Identification of Non-Equilibrium Beliefs in Games of Incomplete Information Using Experimental Data

Victor Aguirregabiria*
University of Toronto and CEPR

Erhao Xie*
Bank of Canada

January 31, 2020

Abstract

This paper studies the identification of players' preferences and beliefs in discrete choice games using experimental data. The experiment comprises a set of games that differ in their matrices of monetary payoffs. The researcher is interested in the identification of preferences (utility of money) and beliefs on the opponents' expected behavior, without imposing equilibrium restrictions or parametric assumptions on utility and belief functions. We show that the hypothesis of unbiased/rational beliefs is testable as long as the set of games in the experiment imply variation in monetary payoffs of other players, keeping the own monetary payoff constant. We present conditions for the full identification of utility and belief functions at the individual level – without restrictions on players' heterogeneity in preferences or beliefs. We apply our method to data from two experiments: a matching pennies game, and a public good game.

Keywords: Games of incomplete information; Biased beliefs; Strategic uncertainty; Coordination game.

JEL classifications: C57, C72.

Victor Aguirregabiria. Department of Economics, University of Toronto. 150 St. George Street, Toronto, Ontario, M5S 3G7.
Email: `victor.aguirregabiria@utoronto.ca`

Erhao Xie. Bank of Canada. 234 Wellington Street, Ottawa, Ontario, K1A 0G9.
Email: `erhaoxie@bankofcanada.ca`

*We would like to thank the editor and two anonymous referees for their helpful and constructive feedback. We have benefited from comments and helpful discussions with Martin Burda, Evan Calford, Robert Gazzale, Tanjim Hossain, Arvind Magesan, Yesim Orhun, David Pacini, Eduardo Souza-Rodrigues, Colin Stewart, Junichi Suzuki, Ruli Xiao, and especially Ryan Webb. We also thank comments from seminar participants at Bristol, Cornell, and Warwick, and at the *Behavioral Industrial Organization & Marketing Symposium* (Ann Arbor, 2019), the *Canadian Economic Association* conference (Banff, 2019), and the *Economics Science Association* North American meeting (Tucson, 2016). This research has been supported by the *Social Sciences and Humanities Research Council of Canada*.

1 Introduction

In games of incomplete information, players' behavior depends on their preferences and on their beliefs about the uncertain actions of other players. For instance, in a game of a public good provision, a player's decision depends on the reduction in utility from the contribution to the public good, but also on her belief about the probability that the other players contribute. Distinguishing between the influence of preferences and beliefs on observed behavior is important for our understanding of strategic and social interactions and for counterfactual predictions. Empirical researchers in different fields in economics are interested in the separate identification of preferences and beliefs using data from the realizations of games.

This paper studies the identification of players' preferences and beliefs in discrete choice games using experimental data. We show that – under simple conditions on the design of the experiment and weak assumptions on the primitives of the model – individuals' belief and utility functions are nonparametrically identified.

We consider a general model that encompasses most of the specifications in the existing applications in the experimental literature. Our model is a discrete choice game of incomplete information where a player's payoff comprises her utility of money plus an additive private information variable. The utility of money is an increasing function but is nonparametrically specified. A player's belief function is a probability distribution over the space of the other players' actions, and it is unrestricted. Players can be heterogeneous in their utility and belief functions. This heterogeneity is unrestricted. Our framework includes as particular cases, among others, the Quantal Response Equilibrium (QRE) by McKelvey and Palfrey (1995), the Cognitive Hierarchy model by Camerer, Ho, and Chong (2004), and the Level-k rationality model by Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007).

In applications of games of incomplete information using field data, a common assumption is that beliefs are in equilibrium, such that they correspond to the actual probability distribution of the other players' actions. This assumption can identify beliefs, but it does not identify players' preferences. Under equilibrium beliefs, a common condition for the identification of payoffs is that — for every player — there is a state variable that affects the payoff of a player but not of the other players (Bajari et al., 2010). Recent papers in this literature relax the assumption of rational beliefs: see Goldfarb and Xiao (2011), Brown, Camerer, and Lovallo (2013), Doraszelski, Lewis, and Pakes (2018), or Aguirregabiria and Magesan (2019), among others. These papers use different identification approaches.

In the experimental economics literature, researchers design laboratory experiments and generate experimental data to study players' behavior in games. From the point of view of identification, there are advantages of having data from a controlled experiment. The design of the experiment determines players' monetary payoffs such that the researcher knows these payoffs. The effect of other players' actions on a player's utility takes place only through monetary payoffs. Therefore, the identification of these effects is not different to the identification of the utility of money. This facilitates the identification of games using laboratory experiments.

Most of the experimental games literature has exploited this advantage together with three alternative identification approaches or restrictions. The first approach imposes the restriction that utility is equal to the monetary payment — the linear utility of money assumption. Then, beliefs are identified using choice data. Examples of this approach include Cheung and Friedman (1997), Nyarko and Schotter (2002), Kline (2018), and Melo, Pogorelskiy and Shum (2019).[1]

A second approach assumes that players form equilibrium beliefs. Under this assumption, a player's belief is equal to the actual probability distribution of other players' actions. Then, the utility function of money can be identified using the player's choice data. An example of this approach is Goeree, Holt and Palfrey (2003) who estimate each player's risk preference under the Quantal Response Equilibrium (QRE) framework (Mckelvey and Palfrey, 1995 and 1998).

In a third approach, experimental researchers focus on subjects' beliefs and measure them using an elicitation mechanism. Karni (2009), Offerman et al. (2009) and Hossain and Okui (2013) developed different methods to elicit players' beliefs regardless of their preferences towards risk.[2] As far as we know, there are no papers in the experimental literature that consider the nonparametric identification of both players' belief and utility functions without the assumption of equilibrium and without data on elicited beliefs.

Before we describe the main features of our identification approach, we first explain the importance of relaxing the assumption of equilibrium beliefs, the parametric restrictions on the utility functions, and the conditions for eliciting beliefs.

There are multiple reasons for players to have biased beliefs. Playing a Bayesian Nash Equilibrium strategy requires a player to figure out other players' equilibrium strategies and to integrate

[1] Cheung and Friedman (1997) estimate a belief learning process in a repeated game. Nyarko and Schotter (2002) compare the estimated beliefs with elicited beliefs. Kline (2018) estimates the solution/equilibrium concept. Melo, Pogorelskiy and Shum (2019) test the restrictions of the Quantal Response Equilibrium (QRE) by McKelvey and Palfrey (1995, 1998).

[2] Karni (2009) and Grether (1992) – following Smith (1961) – propose elicitation mechanisms based on randomized payments. See Schotter and Trevino (2014) and Schlag et al. (2015) for recent reviews of elicitation methods and their practical issues.

them over the distribution of private information. Human cognition limits – and imperfect knowledge on the distribution of private information – may prevent the construction of equilibrium beliefs. Multiple equilibria may also generate biased beliefs. A player may believe that the selected equilibrium is A, while other player may think that it is B. This strategic uncertainty has been studied by Van Huyk, Battalio, and Beil (1990), Crawford and Haller (1990), Morris and Shin (2002, 2004), and Heinemann, Nagel, and Ockenfels (2009), among others.

The linear utility of money assumption places strong restrictions on subjects' preferences, which are at odds with important empirical findings in the experimental literature. Harrison and Rutström (2008) show that risk aversion is prevalent even for the small scale of monetary payoffs in laboratory experiments. Kahneman and Tversky (1979) note that individuals may respond to loss more than to gains. The linear utility of money assumption also rules out heterogeneity across players in their marginal utility of money. Our framework specifies a player's utility as an unrestricted — nonparametric — function of her monetary payoff and can capture both risk preference and loss aversion. Our nonparametric specification of preferences has potential advantages over other methods in the experimental literature. Moreover, we allow the form of the utility function to be different for each subject in the experiment.[3]

Although there is substantial evidence showing that elicited beliefs are consistent with individuals' actions, there are practical issues.[4] Perhaps, the most serious problem is that the mechanism for the elicitation of beliefs can affect players' behavior in games. Schotter and Trevino (2014) denote this problem as Heisenberg problem – using an analogy from physics. A partial list of experimental papers illustrating this issue includes Nyarko and Schotter (2002), Guerra and Zizzo (2004), Ruström and Wilcox (2009), and Palfrey and Wang (2009). Our method does not need elicitation data and avoid these potential issues.[5]

Our identification results build on an exclusion restriction that can be generated by the re-

[3] In an influential study, Roth and Malouf (1979) proposed to linearize the utility function by assigning the payoff as the probability of winning a fixed reward. Ochs (1995) and Feltovich (2000), among others, applied this mechanism. Selten et al. (1999) and Goeree, Holt and Palfrey (2003) raised concerns about the validity of this mechanism. Another common method elicits players' risk preferences using a lottery choice with a known probability distribution. Heinemann, Nagel, and Ockenfels (2009) use this method. A third approach comprises estimating a parametric function for the utility of money, e.g., a CRRA utility function. This is the approach used by Goeree, Holt and Palfrey (2003). Misspecification of the parametric utility function can generate biases in beliefs. For instance, the researcher may spuriously conclude that players' beliefs are biased.

