experimental currency, when \mathbf{a} is the realized outcome. Consequently, given the utility function $u_i(m)$, player i would obtain a utility $u_i[m_i(\mathbf{a})]$ for the profile \mathbf{a} .

Recall that, as described in the main text, the monetary reward $m_i(\mathbf{a})$ is a control variable in the econometric model. This variable has a support $\mathcal{M}_i(\mathbf{a}) \subset \mathbb{R}$. Furthermore, define \mathbf{m}_i as a $\prod_j (K_j + 1) \times 1$ vector, and each element in this vector represents player i 's monetary payoff of the corresponding action profile. Naturally, the vector $\mathbf{m} = (\mathbf{m}'_1, \dots, \mathbf{m}'_N)'$ then summarizes the rewards across profiles and across players. In addition, we use $p_{-i}(\mathbf{a}_{-i} | \mathbf{m})$ to denote the probability that the profile \mathbf{a}_{-i} is chosen by all

players other than i given \mathbf{m} . Similarly, $\mathbf{p}_{-i}(\mathbf{m})$ is a $\prod_{j \neq i} (K_j + 1) \times 1$ vector that consists of the probability of each profile \mathbf{a}_{-i} . With these notations, the function of player i 's expected utility for action $a_i = k$ is expressed as:

$$EU_i[\mathbf{m}_i, a_i = k, \mathbf{p}_{-i}(\mathbf{m})] = \sum_{\mathbf{a}_{-i}} u_i[m_i(a_i = k, \mathbf{a}_{-i})] \cdot p_{-i}(\mathbf{a}_{-i} | \mathbf{m}). \quad (46)$$

In this game with potentially more than two actions, player i 's random perturbations extend to a $(K_i + 1) \times 1$ vector denoted as $\boldsymbol{\epsilon}_i = (\epsilon_i(0), \epsilon_i(1), \dots, \epsilon_i(K_i))'$. Specifically, this vector of errors follows a joint C.D.F. represented by $\Gamma_i(\boldsymbol{\epsilon}_i)$. Moreover, each element in this vector, denoted as $\epsilon_i(a_i)$, represents player i 's calculation error when evaluating the expected utility of the corresponding action. Due to the perturbations of these mistakes, player i will choose $a_i = k$ if and only if

$$EU_i[\mathbf{m}_i, a_i = k, \mathbf{p}_{-i}(\mathbf{m})] + \epsilon_i(k) \geq EU_i[\mathbf{m}_i, a_i = k', \mathbf{p}_{-i}(\mathbf{m})] + \epsilon_i(k'), \quad \forall k' \neq k. \quad (47)$$

Define $\mathbf{p}_i(\mathbf{m}) = (p(a_i = 1 | \mathbf{m}), p(a_i = 2 | \mathbf{m}), \dots, p(a_i = K_i | \mathbf{m}))'$ as a $K_i \times 1$ vector that includes player i 's choice probability of each action. Note that since the sum of probabilities of all actions equals 1, the choice probability of the base action $a_i = 0$ is suppressed in the vector $\mathbf{p}_i(\mathbf{m})$. Cautiously, our treatment of the vector $\mathbf{p}_{-i}(\mathbf{m})$ is slightly different as this vector consists of the probability of every action profile, including the base profile $\mathbf{a}_{-i} = \mathbf{0}$. This slight distinction in the treatment of $\mathbf{p}_i(\mathbf{m})$ and $\mathbf{p}_{-i}(\mathbf{m})$ simplifies the presentation and proofs of our results.

Under QRE, players' decisions are independent conditional on \mathbf{m} . Therefore, the joint choice probabilities $\mathbf{p}_{-i}(\mathbf{m})$ for all players other than i are determined solely by the individual choice probabilities $\mathbf{p}_j(\mathbf{m})$ for each player $j \neq i$. For instance, $p_{-i}(\mathbf{a}_{-i} | \mathbf{m}) = \prod_{j \neq i} p_j(a_j | \mathbf{m})$. If we interpret QRE as BNE in an incomplete information game where ϵ_i represents player i 's private information, the above conditional independence arises from the assumption of independent private information across players (i.e., $\epsilon_i \perp \epsilon_j \quad \forall i \neq j$).

In incomplete information games, when $\text{corr}(\epsilon_i, \epsilon_j) \neq 0$, player i 's private information becomes informative about player j 's payoffs and potential decisions. Therefore, each player should adjust their strategies based on their private information, leading to conditional correlated strategies across players. However, within the QRE framework, where ϵ_i is viewed as player i 's mistakes rather than private information, it raises an issue for the above channel of correlated actions. Specifically, if player i 's strategy depends on the value of ϵ_i , they should be able to distinguish actual utility and calculation errors. However, if such a distinction is clear, player i should not make any mistakes. Due to this contradiction, we are not aware of any studies in the framework of QRE that consider ϵ_i to be correlated across players.

In this general multinomial choice game, we modify our assumptions in the 2×2 game as presented in the main text. These modified assumptions are listed below.