[4] To the best of our knowledge, the only paper that shows a contradictory evidence is Costa-Gomes and Weizsäcker (2008) who found significant discrepancy between elicited beliefs and beliefs inferred from players' actions.

[5] See table 2 in Schlag et al. (2015) for a comprehensive list of empirical evidence on this issue. Other practical issues related to eliciting beliefs include the hedging problem and the complexity of the methods. The empirical evidence on hedging is mixed. See section 3 in Schotter and Trevino (2014) or section 4 in Schlag et al. (2015) for more details.

searcher when she designs the experiment. Suppose that a subject in the experiment plays K different two-player games. The researcher designs the monetary payoff matrices in these games such that the payoff matrix of the column player varies across the K games but the payoff matrix of the row player is the same.[6] This variation across games does not affect the payoff matrix of the row player but it can affect her beliefs about the behavior of the column player.[7]

Under this exclusion restriction, the variation across the K games in the actions of the row player provides information about this player's beliefs. Following an argument similar to Aguirregabiria and Magesan (2019), we show that this exclusion restriction identifies an object that only depends on beliefs and not on preferences. This identification result can be used to test different hypotheses on beliefs, such as equilibrium beliefs, or rationalizability, or the validity of elicited beliefs, among others.

The complete identification of utility and belief functions requires additional restrictions. These conditions are substantially weaker than the linear utility assumption or equilibrium constraints. To achieve full identification of a two-player binary choice game, the researcher needs to restrict either the belief function or the utility function at only two values of the monetary payoffs. For instance, the researcher may assume that elicited beliefs are valid or that beliefs are unbiased at two of the K games. The selection of these two games is an important decision for the researcher. We discuss how our test on beliefs can inform this choice.

We apply our approach to estimate two types of games that have received much attention in the experimental economics literature: the matching pennies game in Goeree and Holt (2001) and the coordination game in Heinemann, Nagel, and Ockenfels (2009). In the matching pennies game, our estimation results cannot reject that players correctly predict other players' behaviors. In the coordination game, at the individual level, we reject the hypothesis of unbiased beliefs for a majority of subjects. We find substantial individual heterogeneity in preferences towards risk and a relationship between individuals' risk loving preferences and biased beliefs. We also use data on players' elicited beliefs and test the null hypothesis that individuals best respond to their self-reported beliefs. When subjects report beliefs on the aggregate behavior of all the other players, we cannot reject the null hypothesis that individuals best respond to their elicited beliefs. However, we reject this hypothesis when subjects report their beliefs about a randomly selected player.

[6] We find this experimental design of monetary payoffs in important studies such as Ochs (1995), Goeree and Holt (2001), and McKelvey, Palfrey and Weber (2000), among others.

[7] There are other experimental designs that can generate this type of exclusion restriction – variation in opponents' behavior without changing a player's payoff matrix. For instance, in a public good provision game, variation in the number of players can change the belief about the probability of successful coordination, but it has no direct impact on a player's payoff matrix. We find such a design in Heinemann, Nagel and Ockenfels (2009).

These empirical results emphasize the importance of relaxing restrictions of unbiased beliefs, linear utility, and individual homogeneity, as well as the different properties of alternative methods to elicit beliefs.

Our paper is mostly related to three recent papers in the literature. Aguirregabiria and Magesan (2019) also study the identification of non-equilibrium beliefs and apply a similar identification approach. However, they do not consider experimental data and study dynamic games. Melo, Pogorelskiy and Shum (2019) focus on testing the Quantal Response Equilibrium model. They consider a nonparametric specification for the distribution of the unobservable but impose the restriction of linear utility of money. Kline (2018) considers games of complete information, and his approach also imposes the restriction of linear utility.

Section 2 describes the model and its assumptions. In section 3, we explain the sampling framework, the experimental design, and the exclusion restriction implied by the design. Section 4 presents our identification results. In section 5, we describe estimation and testing methods based on our identification results. Section 6 presents the two experimental data sets that we use in the applications and our empirical results. We summarize and conclude in Section 7.

2 Model

Consider a two players binary choice game.[8] Each player has a different role. There is a row player (R) and a column player (C). Let a_R and a_C be the actions of the row and the column player, respectively. These actions belong to the choice set $\{0, 1\}$. Given players' actions, the payoff matrix of the game determines the monetary payoff for each player. We use $m_R(a_R, a_C)$ and $m_C(a_R, a_C)$ to represent the monetary payoffs for the row player and for the column player, respectively, when they take actions (a_R, a_C). The complete matrix of monetary payoffs is $\mathbf{m} \equiv \{\mathbf{m}_R, \mathbf{m}_C\} \equiv \{m_R(a_R, a_C), m_C(a_R, a_C) : (a_R, a_C) \in \{0, 1\}^2\}$. This payoff matrix and the role of each player in the game are common knowledge. Players take their actions simultaneously to maximize their respective expected utilities.

There is a population of individuals or subjects – that we index by $i \in \mathcal{I} = \{1, 2, ...\}$ – playing this game. Let $r(i) \in \{R, C\}$ be the role assignment that individual i has in the game. Each individual has a utility function of money: $\pi_i(m) : \mathbb{R} \to \mathbb{R}$. We do not restrict the individual heterogeneity in the utility function of money. The total utility of an individual also includes an additive component that depends on her own action, a_i, and on her role in the game. We represent

[8] Our model and identification results can be generalized to games with multiple players and multiple actions.

these additive utility components as $\varepsilon_{i,R}(a_i)$ and $\varepsilon_{i,C}(a_i)$. The total utility of subject i in the game when playing against individual j is:

$$\Pi_i(a_i,a_j) = \begin{cases} \pi_i\left(m_R(a_i,a_j)\right) + \varepsilon_{i,R}(a_i) & \text{if} \quad r(i) = R \\ \\ \pi_i\left(m_C(a_j,a_i)\right) + \varepsilon_{i,C}(a_i) & \text{if} \quad r(i) = C \end{cases} \tag{1}$$

Define $\varepsilon_{i,R} \equiv \varepsilon_{i,R}(0) - \varepsilon_{i,R}(1)$ and $\varepsilon_{i,C} \equiv \varepsilon_{i,C}(0) - \varepsilon_{i,C}(1)$. We assume that $(\varepsilon_{i,R}, \varepsilon_{i,C})$ is private information of individual i, and is independently distributed across subjects. Let $F_{i,R}$ and $F_{i,C}$ be the cumulative distribution functions for $\varepsilon_{i,R}$ and $\varepsilon_{i,C}$, respectively. These distribution functions can be different across individuals. There are two interpretations, not mutually exclusive, of the additive utilities ε_i. It can be an idiosyncratic non-pecuniary component of the utility and the individual can be fully aware of it. These variables can be also interpreted as *optimization errors*, along the line of the *Quantal Response Equilibrium (QRE)* concept proposed by Mckelvey and Palfrey (1995, 1998). A relevant difference between the utility of money and the ε components is that – for a given individual – $\varepsilon_{i,R}$ and $\varepsilon_{i,C}$ can vary randomly across different realizations of the game while $\pi_i(m)$ is a deterministic function of the amount of money.