Assumption 1'. *Each player i 's utility function $u_i(m)$ is bounded. Moreover, it is strictly increasing and continuously differentiable in m .*

Assumption 2'. *For each player i , let $\hat{A}(i)$ denote the set of action profiles for which the outcome variable has exogenous variations conditional on player i 's outcome variables of other profiles and other players' outcome variables. In other words, $\forall \mathbf{a} \in \hat{A}(i)$, $m_i(\mathbf{a})$ has exogenous variations conditional on $m_i(\mathbf{a}')$ $\forall \mathbf{a}' \in \mathcal{A}$ and $\forall \mathbf{m}_{-i} \in \mathcal{M}_{-i}$. We assume that $\hat{A}(i)$ consists of at least two distinct elements and $\cup_{\mathbf{a} \in \hat{A}(i)} \mathcal{M}_i(\mathbf{a}) = \cup_{\mathbf{a} \in \mathcal{A}} \mathcal{M}_i(\mathbf{a})$. This union is an interval that could be either open or closed.*

Assumption 3'. (a) $\Gamma_i(\epsilon_i)$ has a positive and continuous density function on \mathbb{R}^{K_i+1} , $\forall i$.

(b) $\Gamma_i(\epsilon_i)$ is independent of $(\mathbf{m}_i, \mathbf{m}_{-i})$, $\forall i$.

Assumption 4'. *For each player i , the function of choice probabilities $\mathbf{p}_i(\mathbf{m})$ varies with both \mathbf{m}_i and \mathbf{m}_{-i} . Moreover, $\mathbf{p}_i(\mathbf{m})$ is continuously differentiable for almost every $\mathbf{m} \in \mathcal{M}$. If there are values of \mathbf{m} for which $\mathbf{p}_i(\mathbf{m})$ is not continuously differentiable, the total number of these discontinuous points is finite.*

Assumption 5’. For each player i , there exists a value of the outcome variable, denoted as $m_i^1 \in \cup_{\mathbf{a} \in \hat{A}(i)} \mathcal{M}_i(\mathbf{a})$, such that $u_i'(m_i^1) = 1$.

Assumptions 1’ and 4’ are slightly stronger than Assumptions 1 and 4. These two modified assumptions require the utility function $u_i(m)$ and the choice probability function $\mathbf{p}_i(\mathbf{m})$ not only to be continuous but also to be differentiable. This differentiation simplifies the proofs in this general multinomial choice game with $N \geq 2$ players. Moreover, due to the expanded space of action profiles, Assumption 2’ requires the analyst to exogenously vary the monetary payoffs of at least two action profiles within each player. This is in contrast to most of the identification results in 2×2 games, where the exogenous variation of a single profile’s payoff suffices.

Assumption 3’(a) and Assumption 3’(b) are standard regularity and invariance conditions for the error distributions, adapted for games with more actions. These conditions allow for general error structures. Specifically, the error of each action could follow a heterogeneous marginal distribution and exhibit arbitrary correlation with the error of another action.

Assumption 5’ is an alternative but equivalent scale normalization compared to Assumption 5. Specifically, consider the affine transformation $u_i(m) = c + \beta \hat{u}_i(m)$, Assumption 5’ simply transforms $\hat{u}_i(m)$ to its equivalent form by setting $\beta = \frac{1}{\hat{u}_i'(m_i^1)}$. Since most of the proofs in this generalization are based on derivatives, it is convenient to normalize the marginal utility as in Assumption 5’.

In discrete choice models, the decision rule by Equation (47) implies the following mapping between player i ’s expected utility differences and their choice probabilities:

$$\mathbf{p}_i(\mathbf{m}) = \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))], \quad (48)$$

where $\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))$ is a $K_i \times 1$ vector that represents the difference in expected utilities for player i . In particular, the k^{th} element of this vector, denoted as $\widetilde{EU}_i[\mathbf{m}_i, a_i = k, \mathbf{p}_{-i}(\mathbf{m})]$, is defined as $EU_i[\mathbf{m}_i, a_i = k, \mathbf{p}_{-i}(\mathbf{m})] - EU_i[\mathbf{m}_i, a_i = 0, \mathbf{p}_{-i}(\mathbf{m})]$. As standard

in discrete choice models, these differences of expected utilities completely determine player i 's choice probabilities (Train, 2009). This relationship can be represented by the mapping $\mathbf{F}_i : \mathbb{R}^{K_i} \rightarrow \text{int}(\Delta^{K_i})$, where $\text{int}(\Delta^{K_i})$ denotes the interior of K_i -dimensional simplex. The k^{th} element in this mapping, denoted as $F_{i,k}(\cdot)$, then represents the resulting choice probability of action $a_i = k$. Under Assumption 3'(a), the mapping $\mathbf{F}_i(\cdot)$ is bijective (Norets and Takahashi, 2013; Sørensen and Fosgerau, 2022). Moreover, Hotz and Miller (1993) show that $\mathbf{F}_i(\cdot)$ is differentiable.

In this general multinomial choice game with $N \geq 2$ players, QRE is defined by a fixed-point condition, as summarized by Definition 1'.

Definition 1'. *The vector $(\mathbf{p}_1(\mathbf{m})', \mathbf{p}_2(\mathbf{m})', \dots, \mathbf{p}_N(\mathbf{m})')$ denotes the QRE choice probabilities if and only if the following condition holds:*

$$\mathbf{p}_i(\mathbf{m}) = \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))], \forall i \text{ and } \mathbf{m} \in \mathcal{M}. \quad (49)$$

For any $\mathbf{m} \in \mathcal{M}$, if there are multiple vectors that satisfy Equation (49) (i.e., multiple QRE), there exists a mechanism that selects one of the vectors / equilibria.

In this section, we first prove that player i 's utility function is non-parametrically identified. We then exploit this result to show the over-identification of $\mathbf{F}_i(\cdot)$ and the testable implication of QRE.

Proposition 7. *Suppose that Assumptions 1' to 5' and QRE restrictions hold, then the derivative of the utility function $u_i'(m)$ is point identified $\forall m \in \cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$ and $\forall i$.*

Proof. Under Assumption 2', let $\mathbf{a}' = (a'_i, a'_{-i}) \neq \mathbf{a}'' = (a''_i, a''_{-i})$ be the two action profiles in the set $\hat{A}(i)$. We assume that $a'_i \neq 0$ and $a''_i \neq 0$. This is without loss of generality since the analyst could always relabel player i 's actions. As described above, Assumptions 1', 3', and 4' imply that every function in Equation (49) is differentiable with respect to their arguments. Consequently, we could take derivative on both sides of Equation (49) and

obtain the following relationship:

$$\begin{aligned}\frac{\partial \mathbf{p}_i(\mathbf{m})}{\partial m_i(\mathbf{a}')} &= \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \frac{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial m_i(\mathbf{a}')} + \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \widetilde{\Pi}_i(\mathbf{m}_i) \frac{\partial \mathbf{p}_{-i}(\mathbf{m})}{\partial m_i(\mathbf{a}')}, \\ \frac{\partial \mathbf{p}_i(\mathbf{m})}{\partial m_i(\mathbf{a}'')} &= \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \frac{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial m_i(\mathbf{a}'')} + \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \widetilde{\Pi}_i(\mathbf{m}_i) \frac{\partial \mathbf{p}_{-i}(\mathbf{m})}{\partial m_i(\mathbf{a}'')}.\end{aligned}\tag{50}$$

Note that $\widetilde{\Pi}_i(\mathbf{m}_i)$ is a $K_i \times \prod_{j \neq i} (K_j + 1)$ matrix whose element in cell (a_i, \mathbf{a}_{-i}) is represented by $\tilde{\pi}_i(\mathbf{m}_i, a_i, \mathbf{a}_{-i}) = u_i[m_i(a_i, \mathbf{a}_{-i})] - u_i[m_i(a_i = 0, \mathbf{a}_{-i})]$.

Equation (49) suggests that $\mathbf{p}_i(\mathbf{m})$ could be alternatively represented as $\mathbf{p}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))$. This equivalent form could be consistently estimated from choice data, due to the consistent estimation of $\mathbf{p}_{-i}(\mathbf{m})$ as described in the main text. Consequently, we could take derivative for this equivalent form and obtain:

$$\frac{\partial \mathbf{p}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial \mathbf{p}'_{-i}(\mathbf{m})} = \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \widetilde{\Pi}_i(\mathbf{m}_i).\tag{51}$$

Substituting Equation (51) into Equation (50) leads to:

$$\begin{aligned}\frac{\partial \mathbf{p}_i(\mathbf{m})}{\partial m_i(\mathbf{a}')} - \frac{\partial \mathbf{p}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial \mathbf{p}'_{-i}(\mathbf{m})} \frac{\partial \mathbf{p}_{-i}(\mathbf{m})}{\partial m_i(\mathbf{a}')} &= \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \frac{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial m_i(\mathbf{a}')}, \\ \frac{\partial \mathbf{p}_i(\mathbf{m})}{\partial m_i(\mathbf{a}'')} - \frac{\partial \mathbf{p}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial \mathbf{p}'_{-i}(\mathbf{m})} \frac{\partial \mathbf{p}_{-i}(\mathbf{m})}{\partial m_i(\mathbf{a}'')} &= \frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \frac{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial m_i(\mathbf{a}'')}.\end{aligned}\tag{52}$$

Since each player's choice probability can be consistently estimated, the terms on the left-hand side of Equation (52) are known to the analyst. Consequently, the terms on the right-hand side are identified.

Consider an arbitrary $\mathbf{a} = (a_i, \mathbf{a}_{-i})$ where $a_i \neq 0$. The structure of expected utilities

implies the following expression:

$$\frac{\partial \mathbf{F}_i[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))} \frac{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))}{\partial m_i(\mathbf{a})} = \begin{pmatrix} \frac{\partial F_{i,1}[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{E}U_i(\mathbf{m}_i, a_i, \mathbf{p}_{-i}(\mathbf{m}))} u'_i[m_i(\mathbf{a})] p_{-i}(\mathbf{a}_{-i}|\mathbf{m}) \\ \frac{\partial F_{i,2}[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{E}U_i(\mathbf{m}_i, a_i, \mathbf{p}_{-i}(\mathbf{m}))} u'_i[m_i(\mathbf{a})] p_{-i}(\mathbf{a}_{-i}|\mathbf{m}) \\ \vdots \\ \frac{\partial F_{i,K_i}[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{E}U_i(\mathbf{m}_i, a_i, \mathbf{p}_{-i}(\mathbf{m}))} u'_i[m_i(\mathbf{a})] p_{-i}(\mathbf{a}_{-i}|\mathbf{m}) \end{pmatrix} \quad (53)$$