Suppose, without of loss of generality, that individual i has the role of row player. Individual i does not know the utility function of money, $\pi_j(.)$, and the value of variable $\varepsilon_{j,R}$ for her opponent j, since they are private information. Therefore, even if player i is fully rational, she has uncertainty about the optimal choice of player j in the game. Each player has beliefs about the action that the other player will take. Let $B_{i,R}$ represent the subjective belief that, as row player, individual i has about the probability that the other player chooses action $a_j = 1$. Then, individual i's expected utility of action a_i is:

$$\Pi^e_{i,R}(a_i, B_{i,R}) = [1 - B_{i,R}] \;\; \pi_i\left(m_R(a_i,0)\right) + B_{i,R} \;\pi_i\left(m_R(a_i,1)\right) + \varepsilon_{i,R}(a_i) \tag{2}$$

Players maximize their expected utility. The best response of individual i is alternative $a_i = 1$ if $\Pi^e_{i,R}(1, B_{i,R}) \geq \Pi^e_{i,R}(0, B_{i,R})$, or equivalently,

$$\begin{aligned} &[1 - B_{i,R}] \;\; \pi_i\left(m_R(1,0)\right) + B_{i,R} \;\pi_i\left(m_R(1,1)\right) + \varepsilon_{i,R}(1) \\ &\geq [1 - B_{i,R}] \;\; \pi_i\left(m_R(0,0)\right) + B_{i,R} \;\pi_i\left(m_R(0,1)\right) + \varepsilon_{i,R}(0) \end{aligned} \tag{3}$$

Let $P_{i,R}$ be the probability that individual i, as row player, chooses action $a_i = 1$ given her belief $B_{i,R}$, her utility of money π_i, and the matrix of payoffs $\mathbf{m}$. Integrating the best response condition (3) over the private information $\varepsilon_{i,R}$, we obtain individual i's best response probability function:

$$P_{i,R} = F_{i,R}\left[\alpha_{i.R}(\mathbf{m}_R) + \beta_{i.R}(\mathbf{m}_R) \;\, B_{i,R}\right] \tag{4}$$

where $\alpha_{i.R}(\mathbf{m}_R) \equiv \pi_i(m_R(1,0)) - \pi_i(m_R(0,0))$, and $\beta_{i.R}(\mathbf{m}_R) \equiv \pi_i(m_R(1,1)) - \pi_i(m_R(0,1)) - \pi_i(m_R(1,0)) + \pi_i(m_R(0,0))$.

The model is a game and not a single-agent decision problem. This implies that $\beta_{i.R}(\mathbf{m}_R) \neq 0$. In the model presented above, we assume an individual's utility only depends on her own monetary reward: $\alpha_{i.R}(\mathbf{m}_R)$ and $\beta_{i,R}(\mathbf{m}_R)$ depend only on row player's monetary payoff matrix $\mathbf{m}_R$ but not on the column player's $\mathbf{m}_C$. At the end of this section 2, we present a more general model that allows for *other-regarding preferences*, e.g., preferences with altruism or envy. We have identification results for that extended model.

This framework includes as particular cases – or restricted versions – the models most often used in empirical applications of games. The following examples present common restricted versions of our framework.

Example 1. Quantal Response Equilibrium (QRE) – or Bayesian Nash Equilibrium (BNE) – with homogeneous utilities. Suppose that all the individuals have the same utility function of money, and the same probability distribution for the non-pecuniary component of the utility. That is, we have that $\pi_i = \pi$, $F_{i,R} = F_R$, and $F_{i,C} = F_C$ for every $i \in \mathcal{I}$. Suppose also that π, F_R, and F_C are common knowledge. Then, a BNE in this game can be represented as a pair of choice probabilities (P_R, P_C) that satisfies the following conditions:

$$\begin{aligned} P_R &= F_R\left[\alpha_R(\mathbf{m}_R) + \beta_R(\mathbf{m}_R)\ P_C\right] \\ P_C &= F_C\left[\alpha_C(\mathbf{m}_C) + \beta_C(\mathbf{m}_C)\ P_R\right] \end{aligned} \tag{5}$$

where $\alpha_R(\mathbf{m}_R)$ and $\beta_R(\mathbf{m}_R)$ have the same definition as above, but now they do not vary across individuals.

Under BNE, all the individuals in the role of row player (column player) have the same belief about the choice probability of a column player (row player) and this belief is equal to the actual choice probability of a column player (row player). That is, $B_{i,R} = P_C$ and $B_{i,C} = P_R$ for any individual i. It is clear that this model imposes substantial restrictions on our framework: homogeneous preferences and beliefs, and equilibrium beliefs.

Examples of empirical applications using experimental data are, among others, McKelvey and Palfrey (1995) – using a variety of experimental datasets – and Anderson, Goeree, and Holt (2001) – for a coordination game. ■

Example 2. BNE with heterogeneous common knowledge utilities. Suppose that the population is relatively small such that individuals know each other and they know the utility function of every subject in the population. Utility functions π_i and distribution functions $F_{\varepsilon,i,R}$

and $F_{\varepsilon,i,C}$ are unrestricted but they are assumed to be common knowledge. Consider the game where individual i is the row player and individual j is the column player. A BNE in this game can be represented as a pair of choice probabilities $(P_{i,R}, P_{j,C})$ that satisfies the equations:

$$\begin{array}{rcl} P_{i,R} & = & F_{i,R}\left[\alpha_{i.R}(\mathbf{m}_R) + \beta_{i.R}(\mathbf{m}_R)\ P_{j,C}\right] \\ \\ P_{j,C} & = & F_{j,C}\left[\alpha_{j.C}(\mathbf{m}_C) + \beta_{j.C}(\mathbf{m}_C)\ P_{i,R}\right] \end{array} \tag{6}$$

Though this model allows for a general form of individual heterogeneity in the utilities, it imposes the restriction of common knowledge utilities and equilibrium beliefs. That is, $B_{i,R} = P_{j,C}$ and $B_{j,C} = P_{i,R}$.

This model has been used in applications of discrete choice games of oligopoly competition using field data. Some examples are Ellickson and Misra (2008), Sweeting (2009), or Aguirregabiria and Ho (2012), among others. See Bajari, Hong and Nekipelov (2013) for a survey. ■

Example 3. Heterogeneous QRE. Consider a similar model as in Example 2, but now the functions $\{\pi_i,\ F_{i,R},\ F_{i,C}\}$ are private information of each individual. Suppose that there is a finite number of individual types – L – for the form of these functions. Let ℓ be the index that represents a type; let λ_ℓ be the proportion of individuals of type-ℓ; and let $\{\pi^{(\ell)},\ F_R^{(\ell)},\ F_C^{(\ell)}\}$ represent these functions for type ℓ. The probability distribution $\lambda = \{\lambda_\ell : \ell = 1, 2, ..., L\}$ is common knowledge. A BNE in this game can be represented as L pairs of choice probabilities $(P_R^{(\ell)}, P_C^{(\ell)})$ – one for each type – satisfying the following conditions:

$$\begin{array}{rcl} P_R^{(\ell)} & = & F_R^{(\ell)}\left[\alpha_R^{(\ell)}(\mathbf{m}_R) + \beta_R^{(\ell)}(\mathbf{m}_R)\ \sum_{\ell'=1}^{L}\lambda_{\ell'}\ P_C^{(\ell')}\right] \\ \\ P_C^{(\ell)} & = & F_C^{(\ell)}\left[\alpha_C^{(\ell)}(\mathbf{m}_C) + \beta_C^{(\ell)}(\mathbf{m}_C)\ \sum_{\ell'=1}^{L}\lambda_{\ell'}\ P_R^{(\ell')}\right] \end{array} \tag{7}$$

This model imposes several restrictions with respect to our general framework. First, it implies that – in the population of individuals – there are only L different types of choice probabilities. Second, it imposes the restriction that every individual has the same beliefs: for every individual i, $B_{i,R} = B_R$ and $B_{i,C} = B_C$. Finally, it imposes the restriction that beliefs are in equilibrium: $B_R = \sum_{\ell=1}^{L}\lambda_{\ell'}\ P_C^{(\ell)}$ and $B_C = \sum_{\ell=1}^{L}\lambda_\ell P_R^{(\ell)}$.

Rogers, Palfrey, and Camerer (2009) introduced this heterogeneous QRE model and applied it to 17 different experimental datasets. ■

Example 4. Cognitive Hierarchy (CH) and Level-k Rationality.[9] In these models, players are heterogeneous in their beliefs. There is a finite number of belief types: for every individual i,

[9] Here we present a version of Level-K and CH models where individuals have homogeneous preferences. It is possible to extend these models to allow for preferences that are heterogeneous (see Rogers, Palfrey, and Camerer, 2009). For a survey of papers in this field, see Crawford, Costa-Gomes and Iriberri (2013).