Substituting Equation (53) into Equation (52) implies that the terms $\frac{\partial F_{i,k}(\widetilde{\mathbf{E}}\mathbf{U}_i)}{\partial \widetilde{E}U_i(\cdot, a'_i)} u'_i[m_i(\mathbf{a}')] p_{-i}(\mathbf{a}'_{-i}|\mathbf{m})$ and $\frac{\partial F_{i,k}(\widetilde{\mathbf{E}}\mathbf{U}_i)}{\partial \widetilde{E}U_i(\cdot, a''_i)} u'_i[m_i(\mathbf{a}'')] p_{-i}(\mathbf{a}''_{-i}|\mathbf{m})$ are identified for each k . It further implies the following result:

$$\frac{\frac{\partial F_{i,a''}[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{E}U_i(\mathbf{m}_i, a''_i, \mathbf{p}_{-i}(\mathbf{m}))} u'_i[m_i(\mathbf{a}'')] p_{-i}(\mathbf{a}''_{-i}|\mathbf{m})}{\frac{\partial F_{i,a'}[\widetilde{\mathbf{E}}\mathbf{U}_i(\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m}))]}{\partial \widetilde{E}U_i(\mathbf{m}_i, a'_i, \mathbf{p}_{-i}(\mathbf{m}))} u'_i[m_i(\mathbf{a}')] p_{-i}(\mathbf{a}'_{-i}|\mathbf{m})} = \frac{u'_i[m_i(\mathbf{a}')] p_{-i}(\mathbf{a}'_{-i}|\mathbf{m})}{u'_i[m_i(\mathbf{a}'')] p_{-i}(\mathbf{a}''_{-i}|\mathbf{m})} \text{ is identified.} \quad (54)$$

The equality of Equation (54) follows the results in discrete choice models such that $\frac{\partial \mathbf{F}_i(\widetilde{\mathbf{E}}\mathbf{U}_i)}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\cdot)}$ can be seen as the hessian matrix of the social welfare function (Sørensen and Fosgerau, 2022). The social welfare function is strictly convex. Therefore, the matrix $\frac{\partial \mathbf{F}_i(\widetilde{\mathbf{E}}\mathbf{U}_i)}{\partial \widetilde{\mathbf{E}}\mathbf{U}_i(\cdot)}$ is symmetric and positive definite. This symmetry implies that $\frac{\partial F_{i,a''}(\cdot)}{\partial \widetilde{E}U_i(\cdot, a''_i)} = \frac{\partial F_{i,a'}(\cdot)}{\partial \widetilde{E}U_i(\cdot, a'_i)}$ and can be canceled out in Equation (54). Moreover, each term in the denominator is strictly positive so that the ratio is well defined. In particular, $\frac{\partial F_{i,a''}(\cdot)}{\partial \widetilde{E}U_i(\cdot, a''_i)} > 0$ due to the strictly positive density in Assumption 3'(a); $u'_i(\cdot) > 0$ due to the strict monotonicity as per Assumption 1'; $p_{-i}(\mathbf{a}''_{-i}|\mathbf{m}) > 0$ due to the full support condition of ϵ_i in Assumption 3'(a).

Equation (54) identifies $\frac{u'_i(m')}{u'_i(m'')}$ for any $m', m'' \in \cup_{\mathbf{a} \in \hat{\mathcal{A}}(i)} \mathcal{M}_i(\mathbf{a}) = \cup_{\mathbf{a} \in \mathcal{A}} \mathcal{M}_i(\mathbf{a})$. Assumption 5' normalizes the scale of $u'_i(m'_i)$ at one arbitrary value m'_i . This normalization then identifies $u'_i(m) \forall m \in \cup_{\mathbf{a} \in \mathcal{A}} \mathcal{M}_i(\mathbf{a})$ and completes the proof. \square

Proposition 7 identifies the marginal utility $u'_i(m)$ and consequently identifies the utility function in the class of $u_i(m) + c$. As in the main text, when the analyst considers

the normalization, such as $u_i(0) = 0$ or $u_i[\min\{\cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})\}] = 0$, the value of c is identified. It further identifies the utility function $u_i(m) \forall m \in \cup_{\mathbf{a}} \mathcal{M}_i(\mathbf{a})$.

As described above, Assumption 3'(a) implies that the mapping $\mathbf{F}_i(\cdot)$ is bijective (Norets and Takahashi, 2013; Sørensen and Fosgerau, 2022). Therefore, we could invert $\mathbf{F}_i(\cdot)$ and the QRE restriction by Equation (49) becomes:

$$\mathbf{F}_i^{-1}[\mathbf{p}_i(\mathbf{m})] = \widetilde{\mathbf{E}}\mathbf{U}_i[\mathbf{m}_i, \mathbf{p}_{-i}(\mathbf{m})] = \widetilde{\mathbf{\Pi}}_i(\mathbf{m}_i) \cdot \mathbf{p}_{-i}(\mathbf{m}). \quad (55)$$

This equation implies the non-parametric identification of the error distribution, as established by Proposition 1'.

Proposition 1'. *Suppose that Assumptions 1' to 5' and QRE restrictions hold; therefore, the marginal utility $u_i'(m)$ is point identified for each player i by Proposition 7. In the next step, suppose that the analyst fixes \mathbf{m}_i at an arbitrary value \mathbf{m}_i^1 and only considers the variation of \mathbf{m}_{-i} , then $\mathbf{F}_i^{-1}(\mathbf{p})$ is point identified $\forall \mathbf{p} \in \mathcal{P}_i(\mathbf{m}_i^1)$.*

Proof. As described above, the identification of $u_i'(m)$ implies that the utility function is identified in the class of $u_i(m) + c$. Consequently, the difference of utilities $\tilde{\pi}_i(\mathbf{m}_i, a_i, \mathbf{a}_{-i}) = u_i[m_i(a_i, \mathbf{a}_{-i})] - u_i[m_i(a_i = 0, \mathbf{a}_{-i})]$ is uniquely determined as the constant c is canceled out. It further implies that the matrix $\widetilde{\mathbf{\Pi}}_i(\mathbf{m}_i)$ is known to the analyst for each $\mathbf{m}_i \in \mathcal{M}_i$. For an arbitrary value \mathbf{m}_i^1 , Equation (55) turns to:

$$\mathbf{F}_i^{-1}[\mathbf{p}_i(\mathbf{m}_i^1, \mathbf{m}_{-i})] = \widetilde{\mathbf{\Pi}}_i(\mathbf{m}_i^1) \cdot \mathbf{p}_{-i}(\mathbf{m}_i^1, \mathbf{m}_{-i}). \quad (56)$$

The terms on the right-hand side are known to the analyst. Consequently, the exogenous variation of \mathbf{m}_{-i} then identifies $\mathbf{F}_i^{-1}(\mathbf{p})$ for all values of \mathbf{p} in the support of $\mathcal{P}_i(\mathbf{m}_i^1)$. This completes the proof. \square

Due to the inverse relationship between $\mathbf{F}_i^{-1}(\cdot)$ and $\mathbf{F}_i(\cdot)$, Proposition 1' implies the non-parametric identification of the mapping $\mathbf{F}_i(\cdot)$. It further implies that the distribution

of the difference of errors $\tilde{\epsilon}_i = (\epsilon_i(1) - \epsilon_i(0), \epsilon_i(2) - \epsilon_i(0), \dots, \epsilon_i(K_i) - \epsilon_i(0))'$ is uniquely determined (Train, 2009).

Recall that $\hat{\mathbf{F}}_i^{-1}(\mathbf{p}|\mathbf{m}_i^1)$ represents the inverted choice probability function that satisfies the QRE restrictions when the analyst fixes \mathbf{m}_i at \mathbf{m}_i^1 . Proposition 1' directly implies a testable restriction of QRE, as established below.

Proposition 2'. *Suppose that Assumptions 1' to 5' and QRE restrictions hold. Consider any two realizations of \mathbf{m}_i , denoted as \mathbf{m}_i^1 and \mathbf{m}_i^2 . Suppose that $\mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2) \neq \emptyset$, then the null hypothesis of QRE implies the following testable restriction:*

$$\hat{\mathbf{F}}_i^{-1}(\mathbf{p}|\mathbf{m}_i^1) = \hat{\mathbf{F}}_i^{-1}(\mathbf{p}|\mathbf{m}_i^2), \forall \mathbf{p} \in \mathcal{P}_i(\mathbf{m}_i^1) \cap \mathcal{P}_i(\mathbf{m}_i^2). \quad (57)$$

Proof. A direct implication of Proposition 1'. □

Comparison to Xie (2022) In multinomial choice games with field data, Xie (2022) establishes the non-parametric identification of the utility and the error distribution. However, to prove these results, he imposes two strong restrictions: one on the model primitive and the other on the data. In contrast, this paper shows that in experimental datasets where the outcome variable is observed, the non-parametric identification results can be achieved without imposing these two strong restrictions. We elaborate these two restrictions below.

The first restriction in Xie (2022) is that the error distribution must satisfy a *rank ordering property*. Under this property, one action is chosen more frequently than another if and only if it yields a strictly higher expected utility. While this assumption is often made in empirical applications of QRE, it imposes strong restrictions on the error distribution. In particular, this property rules out error structures with flexible correlations and features of heteroskedasticity. To illustrate this point, consider an agent facing three actions labeled as 1, 2, and 3, with associated expected utilities denoted as

$EU(1)$, $EU(2)$, and $EU(3)$. Without loss of generality, suppose that $EU(1) > EU(2) > EU(3)$. Consider one structure where the errors $\varepsilon(k)$ follow the same marginal distribution but can be correlated across actions. In particular, $\varepsilon(1)$ and $\varepsilon(2)$ exhibit strong positive correlation, and they are independent of $\varepsilon(3)$. This error structure implies that $EU(1) + \varepsilon(1) > EU(2) + \varepsilon(2)$ would hold with high probability, leading to a low choice probability of action 2. However, action 3 might be chosen more frequently than action 2 due to its independent error $\varepsilon(3)$. Next, consider another structure where $\varepsilon(k)$ are independent across actions but are heterogeneous in their scales. In particular, the agent only makes minor mistakes for the first two actions, resulting in small $Var(\varepsilon(1))$ and $Var(\varepsilon(2))$. Consequently, action 2 is still chosen with a small probability since its perturbed expected utility is likely to be smaller than that of action 1. In contrast, suppose $\varepsilon(3)$ has a large variance; then the third action could be chosen more frequently than action 2, since this agent may frequently and mistakenly evaluate action 3 as highly attractive. In summary, under these two reasonable error structures, the agent will choose action 3 more frequently than action 2 even though action 3 has a lower expected utility. Clearly, the rank ordering property is violated.