$(B_{i,R}, B_{i,C}) \in \{(B_R^{(k)}, B_C^{(k)}) : k = 1, 2, ..., K\}$. These types correspond to different levels of strategic sophistication. Beliefs have a hierarchical structure. The CH model (Camerer et al., 2004) and the Level-k model (Costa-Gomes and Crawford, 2006; Crawford and Iriberri, 2007) propose similar hierarchical structures of beliefs. A type-0 individual has some beliefs $(B_R^{(0)}, B_C^{(0)})$. A type-1 individual believes that all the other individuals are type-0. Therefore, she has beliefs:

$$\begin{aligned} B_R^{(1)} &= F_C\left[\alpha_C(\mathbf{m}_C) + \beta_C(\mathbf{m}_C)\ B_C^{(0)}\right] \\ B_C^{(1)} &= F_R\left[\alpha_R(\mathbf{m}_R) + \beta_R(\mathbf{m}_R)\ B_R^{(0)}\right] \end{aligned} \tag{8}$$

The CH and the Level-k models differ in the specification of beliefs for types greater than one. In the Level-k model, a type-k player believes that the other players are type-(k-1). In the CH model, a type-k player believes that the other players come from a probability distribution over types 0 to (k-1). Therefore, for the CH model, beliefs are:

$$\begin{aligned} B_R^{(k)} &= \textstyle\sum_{k'=0}^{k-1} \lambda_{k'}^{(k)}\ F_C\left[\alpha_C(\mathbf{m}_C) + \beta_C(\mathbf{m}_C)\ B_C^{(k')}\right] \\ B_C^{(k)} &= \textstyle\sum_{k'=0}^{k-1} \lambda_{k'}^{(k)}\ F_R\left[\alpha_R(\mathbf{m}_R) + \beta_R(\mathbf{m}_R)\ B_R^{(k')}\right] \end{aligned} \tag{9}$$

where – for every value of k – the vector $\{\lambda_{k'}^{(k)} : k' = 0, 1, ..., k-1\}$ is a probability distribution. Compared to our general framework, this model imposes important restrictions on individuals' beliefs. There is a finite number K of types. Most importantly, beliefs satisfy the equilibrium hierarchical structure described by equations (8) and (9).

These models have received much attention in experimental economics, with many applications using data from lab experiments. There are also applications using field data, such as Goldfarb and Xiao (2011) and Brown, Camerer, and Lovallo (2013). ■

Our framework relaxes important restrictions in the models presented in Examples 1 to 4. We do not restrict beliefs to be in equilibrium or to satisfy a hierarchical structure. The payoff functions π_i can differ from the monetary payoff. We do not restrict the form of the heterogeneity – in preferences or beliefs – between the subjects in the experiment.

An extension: Other-regarding preferences (e.g., altruism / envy). Though most of the analysis in this paper focuses on the model presented above, we also present identification results for a more general model. This extended model allows for a subject's utility to depend on the amount of money received by the other player. The utility of subject i in the game depends on her own amount of money and on the amount of money obtained by the other player, $\pi_i(m_i, m_j)$.

That is, if subject i is a row player:

$$\Pi_i(a_i, a_j) = \pi_i \left(m_R(a_i, a_j),\, m_C(a_i, a_j) \right) + \varepsilon_{i,R}(a_i) \tag{10}$$

where $\pi_i(m_i, m_j) : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a function that is increasing in the own monetary payoff m_i, and it can be increasing – under altruism – or decreasing – under envy – in the other player's payoff, m_j.

The general form of the utility function in equation (10) includes – as particular cases – the specifications in most papers in experimental economics allowing for other-regarding preferences. For instance, in a very influential paper, Fehr and Schmidt (1999) study preferences for inequality aversion and propose a utility function $\pi_i(m_i, m_j)$ that is equal to $m_i - \alpha_i(m_i - m_j)$ if $m_i \geq m_j$ and to $m_i - \beta_i \ (m_j - m_i)$ if $m_j \geq m_i$, where α_i and β_i are parameters.

3 Experimental design and data generating process

The experimental researcher has panel data of N individuals where each individual plays the game T different times or rounds, indexed by $t \in \{1,\, 2,\, ...,\, T\}$. The dataset has NT observations $\{d_{it},$ $a_{it} : i = 1,\, 2,\, ...,\, N;\ t = 1,\, 2,\, ...,\, T\}$, where d_{it} is the *treatment* received by subject i in her game t, and a_{it} is her action in that game. We describe below the different *treatments* and how players are assigned to treatments.

The researcher chooses M different matrices of monetary payoffs. Let $\mathcal{M}$ be the set of monetary payoff matrices in the experiment; and $\mathbf{m} = (\mathbf{m}_R,\, \mathbf{m}_C) \in \mathcal{M}$ is an element in this set. A *treatment* in this experiment is a pair $(\mathbf{m}$,$r)$, where $\mathbf{m}$ is a payoff matrix in the set $\mathcal{M}$ and $r \in \{R, C\}$ is the subject's role in that game – either row or column player. Therefore, there are $2M$ possible treatments. At each round t, the researcher randomly assigns each of the N subjects to one of the $2M$ treatments. The random allocation of players to treatments is anonymous: each subject does not have any information about who is the other subject she is playing against. Once subjects have been allocated to treatments, they play their respective games. We use the categorical variable $d_{it} \in \mathcal{M} \times \{R,\, C\}$ to represent subject i' treatment at round t. The binary variable $a_{it} \in \{0, 1\}$ represents the subject's actual choice in the game.

Assumption 1 establishes that – for every individual – the utility function of money and the probability distribution of non-pecuniary payoffs are invariant over the T rounds of the experiment.

Assumption 1. *An individual's utility function of money, π_i, and the distribution functions $F_{i,R}$ and $F_{i,C}$ are invariant over the T rounds that an individual plays the game.* ■

Though the conditions in Assumption 1 are standard, there may be applications of experimental economics that violate this assumption.

For instance, suppose that a subject maximizes the expected utility of the sum of monetary rewards over the T rounds — instead of maximizing her expected utility in each of the T games. This subject has an incentive to hedge her payoff across rounds. Moreover, her utility function could change over the T rounds as she accumulates monetary reward due to – for instance – endowment effects. This behavior violates Assumption 1. Experimental economists have been well aware of this problem (e.g., Schotter and Trevino, 2014). One way to avoid this issue is to design the experiment such that the reward that a subject receives is determined — ex post — by randomly selecting her realized payoff in one of the T rounds. This design eliminates hedging incentives and endowment effects, such that we can treat each round as an independent decision. The experimental designs in the applications we present in this paper incorporate this feature to avoid hedging incentives and endowment effects.

Another departure from Assumption 1 occurs when individuals have other-regarding preferences but the specification of the utility function does not take them into account. In that case, variables $(\varepsilon_{it,R}, \varepsilon_{it,C})$ include individual i's preference for the monetary payoff of the other player. Since these monetary payoffs vary over treatments, the distribution functions of $(\varepsilon_{it,R}, \varepsilon_{it,C})$ do as well. To deal with this potential problem, we also present results on the identification of beliefs that apply to the model with *other-regarding preferences* of the general form $\pi_i(m_i, m_j)$.

Assumption 2. *The treatment variable d_{it} and the non-pecuniary utility components $(\varepsilon_{it,R},\, \varepsilon_{it,C})$ are independently distributed.* ■

Given the random allocation of individuals to treatments and the non-pecuniary nature of the variables $(\varepsilon_{it,R},\, \varepsilon_{it,C})$, Assumption 2 seems plausible.

Assumption 3 establishes that the beliefs function is invariant over rounds/games with the same matrix of monetary payoffs – with the same treatment.

Assumption 3. *An individual's belief in round t depends on her treatment $(\mathbf{m},r)$ in that round, but conditional on the same treatment it is invariant over rounds: $B_{it,r}(\mathbf{m}) = B_{it',r'}(\mathbf{m}')$ for every two rounds t and t' with $(\mathbf{m},r) = (\mathbf{m}',r')$.* ■

Assumption 3 imposes the restriction that subjects do not learn over the T rounds. Such an assumption is more plausible if subjects do not receive any information on the outcome of the game after each round. The experimental designs in the applications that we present in this paper satisfy this condition to avoid learning.

Most of our main identification results use the assumption that the researcher knows the distribution of non-pecuniary components $(\varepsilon_{it,R}, \varepsilon_{it,C})$ — up to location and scale parameters that are individual-specific. This is our Assumption 4, and it is common in applications of discrete choice games.

Assumption 4. *For $r \in \{R, C\}$, define $\mu_{i,r} \equiv \mathbb{E}(\varepsilon_{it,r})$ and $\sigma^2_{i,r} \equiv \mathbb{V}(\varepsilon_{it,r})$. (A) The standardized variable $(\varepsilon_{it,r} - \mu_{i,r})/\sigma_{i,r}$ has CDF $F(.)$ that is the same for every individual and is strictly increasing over the real line. (B) The distribution function $F(.)$ is known to the researcher.* ■

Assumption 4 does not restrict the location and scale parameters $\mu_{i,r}$ and $\sigma_{i,r}$ which are unknown to the researcher and can vary over subjects unrestrictedly.

The location parameter $\mu_{i,r}$ deserves some explanation. This parameter represents individual i's non-pecuniary preference for action 0 compared to action 1. For instance, in a public good game, this parameter represents an individual's taste to contribute to the public good for non-pecuniary reasons. Many experimental studies under the QRE framework assume that this parameter is zero, e.g., McKelvey and Palfrey (1995), Goeree, Holt and Palfrey (2005), and Melo, Pogorelskiy and Shum (2019), among others. The work of Nyarko and Schotter (2002) is one of the few experimental studies that estimates this parameter — under the assumption of linear utility of money. Our identification of beliefs and our tests do not impose any restriction on the parameter $\mu_{i,r}$.

Assumption 5 establishes a condition on the set $\mathcal{M}$ of monetary payoff matrices in the experiment. This condition plays a fundamental role in our identification results. For notational simplicity, we present this assumption focusing on the identification of beliefs for the row player, R. However, we can consider a symmetric version of the assumption that applies to the identification of beliefs for the column player, C.

Assumption 5. *The set $\mathcal{M}$ of monetary payoff matrices in the experiment contains at least two matrices, say $\mathbf{m}^1$ and $\mathbf{m}^2$, such that: (A) the matrix for the row player does not vary over the two treatments – $\mathbf{m}^1_R = \mathbf{m}^2_R$ – but the matrix for the column player varies – $\mathbf{m}^1_C \neq \mathbf{m}^2_C$; (B) for an individual i, her conditional choice probability as a row player varies across the two treatments – that is, $P_{i,R}(\mathbf{m}^1) \neq P_{i,R}(\mathbf{m}^2)$.* ■

Under Assumption 5(A), the experimental design generates a particular variation in monetary payoffs across treatments: the payoff matrix of the column player C varies while the payoff matrix of the row player R remains constant. More than assumption, 5(A) is a condition that experimental design should satisfy. We show below that this condition provides an *exclusion restriction* that identifies the beliefs of individual i as a row player, $B_{i,R}$.

Assumption 5(B) is a rank condition for identification. In the next section, we establish that the conditional choice probabilities $P_{i,R}(\mathbf{m}^1)$ and $P_{i,R}(\mathbf{m}^2)$ are identified from the data under mild conditions (see section 4.1 below). Therefore, Assumption 5(B) is testable from the data. This assumption can be also interpreted as an implication of condition 5(A) under *Rationalizability*. That is, player R knows that player C maximizes expected utility given beliefs. Since player C's payoff matrix varies across treatments d^1 and d^2, then player R's beliefs about player C's behavior also vary between these two treatments. Given that player $R's$ own monetary payoff matrix did not change, her actual behaviors should be different as long as her behavior depends on her beliefs.

Though the focus of the paper is on the model where an individual's utility depends only on her own monetary payoff, we also present identification results for the model with altruism/envy that we have described at the end of section 2. For the identification of beliefs in this model we need additional conditions on the design of the experiment. Assumption 5* presents these conditions.

Assumption 5*. *The set $\mathcal{M}$ of monetary payoff matrices in the experiment contains at least two matrices, say $\mathbf{m}^1$ and $\mathbf{m}^2$, that satisfy the following conditions: (A) for t=1,2, $m_R^t(0,0) = m_R^t(1,0)$ and $m_C^t(0,0) = m_C^t(1,0)$; (B) $m_R^t(0,1)$, $m_R^t(1,1)$, $m_C^t(0,1)$, and $m_C^t(1,1)$ do not vary over t=1,2; (C) for an individual i, her conditional choice probability as a row player varies across the two treatments – that is, $P_{i,R}(\mathbf{m}^1) \neq P_{i,R}(\mathbf{m}^2)$.* ■

Example 5. The following matrices provide an example that satisfies Assumption 5*.

Matrix $\mathbf{m}^1$

		Player C	
		$a_C = 0$	$a_C = 1$
Player R	$a_R = 0$	[3 , 6]	[2 , 8]
	$a_R = 1$	[3 , 6]	[4 , 7]

Matrix $\mathbf{m}^2$

		Player C	
		$a_C = 0$	$a_C = 1$
Player R	$a_R = 0$	[5 , 9]	[2 , 8]
	$a_R = 1$	[5 , 9]	[4 , 7]

This example shows that Assumption 5* still allows for variation in the matrices across treatments that can generate substantial changes in players' behavior. For instance, outcome $(a_R, a_C) = (1, 1)$ is a pure strategy Nash equilibrium in the game with matrix $\mathbf{m}^1$ but it is not a Nash equilibrium in the game with matrix $\mathbf{m}^2$. ■

For the asymptotic distribution of the estimators and tests that we propose, we consider panel data where the number of rounds T is large. This sampling framework corresponds to most empirical applications of discrete choice games of oligopoly competition in empirical IO (Berry and Tamer, 2006). In that setting, N is the number of firms in an industry, and T is the number of geographic markets where these firms compete. This sampling framework is not uncommon in empirical applications in experimental economics. In many experiments, T can be larger than 50 and some recent experiments have T about 200, for instance, Selten and Chmura (2008) and Chmura, Goerg and Selten (2012). Furthermore, in many of these experiments, subjects do not receive any feedback after each treatment to avoid learning.

Assumption 6. *(A) The experiment is such T is large – e.g., larger than 20 – and each subject i is observed many times playing at least two treatments, i.e., choice probabilities can be estimated consistently for each subject i. (B) After each treatment, players do not receive any information on the outcome of the game.* ■

4 Identification

Under Assumptions 1 to 5, we show the identification of each individual's beliefs. This identification result can be used to test different belief restrictions, such as equilibrium beliefs or the validity of elicited beliefs. Then, we present additional conditions for the full nonparametric identification of preferences and beliefs.

Suppose that the treatment of individual i at round t is $d_{it} = (\mathbf{m},r)$. Given this treatment, her belief probability is $B_{i,r}(\mathbf{m})$ and her best response *Conditional Choice Probability (CCP)* function is:

$$P_{i,r}(\mathbf{m}) = F\left[\alpha_{i.r}(\mathbf{m}_r) + \beta_{i.r}(\mathbf{m}_r)\ B_{i,r}(\mathbf{m})\right] \tag{11}$$

where $\alpha_{i,r}(\mathbf{m}_r) \equiv [\pi_i\left(m_r(1,0)\right) - \pi_i\left(m_r(0,0)\right) - \mu_{i,r}]/\sigma_{i,r}$ and $\beta_{i,r}(\mathbf{m}_r) \equiv [\pi_i\left(m_R(1,1)\right) - \pi_i\left(m_R(0,1)\right) - \pi_i\left(m_R(1,0)\right) + \pi_i\left(m_R(0,0)\right)]/\sigma_{i,r}$. Let $F^{-1}(\cdot)$ be the inverse function of $F(\cdot)$. This inverse function exists because the strict monotonicity of CDF, as stated in Assumption 4. Consequently, the inversion of equation (11) is written as:

$$F^{-1}[P_{i,r}(\mathbf{m})] = \alpha_{i,r}(\mathbf{m}_r) + \beta_{i,r}(\mathbf{m}_r)\ B_{i,r}(\mathbf{m}) \tag{12}$$

Equation (12) is the key restriction of the model that we use to identify and estimate individuals' utility and belief functions. As implied by Assumption 4, F^{-1} is known to the researcher. Moreover, $P_{i,r}(\mathbf{m})$ can be consistently estimated using data from the experiment. Therefore, the left hand side of equation (12) is known to the researcher.

Let $\boldsymbol{\pi}_i$ be the vector of utility parameters for individual i in the experiment, that is, $\boldsymbol{\pi}_i \equiv \{\alpha_{i,r}(m_r(a_R, a_C)), \beta_{i,r}(m_r(a_R, a_C)) : (a_R, a_C) \in \{0,1\}^2$ and $(\mathbf{m},r) \in \mathcal{M} \times \{R, C\}\}$. Similarly, let $\mathbf{B}_i$ be the vector of beliefs for individual i, that is, $\mathbf{B}_i \equiv \{B_{i,r}(\mathbf{m}) : (\mathbf{m},r) \in \mathcal{M} \times \{R, C\}\}$. The researcher is interested in using the experimental data to identify $\boldsymbol{\pi}_i$ and $\mathbf{B}_i$ for each of the N individuals in the sample.

4.1 Identification of beliefs and tests

Let d^1 and d^2 be the two treatments in Assumption 5. Let $\mathcal{D}_{d^1}$ be the subset of treatments d where the row player has the same payoff matrix as in treatment d^1: $\mathcal{D}_{d^1} \equiv \{d = (\mathbf{m}^d, R) : \mathbf{m}^d_R = \mathbf{m}^1_R\}$. By Assumption 5(A), $d^1, d^2 \in \mathcal{D}_{d^1}$ such that this set contains at least two elements. For any treatment $d \in \mathcal{D}_{d^1}$, we have that $\alpha_{i,R}(\mathbf{m}^d_R) = \alpha_{i,R}(\mathbf{m}^1_R)$ and $\beta_{i,R}(\mathbf{m}^d_R) = \beta_{i,R}(\mathbf{m}^1_R)$. This implies that, for any $d \in \mathcal{D}_{d^1}$,

$$F^{-1}[P_{i,R}(\mathbf{m}^d)] - F^{-1}[P_{i,R}(\mathbf{m}^1)] = \beta_{i,R}(\mathbf{m}^1_R)\ [B_{i,R}(\mathbf{m}^d) - B_{i,R}(\mathbf{m}^1)] \tag{13}$$

Intuitively, individual i faces the same monetary payoff matrix in d^1 and d; therefore, the variation of her choice probabilities provides information on her adjustment of beliefs, as represented by $B_{i,R}(\mathbf{m}^d) - B_{i,R}(\mathbf{m}^1)$.