Xie (2022) also considers another strong restriction on the data. In particular, he defines two actions to be *connected* if these two actions can be chosen with equal probability. His identification results require each pair of two actions to be either connected or linked through a sequence of connected actions. This particular data structure is challenging to construct in an experiment.

Implications of Assumption 2'

This section focuses on 2×2 games and discusses the implications of Assumption 2' on the results presented in the main text. First, under this assumption, the previous section establishes the identification results without assuming prior knowledge of the values of $F_i^{-1}(p^1)$ and $F_i^{-1}(p^2)$. Consequently, when player i 's monetary payoffs exhibit

exogenous variations for at least two action profiles, the analyst does not need to impose the payoff structures as required by Assumption 6 to identify the median and the mean of the errors.

Assumption 2' itself also implies that our over-identification test incorporates more restrictions than the test by Xie (2022). Specifically, let $\mathbf{a}' \neq \mathbf{a}''$ be the two action profiles in $\hat{\mathcal{A}}(i)$ as required by Assumption 2'. When $\mathcal{M}_i(\mathbf{a}') \cap \mathcal{M}_i(\mathbf{a}'')$ is an interval, there exist infinite pairs where each pair contains two values \mathbf{m}_i^1 and \mathbf{m}_i^2 such that $m_i^1(\mathbf{a}') = m_i^2(\mathbf{a}'')$ and $m_i^2(\mathbf{a}') = m_i^1(\mathbf{a}'')$. Moreover, between \mathbf{m}_i^1 and \mathbf{m}_i^2 , let player i 's monetary rewards hold constant for all action profiles other than \mathbf{a}' and \mathbf{a}'' . We claim that these pairs imply testable restrictions of QRE in addition to Xie (2022).

Consider the first scenario that $a'_{-i} \neq a''_{-i}$. The pair \mathbf{m}_i^1 and \mathbf{m}_i^2 implies that $\tilde{\pi}_i(\mathbf{m}_i^1, a'_{-i}) - \tilde{\pi}_i(\mathbf{m}_i^2, a'_{-i}) = \tilde{\pi}_i(\mathbf{m}_i^1, a''_{-i}) - \tilde{\pi}_i(\mathbf{m}_i^2, a''_{-i})$ if $a'_i \neq a''_i$ and $\tilde{\pi}_i(\mathbf{m}_i^1, a'_{-i}) - \tilde{\pi}_i(\mathbf{m}_i^2, a'_{-i}) = \tilde{\pi}_i(\mathbf{m}_i^2, a''_{-i}) - \tilde{\pi}_i(\mathbf{m}_i^1, a''_{-i})$ if $a'_i = a''_i$. Consider the other scenario that $a'_{-i} = a''_{-i}$. The pair \mathbf{m}_i^1 and \mathbf{m}_i^2 then implies that $\tilde{\pi}_i(\mathbf{m}_i^1, a'_{-i}) = -\tilde{\pi}_i(\mathbf{m}_i^2, a'_{-i})$. These linear restrictions on utility differences $\tilde{\pi}_i(\cdot)$ lead to additional linear restrictions on $\hat{F}_i^{-1}(p|\mathbf{m}_i)$ through the definition of $\hat{F}_i^{-1}(p|\mathbf{m})$ by Equation (6). These restrictions are therefore included in our over-identification test that is based on $\hat{F}_i^{-1}(p|\mathbf{m}_i)$. However, they are excluded from the test by Xie (2022).

Relaxing the Invariance Assumption When there are at least two action profiles with varying monetary payoffs as per Assumption 2', the analyst is able to relax the invariance condition as in Assumption 3(b) and still tests the null hypothesis of QRE. In particular, we allow player i 's distribution function $F_i(\tilde{\epsilon}_i)$ to depend on their own rewards \mathbf{m}_i and restrict it to be independent of other players' payoffs \mathbf{m}_{-i} . To capture this dependence, we adapt the notation $F_i(\tilde{\epsilon}_i|\mathbf{m}_i)$.

Let \mathbf{a}' and \mathbf{a}'' be the two distinct profiles in $\hat{\mathcal{A}}(-i)$. Suppose that the analyst fixes $m_{-i}(\mathbf{a}') = m_{-i}^1$ and varies player $-i$'s payoffs for the other profile \mathbf{a}'' . In alignment with the main text, we denote $\hat{F}_i^{-1}(p|\mathbf{m}_i, m_{-i}(\mathbf{a}') = m_{-i}^1)$ as the quantile function that satisfies

QRE restrictions when $m_{-i}(\mathbf{a}')$ is fixed at m_{-i}^1 . The proof of Proposition 1 indicates that this quantile function is identified for any value of $m_{-i}(\mathbf{a}')$ based on the variation provided by $m_{-i}(\mathbf{a}'')$. Consequently, the null hypothesis of QRE implies the following testable implication:

$$\hat{F}_i^{-1}(p|\mathbf{m}_i, m_{-i}(\mathbf{a}') = m_{-i}^1) = \hat{F}_i^{-1}(p|\mathbf{m}_i, m_{-i}(\mathbf{a}') = m_{-i}^2), \forall \mathbf{m}_i, m_{-i}^1, m_{-i}^2.$$

Non-Parametric Identification and Testing of Quantal Response Equilibrium

Online Appendix: Experimental Interface

Johannes Hoelzemann Ryan Webb Erhao Xie

April 22, 2024

Instructions

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money that will be paid to you at the end of the experiment in cash.