Assumption 5(B) and the strict monotonicity of F imply that $F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]$ is different to zero. Therefore, for any treatment $d \in \mathcal{D}_{d^1}$ we have:

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]}{F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]} = \frac{B_{i,R}(\mathbf{m}^d) - B_{i,R}(\mathbf{m}^1)}{B_{i,R}(\mathbf{m}^2) - B_{i,R}(\mathbf{m}^1)} \tag{14}$$

All the terms in the left hand side of equation (14) are known to the researcher. This expression shows that, under Assumptions 1-5, the observed behavior of subject i identifies her beliefs ratio $[B_{i,R}(\mathbf{m}^d) - B_{i,R}(\mathbf{m}^1)]/[B_{i,R}(\mathbf{m}^2) - B_{i,R}(\mathbf{m}^1)]$ for any treatment $d \in \mathcal{D}_{d^1}$. Observed behavior identifies an object that depends only on beliefs and not on preferences. It also implies that different hypotheses on beliefs are testable.

Proposition 1. *Suppose that Assumptions 1 to 5 hold. (A) Equation (14) characterizes the identified set of individual i's beliefs in treatments $d \in \mathcal{D}_{d^1}$. (B) Consider the hypothesis that $B_{i,R}(\mathbf{m}^d) = B^*(\mathbf{m}^d)$ for $d \in \mathcal{D}_{d^1}$, where the values $B^*(\mathbf{m}^d)$ are known to the researcher. If the set $\mathcal{D}_{d^1}$ contains at least three treatments, this hypothesis can be tested by testing the following restrictions:*

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]}{F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]} = \frac{B^*(\mathbf{m}^d) - B^*(\mathbf{m}^1)}{B^*(\mathbf{m}^2) - B^*(\mathbf{m}^1)} \quad \text{for } d \in \mathcal{D}_{d^1}. \qquad \blacksquare \tag{15}$$

The testable implications in equation (15) have power. There exist invalid beliefs – i.e., with $B_{i,R}(\mathbf{m}) \neq B^{*}(\mathbf{m})$ – that can satisfy the restrictions in equation (15). However, within the space of possible beliefs, the set of beliefs $B_{i,R}(\mathbf{m}) \neq B^{*}(\mathbf{m})$ that satisfy equation (15) has measure zero. Importantly, this set shrinks when the number of treatment in $\mathcal{D}_{d^1}$ increases.[10]

Researchers can be interested in testing the null hypothesis of equilibrium beliefs. Under BNE and the condition that the opponent's identity is anonymous, individual i's beliefs – as a row player – are in equilibrium if $B_{i,R}(\mathbf{m}) = \mathbb{E}_{\pi_j}\left(P_{j,C}(\mathbf{m})\right)$ where the expectation $\mathbb{E}_{\pi_j}(.)$ is taken over the distribution of subjects' preferences π_j. Subject j's choice probabilities $P_{j,C}$ are identified, and the expectation $\mathbb{E}_{\pi_j}\left(P_{j,C}(\mathbf{m})\right)$ can be estimated consistently using $(1/N-1)\sum_{j\neq i} P_{j,C}(\mathbf{m})$. Therefore, we can test the null hypothesis of equilibrium beliefs for subject i by testing the restrictions:

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^{d})] - F^{-1}[P_{i,R}(\mathbf{m}^{1})]}{F^{-1}[P_{i,R}(\mathbf{m}^{2})] - F^{-1}[P_{i,R}(\mathbf{m}^{1})]} = \frac{\sum_{j\neq i} P_{j,C}(\mathbf{m}^{d}) - P_{j,C}(\mathbf{m}^{1})}{\sum_{j\neq i} P_{j,C}(\mathbf{m}^{2}) - P_{j,C}(\mathbf{m}^{1})} \tag{16}$$

A researcher may be also interested in testing the validity of the beliefs $\{B_R^{(k)} : k = 0, 1, ..., K\}$ obtained after the estimation of the Cognitive Hierarchy (CH) or the Level-k Rationality models presented in Example 4 above. There are two null hypothesis that the researcher may test. A first null hypothesis is individual i is type k. This corresponds to test the restrictions in equation (15) with $B^{*}(\mathbf{m}) = B_R^{(k)}(\mathbf{m})$. A second null hypothesis is that individual i's beliefs function belongs to one of the K types – $B_{i,R}(\mathbf{m}) \in \{B_R^{(k)}(\mathbf{m}) : k = 0, 1, ..., K\}$. This is equivalent to testing whether CH (or Level-K) is a valid model to represent the beliefs of individual i.

In this context, it is important to note that Level-1 rationality may not be testable. Under Level-1, a player has arbitrary beliefs that may not depend on the payoff matrix of the other players. This implies that the (rank) identification condition in Assumption 5(B) may not hold. That is, under Level-1, a player's conditional choice probability may not vary across the two treatments such that $P_{i,R}(\mathbf{m}^1) = P_{i,R}(\mathbf{m}^2)$.

The researcher can also test the validity of elicited beliefs. In that case, $B^{*}(\mathbf{m})$ is the subject's reported belief.

The tests based on equation (15) require that the set $\mathcal{D}_{d^1}$ has at least three treatments. A particular structure of the monetary payoff matrix – that we find in some experiments – implies testable restrictions on beliefs even with only two treatments in $\mathcal{D}_{d^1}$. Suppose that the matrix of monetary payoffs of the row player is symmetric and diagonal-constant – a Toeplitz matrix. That is, $m_R(0,0) = m_R(1,1)$ and $m_R(0,1) = m_R(1,0)$. For instance, this is form of the payoff matrix

[10] Aguirregabiria and Magesan (2019) present a Monte Carlo study that shows the power of a similar test as ours, in context of dynamic games.

in a matching pennies game. Under this condition, and the additional restriction $\mu_{i,R} = 0$, we have that $\beta_{i,R}(\mathbf{m}_R) = -2\ \alpha_{i,R}(\mathbf{m}_R)$ and equation (12) becomes $F^{-1}[P_{i,R}(\mathbf{m})] = \alpha_{i,R}(\mathbf{m}_R)[1 - 2\ B_{i,R}(\mathbf{m})]$. Under Assumption 5, for any treatment d in $\mathcal{D}_{d^1}$ we have that:

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)]}{F^{-1}[P_{i,R}(\mathbf{m}^1)]} = \frac{1 - 2\ B_{i,R}(\mathbf{m}^d)}{1 - 2\ B_{i,R}(\mathbf{m}^1)} \tag{17}$$

provided that $F^{-1}[P_{i,R}(\mathbf{m}^1)] \neq 0$. This condition characterizes the identified set and testable restrictions of beliefs for Toeplitz monetary payoff matrix. It is summarized by Proposition 2.

Proposition 2. *Suppose that Assumptions 1 to 5 hold and: (i) the matrix of monetary payoffs of the row player is symmetric and diagonal-constant (Toeplitz matrix); and (ii)* $\mu_{i,R} = 0$. *(A) Equation (17) characterizes the identified set of individual i's beliefs in treatments* $d \in \mathcal{D}_{d^1}$. *(B) Let* $B^*(.)$ *be a beliefs function, and consider the hypothesis represented by the restrictions* $B_{i,R}(\mathbf{m}^d) = B^*(\mathbf{m}^d)$ *for* $d \in \mathcal{D}_{d^1}$. *If the set* $\mathcal{D}_{d^1}$ *contains at least two treatments, this hypothesis can be tested by testing the following restrictions:*

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)]}{F^{-1}[P_{i,R}(\mathbf{m}^1)]} = \frac{1 - 2\ B^*(\mathbf{m}^d)}{1 - 2\ B^*(\mathbf{m}^1)} \quad \textit{for } d \in \mathcal{D}_{d^1}. \qquad \blacksquare \tag{18}$$

Proposition 2 requires the additional restriction $\mu_{i,R} = 0$. This restriction is not innocuous. In some experiments, individuals can have a non-pecuniary preference to choose a particular action – i.e., $\mu_{i,R}$ is different to zero. While the identification result in Proposition 1 is robust to these non-pecuniary biases, this is not the case for Proposition 2.