In this experiment, you will face 200 rounds of decision-making problems. During the experiment, and in order to determine your payment, you will be randomly matched with other participants in this session.

The Basic Idea

There will be 200 different rounds. In each round, you will be presented with an interactive decision problem similar to the one below.

Your Earnings

		Opponent's action	
		Top	Bottom
Your action	Top	16	8
	Bottom	8	16

Opponent's Earnings

		Your action	
		Top	Bottom
Opponent's action	Top	8	16
	Bottom	16	8

Your earnings in each problem depend on your choice of action (between "Top" and "Bottom") and your opponent's choice of action (between "Top" and "Bottom"). Your earnings possibilities are presented in tables like the ones above. In each problem, your earnings are given by the black numbers in the left table, labelled 'Your Earnings'. Your opponent's earnings are given in the right table, labelled 'Opponent's Earnings'. Your choice of action determines the row in 'Your Earnings' table and your opponent's choice of action determines the column in the same table. The black number in the cell corresponding to any combination of actions (yours and your opponent's) represent your earnings. Similarly, your opponent's choice of action determines the row in the 'Opponent's Earnings' table, while your choice of action determines the column in this table. The number in the cell corresponding to any combination of actions (yours and your opponent's) represents your opponent's earnings. In summary, your choice of action AND your opponent's choice of action affect both your earnings and your opponent's earnings.

For example, if you choose action "Top" and your opponent chooses action "Bottom" your earnings would be €8 and your opponent's earnings would be €16. If you choose action "Bottom" and your opponent chooses action "Bottom", your earnings would be €16 and your opponent's earnings would be €8. Numbers in the example are just an example and do not intend to suggest how anyone should make their choices.

Note that in each round only the bold numbers will change and the other numbers will not change.

In order to assist you to choose an action, when you move your mouse over a row in the 'Your Earnings' table on the left, the action will be highlighted in yellow in both tables: a row on the left table, and a column on the right table. By left clicking your mouse over a row it will remain highlighted, and you can unhighlight it by clicking your mouse again or clicking another row. Similarly, when you move your mouse over a row that corresponds to an action of your opponent in 'Opponent's Earnings' on the right, the row will be highlighted in yellow on the right table and the corresponding column will be highlighted in yellow on the left table. Clicking your mouse over the row will keep it highlighted, and clicking it again (or clicking another action) will unhighlight it.

Please try to highlight actions for you and your opponent in the earnings tables above.

Finally, to choose an action you must click on the button around the action name (the text next to the row, on the margin of the left table).

The Rounds

There will be 200 rounds. You will need to choose an action in each round, as described above. After you have made your choice of action you will be informed about the outcome and then advance to the next screen and play a new round.

The earnings tables in each round are different, so you should look carefully at them before making your choice. In each round only the bold numbers will change and the other numbers will not change.

Your Opponents

In each round, you will be randomly matched with another participant. All your matches are participating in this session. Your randomly matched opponent's screen displays the same two earnings tables, but in reverse order. You do not know which actions your opponent chooses when you make your choices of actions. You can, however, attempt to reason about the actions the other participant will choose.

Payment

You will earn a participation payment of €5 for participating in this experiment.

In addition to the participation payment, one round will be randomly selected for payment at the end of the experiment. You will be paid your earnings in that round as described above. Any of the 200 rounds could be the one

selected. This protocol of determining payments suggests that you should choose in each round as if it is the only round that determines your payment.

Before the actual experiment starts, you will be asked to answer some questions. You must answer these correctly in order to proceed to the next question.

You will be informed of your payment, the round chosen for payment, and the choices of you and of the other participant only at the end of the experiment. You will not learn any other information about the choices of other participants during the experiment. The identity of the other participants to which you will be matched will never be revealed.

Finally, after completing the experiment you will be paid your earnings in cash.

Frequently Asked Questions

Q1. Is this some kind of psychology experiment with an agenda you haven't told us? Answer. No. It is an economics experiment. If we do anything deceptive or don't pay you cash as described, then you can complain to University of Vienna Research Ethics Board and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make decisions.

Q2. Is there a "correct" choice of action? Is this kind of a test? No. Your optimal action depends on your belief which actions will other participants choose. Different people may hold different beliefs.

Next

This button will be activated after 57 seconds. Please take your time to read through the instructions.

Quiz Time!

Q1. The computer will randomly select one of the rounds to determine your payment.

A Correct

B Incorrect

Quiz Time!

Q1. **Correct Answer** ×

A

B

Quiz Time!

Q2. In each round you will be randomly matched with another participant in this session.

A Correct

B Incorrect

Quiz Time!

Q2. **Wrong Answer** ×

A

The answer is incorrect. Please try again.

B Incorrect

Quiz Time!

Q3. In each round the earning possibilities never change and stay always the same.

A Correct

B Incorrect

Quiz Time!

Q3. **Correct Answer** x

A

✓ Proceed to the experiment

B

Your Choice

You are in Round 1 of 200

You are randomly matched with **another** participant.