4.2 Identification of beliefs: Extensions

In this section, we present two extensions of the previous identification result. First, we show the identification of beliefs in a model when subjects may have preferences with altruism/envy. Second, we establish the identification of beliefs in a model with more than two players that corresponds to the game in our public good application.

4.2.1 Identification of beliefs when preferences have altruism/envy

The identification results in Propositions 1 and 2 rely on the variation in a subject's behavior when the payoff matrix of the other player changes. A concern with this approach is that the observed variation in the subject's behavior could capture not only changes in her beliefs but also in her utility if this depends on the monetary payoff of the other player, i.e., utility with altruism or envy.

Consider the model where the utility function for the subject i as a row player is $\pi_i(m_R(a_i, a_j), m_C(a_i, a_j))$. Under Assumptions 1-4, the best response of subject i in the model implies the equation:

$$F^{-1}[P_{i,R}(\mathbf{m})] = \alpha_{i,R}(\mathbf{m}) + \beta_{i,R}(\mathbf{m})\ B_{i,R}(\mathbf{m}) \tag{19}$$

But now we have that, $\alpha_{i,R}(\mathbf{m}) \equiv \pi_i(m_R(1,0),\ m_C(1,0)) - \pi_i(m_R(0,0),\ m_C(0,0))$, and $\beta_{i,R}(\mathbf{m}) \equiv \pi_i(m_R(1,1),\ m_C(1,1)) - \pi_i(m_R(0,1),\ m_C(0,1)) - \pi_i(m_R(1,0),\ m_C(1,0)) + \pi_i(m_R(0,0),\ m_C(0,0))$. Under Assumptions 1-5, for any treatment d in $\mathcal{D}_{d^1}$ we have that:

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]}{F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]} = \frac{\alpha_{i,R}(\mathbf{m}^d) + \beta_{i,R}(\mathbf{m}^d)B_{i,R}(\mathbf{m}^d) - \alpha_{i,R}(\mathbf{m}^1) - \beta_{i,R}(\mathbf{m}^1)B_{i,R}(\mathbf{m}^1)}{\alpha_{i,R}(\mathbf{m}^2) + \beta_{i,R}(\mathbf{m}^2)B_{i,R}(\mathbf{m}^2) - \alpha_{i,R}(\mathbf{m}^1) - \beta_{i,R}(\mathbf{m}^1)B_{i,R}(\mathbf{m}^1)} \tag{20}$$

Without further restrictions, $\alpha_{i,R}(\mathbf{m}^d) \neq \alpha_{i,R}(\mathbf{m}^1)$ and $\beta_{i,R}(\mathbf{m}^d) \neq \beta_{i,R}(\mathbf{m}^1)$ such that this equation does not identify beliefs. More specifically,

$$\begin{array}{rcl} \alpha_{i,R}(\mathbf{m}^d) - \alpha_{i,R}(\mathbf{m}^1) & = & \left[\pi_i(m_R^d(1,0), m_C^d(1,0)) - \pi_i(m_R^d(0,0), m_C^d(0,0))\right] \\ & - & \left[\pi_i(m_R^1(1,0), m_C^1(1,0)) - \pi_i(m_R^1(0,0), m_C^1(0,0))\right] \end{array} \tag{21}$$

and

$$\begin{array}{rcl} \beta_{i,R}(\mathbf{m}^d) - \beta_{i,R}(\mathbf{m}^1) & = & \left[\pi_i(m_R^d(1,1), m_C^d(1,1)) - \pi_i(m_R^d(0,1), m_C^d(0,1))\right] \\ & - & \left[\pi_i(m_R^1(1,1), m_C^1(1,1)) - \pi_i(m_R^1(0,1), m_C^1(0,1))\right] \\ & - & \left[\alpha_{i,R}(\mathbf{m}^d) - \alpha_{i,R}(\mathbf{m}^1)\right] \end{array} \tag{22}$$

Assumption 5* establishes conditions on the payoff matrix that imply $\alpha_{i,R}(\mathbf{m}^d) = \alpha_{i,R}(\mathbf{m}^1)$ and $\beta_{i,R}(\mathbf{m}^d) = \beta_{i,R}(\mathbf{m}^1)$ but they still allow for variation in the subject's behavior that provides identification of beliefs.

Let $\mathcal{D}_{d^1}$ be the set of treatments that satisfy the conditions in Assumption 5*. Assumption 5*(A) establishes that for any treatment d in $\mathcal{D}_{d^1}$, we have that $m_R^d(0,0) = m_R^d(1,0)$ and $m_C^d(0,0) = m_C^d(1,0)$. Looking at equation (21) we can see that this condition implies that $\alpha_{i,R}(\mathbf{m}^d) - \alpha_{i,R}(\mathbf{m}^1) = 0$. Assumption 5*(B) establishes that the monetary payoffs $m_R^d(0,1)$, $m_R^d(1,1)$, $m_C^d(0,1)$, and $m_C^d(1,1)$ do not vary over the treatments d in $\mathcal{D}_{d^1}$. Using equation (22), we can see that this restriction implies that $\beta_{i,R}(\mathbf{m}^d) - \beta_{i,R}(\mathbf{m}^1) = -[\alpha_{i,R}(\mathbf{m}^d) - \alpha_{i,R}(\mathbf{m}^1)]$, from Assumption 5*(A) this is equal to zero.

As shown in Example 5 above, Assumption 5* allows for variation in the matrices across treatments that can generate enough variation in players' behavior.

Proposition 3. *Consider the model that allows for preferences with altruism/envy. Suppose that Assumptions 1-4 and 5* hold. Then, the identification results Proposition 1(A) and 1(B) apply.* ■

4.2.2 Identification of beliefs in a public good game

So far, our identification of beliefs is based on exogenous variation in the payoff matrices. In games with more than two players, there are other possibilities to generate exogenous variation that can identify beliefs. We illustrate this in the context of public good binary choice game.

There are G players. The game is symmetric in the sense that every player has the same role. The binary choice consists in contributing to the public good ($a_i = 1$) or not ($a_i = 0$). Each individual has a utility function of money, $\pi_i(m_i)$, and the total utility is $\pi_i(m) + \varepsilon_i(a_i)$. The monetary payoff that player i receives depends on the other players' actions according to the following rule:

$$m_i\left(a_i, \sum_{j=1,j\neq i}^{G} a_j\right) = \begin{cases} m_s & \text{if} \quad a_i = 0 \\ m_s - c & \text{if} \quad a_i = 1 \text{ and } \sum_{j=1,j\neq i}^{G} a_j < \phi(G-1) \\ m_s + g & \text{if} \quad a_i = 1 \text{ and } \sum_{j=1,j\neq i}^{G} a_j \geq \phi(G-1) \end{cases} \tag{23}$$

If player i does not contribute to the funding of the public good ($a_i = 0$), she receives a safe amount of money m_s regardless the choices of the other players. If she contributes to the public good ($a_i = 0$), she has to pay a contribution $c > 0$, and her total monetary payoff depends on the fraction of other players that decide to contribute: $(G-1)^{-1} \sum_{j=1,j\neq i}^{G} a_j$. If this fraction is smaller than constant ϕ, then the public good project fails and player i receives a monetary payoff $m_s - c$. If this fraction is greater or equal than ϕ, then the project generates a net return $g > 0$ per contributor such that the monetary payoff of player i is $m_s + g$. We describe parameter ϕ as the *coordination difficulty*.

In this game, the exogenous characteristics that define players' monetary payoffs are $\mathbf{m} \equiv (m_s, c, g, \phi, G)$.

Player i needs to form beliefs about the probability of the event $\sum_{j=1,j\neq i}^{G} a_j \geq \phi(G-1)$. Let $B_i(\mathbf{m})$ be the probability that represents the belief of subject i. Our notation emphasizes that players' beliefs may depend on the parameters that characterize the monetary payoffs: (m_s, ϕ, G). For instance, one would expect $B_i(\mathbf{m})$ to decline with ϕ. The best response of individual i is to choose alternative $a_i = 1$ if

$$[1 - B_i(\mathbf{m})] \ \pi_i(m_s - c) + B_i(\mathbf{m}) \ \pi_i(m_s + g) + \varepsilon_i(1) \geq \pi_i(m_s) + \varepsilon_i(0) \tag{24}$$

And the best response choice probability is:

$$P_i(\mathbf{m}) = F\left[\alpha_i(m_s, c, g) + \beta_i(m_s, c, g) \ B_i(\mathbf{m})\right] \tag{25}$$

where $\alpha_i(m_s, c, g)$ and $\beta_i(m_s, c, g)$ have the following definitions: $\alpha_i(m_s, c, g) \equiv [\pi_i(m_s - c) - \pi_i(m_s) - \mu_i]/\sigma_i$; and $\beta_i(m_s, c, g) \equiv [\pi_i(m_s + g) - \pi_i(m_s - c)]/\sigma_i$.