Please make your choice by clicking on one of the two buttons on the left in "Your Earnings" table.

Your Earnings

		Opponent's action	
		Top	Bottom
Your action	Top	22	8
	Bottom	8	16

Opponent's Earnings

		Your action	
		Top	Bottom
Opponent's action	Top	8	18
	Bottom	16	8

Instructions

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money that will be paid to you at the end of the experiment in cash.

In this experiment, you will face 200 rounds of decision-making problems. During the experiment, and in order to determine your payment, you will be randomly matched with other participants in this session.

Please wait

Waiting for the other participant.



Your Choice

You are in Round 1 of 200

You are randomly matched with **another** participant.

Please make your choice by clicking on one of the two buttons on the left in "Your Earnings" table.

Your Earnings

		Opponent's action	
		Top	Bottom
Your action	Top	8	18
	Bottom	16	8

Opponent's Earnings

		Your action	
		Top	Bottom
Opponent's action	Top	22	8
	Bottom	8	16

Instructions

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money that will be paid to you at the end of the experiment in cash.

In this experiment, you will face 200 rounds of decision-making problems. During the experiment, and in order to determine your payment, you will be randomly matched with other participants in this session.

Please wait

Waiting for the other participant.



Round Outcome

You chose **TOP** and your opponent chose **TOP**.

As a result, your payoff is **€22.00**.

Next

Round Outcome

You chose **TOP** and your opponent chose **TOP**.

As a result, your payoff is **€8.00**.

Next

Your Choice

You are in Round 2 of 200

You are randomly matched with **another** participant.

Please make your choice by clicking on one of the two buttons on the left in "Your Earnings" table.

Your Earnings

		Opponent's action	
		Top	Bottom
Your action	Top	46	8
	Bottom	8	16

Opponent's Earnings

		Your action	
		Top	Bottom
Opponent's action	Top	8	34
	Bottom	16	8

Instructions

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money that will be paid to you at the end of the experiment in cash.

In this experiment, you will face 200 rounds of decision-making problems. During the experiment, and in order to determine your payment, you will be randomly matched with other participants in this session.

Please wait

Waiting for the other participant.



Your Choice

You are in Round 2 of 200

You are randomly matched with **another** participant.

Please make your choice by clicking on one of the two buttons on the left in "Your Earnings" table.

Your Earnings

		Opponent's action	
		Top	Bottom
Your action	Top	8	34
	Bottom	16	8

Opponent's Earnings

		Your action	
		Top	Bottom
Opponent's action	Top	46	8
	Bottom	8	16

Instructions

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money that will be paid to you at the end of the experiment in cash.

In this experiment, you will face 200 rounds of decision-making problems. During the experiment, and in order to determine your payment, you will be randomly matched with other participants in this session.

Please wait

Waiting for the other participant.



Round Outcome

You chose **TOP** and your opponent chose **TOP**.

As a result, your payoff is **€46.00**.

Next

Round Outcome

You chose **TOP** and your opponent chose **TOP**.

As a result, your payoff is **€8.00**.

Next

-
-
-
-
-
-

Thank you for participating in this experiment!

Your Earnings

The randomly selected round for payment is Round **58**. Your total earnings are therefore **€21**. This includes your participation payment of €5 for taking part in this experiment.

Your Results in More Detail

Round	Choices and Outcomes
1	You chose TOP and your opponent chose TOP . Your payoff is €22 (and your opponent's payoff is €8).
2	You chose TOP and your opponent chose TOP . Your payoff is €46 (and your opponent's payoff is €8).
3	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).
4	You chose TOP and your opponent chose Bottom . Your payoff is €8 (and your opponent's payoff is €16).
5	You chose TOP and your opponent chose Bottom . Your payoff is €8 (and your opponent's payoff is €16).
6	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).
7	You chose TOP and your opponent chose Bottom . Your payoff is €8 (and your opponent's payoff is €16).
8	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).
9	You chose Bottom and your opponent chose TOP . Your payoff is €8 (and your opponent's payoff is €26).
10	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).
11	You chose TOP and your opponent chose TOP . Your payoff is €16 (and your opponent's payoff is €8).
12	You chose TOP and your opponent chose TOP . Your payoff is €22 (and your opponent's payoff is €8).
13	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).
14	You chose TOP and your opponent chose Bottom . Your payoff is €8 (and your opponent's payoff is €16).
15	You chose Bottom and your opponent chose Bottom . Your payoff is €16 (and your opponent's payoff is €8).

-
-
-
-
-
-

194 You chose **Bottom** and your opponent chose **Bottom**. Your payoff is **€16** (and your opponent's payoff is €8).

195 You chose **TOP** and your opponent chose **TOP**. Your payoff is **€18** (and your opponent's payoff is €8).

196 You chose **TOP** and your opponent chose **Bottom**. Your payoff is **€8** (and your opponent's payoff is €16).

197 You chose **Bottom** and your opponent chose **TOP**. Your payoff is **€8** (and your opponent's payoff is €34).

198 You chose **Bottom** and your opponent chose **Bottom**. Your payoff is **€16** (and your opponent's payoff is €8).

199 You chose **TOP** and your opponent chose **TOP**. Your payoff is **€36** (and your opponent's payoff is €8).

200 You chose **Bottom** and your opponent chose **TOP**. Your payoff is **€8** (and your opponent's payoff is €24).