Equation (25) shows that the model has exclusion restrictions: the number of players G and the *coordination difficulty* ϕ can affect players' beliefs but they are not arguments in the functions

$\alpha_i(m_s, c, g)$ and $\beta_i(m_s, c, g)$. These exclusion restrictions can be used to design an experiment that identifies players' beliefs. If the experiment includes treatments with the same value of (m_s, c, g) but different values of (G, ϕ), the variation in players' behavior between these treatments identifies players' beliefs. Proposition 4 formalizes this result.

Proposition 4. *Suppose that Assumptions 1 to 4 hold and that the experiment contains at least two treatments* $\mathbf{m}^1 = (m_s^1, c^1, g^1, \phi^1, G^1)$ *and* $\mathbf{m}^2 = (m_s^2, c^2, g^2, \phi^2, G^2)$ *such that: (i)* $(m_s^1, c^1, g^1) = (m_s^2, c^2, g^2)$*; (ii)* $(\phi^1, G^1) \neq (\phi^2, G^2)$*; and (iii)* $P_i(\mathbf{m}^1) \neq P_i(\mathbf{m}^2)$*. Let* $\mathcal{D}_{d^1}$ *the set of treatments that satisfy conditions (i) to (iii). Then: (A) The following equation characterizes the identified set of individual i's beliefs in treatments* $d \in \mathcal{D}_{d^1}$*:*

$$\frac{F^{-1}[P_i(\mathbf{m}^d)] - F^{-1}[P_i(\mathbf{m}^1)]}{F^{-1}[P_i(\mathbf{m}^2)] - F^{-1}[P_i(\mathbf{m}^1)]} = \frac{B_i(\mathbf{m}^d) - B_i(\mathbf{m}^1)}{B_i(\mathbf{m}^2) - B_i(\mathbf{m}^1)} \quad \text{for } d \in \mathcal{D}_{d^1}. \quad \blacksquare \tag{26}$$

(B) Let $B^*(.)$ *be a beliefs function, and consider the hypothesis represented by the restrictions* $B_i(\mathbf{m}^d) = B^*(\mathbf{m}^d)$ *for* $d \in \mathcal{D}_{d^1}$*. If the set* $\mathcal{D}_{d^1}$ *contains at least three treatments, this hypothesis can be tested by testing the restrictions implied by equation (26) with* $B_i(\mathbf{m}^d) = B_i^*(\mathbf{m}^d)$ *for* $d \in \mathcal{D}_{d^1}$*.* $\blacksquare$

4.3 Complete identification of utility and beliefs

We now consider the full identification of the model. More precisely, we consider the identification of utility parameters $\alpha_{i,R}(\mathbf{m}_R)$ and $\beta_{i,R}(\mathbf{m}_R)$ and belief parameters $B_{i,R}(\mathbf{m}_R, \mathbf{m}_C^d)$ for every treatment d in the set $\mathcal{D}_{d^1}$. The focus on the set $\mathcal{D}_{d^1}$ is without loss of generality: we can consider treatments d^1 with different values of $\mathbf{m}_R$ such that functions $\alpha_{i,R}(.)$, $\beta_{i,R}(.)$, and $B_{i,R}(.)$ are identified everywhere.

Proposition 1 shows that, for an individual i, we can identify beliefs $B_{i,R}(\mathbf{m}^d)$ for every treatment d in $\mathcal{D}_{d^1}$ up to location and up to scale. That is, given beliefs at points $B_{i,R}(\mathbf{m}^1)$ and $B_{i,R}(\mathbf{m}^2)$, equation (14) provides the value of $B_{i,R}(\mathbf{m}^d)$ at every treatment d in $\mathcal{D}_{d^1}$. Therefore, if we impose two restrictions in the beliefs function $B_{i,R}(.)$, the model can be fully identified.

Suppose that the researcher is willing to restrict the value of subject i's beliefs in treatments d^1 and d^2 such that $B_{i,R}(\mathbf{m}^1)$ and $B_{i,R}(\mathbf{m}^2)$ are known to the researcher. For instance, the researcher could be willing to impose the restriction of rational beliefs in these treatments – $B_{i,R}(\mathbf{m}^d) = (N-1)^{-1} \sum_{j \neq i} P_{j,C}(\mathbf{m}^d)$ for $d = d^1, d^2$ – or use information on elicited beliefs. Under these restrictions, we can use (12) evaluated at $\mathbf{m}^1$ and $\mathbf{m}^2$ to identify $\alpha_{i,R}(\mathbf{m}_R)$ and $\beta_{i,R}(\mathbf{m}_R)$. That is,

$$\beta_{i,R}(\mathbf{m}_R) = \frac{F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]}{B_{i,R}(\mathbf{m}^2) - B_{i,R}(\mathbf{m}^1)} \tag{27}$$

And,

$$\alpha_{i,R}(\mathbf{m}_R) = \frac{B_{i,R}(\mathbf{m}^2)\ F^{-1}[P_{i,R}(\mathbf{m}^1)] - B_{i,R}(\mathbf{m}^1)\ F^{-1}[P_{i,R}(\mathbf{m}^2)]}{B_{i,R}(\mathbf{m}^2) - B_{i,R}(\mathbf{m}^1)} \tag{28}$$

Then, given $\alpha_{i,R}(\mathbf{m}_R)$ and $\beta_{i,R}(\mathbf{m}_R)$, we can use equation (12) to obtain beliefs at any treatment $d \in \mathcal{D}_{d^1}$:

$$B_{i,R}(\mathbf{m}^d) = \frac{F^{-1}[P_{i,R}(\mathbf{m}^d)] - \alpha_{i,R}(\mathbf{m}_R)}{\beta_{i,R}(\mathbf{m}_R)} \tag{29}$$

Proposition 5. *Suppose that Assumptions 1 to 5 hold and that $B_{i,R}(\mathbf{m}^1)$ and $B_{i,R}(\mathbf{m}^2)$ are known to the researcher. Then, $\alpha_{i,R}(\mathbf{m}_R)$, $\beta_{i,R}(\mathbf{m}_R)$, and $\{B_{i,R}(\mathbf{m}^d) : d \in \mathcal{D}_{d^1}\}$ are point identified.* ■

Proposition 5 requires that subject i's beliefs at d^1 and d^2 are observed to the researcher. This condition holds true if the researcher has access to data of credible elicited belief at these two treatments. It also holds if the researcher is willing to assume subject i has unbiased belief of her opponents' behaviors at d^1 and d^2 so that subject i's beliefs can be identified from choice probabilities of the other subjects' as column players.

If the researcher is willing to assume equilibrium beliefs – or some other beliefs – in more than two treatments, then the model is over-identified and a test of over-identifying restrictions can be used to evaluate internal consistency of the model. For instance, in our empirical application to a coordination game, this test is implicit in Table 4 where we show that estimated beliefs all appear to be within confidence bands regardless of the over-identifying restrictions.

The selection of the two treatments where to impose the unbiased beliefs assumption is an important decision. The selection, if misspecified, can generate bias estimates on both preferences and beliefs. However, this condition is substantially weaker than those imposed by the models presented in Examples 1 to 4 which are commonly used in empirical applications. The identification result in Proposition 1 can provide a guidance on the selection of treatments.

When the game is such that the matrix of monetary payoffs for the row player has a Toeplitz structure – as in Proposition 2 – subject i's preferences and beliefs are identified under only one restriction on beliefs.

Proposition 6. *Suppose that Assumptions 1 to 5 hold and: (i) the matrix of monetary payoffs of the row player is symmetric and diagonal-constant (Toeplitz matrix); (ii) $\mu_{i,R} = 0$; and (iii) $B_{i,R}(\mathbf{m}^1)$ is known to the researcher. Then, $\alpha_{i,R}(\mathbf{m}_R)$, $\beta_{i,R}(\mathbf{m}_R)$, and $\{B_{i,R}(\mathbf{m}^d) : d \in \mathcal{D}_{d^1}\}$ are point identified.* ■

5 Estimation and inference

In this section, we describe a constrained maximum likelihood approach to estimate the parameters of the model and to test different types of restrictions on beliefs and/or preference.

Consider an experiment as described in Section 3 with K treatments where each treatment d is an element of the set $\mathcal{D} \equiv \mathcal{M} \times \{R, C\}$. The unconstrained log-likelihood for this model and data has the following form:

$$\ell(\mathbf{P}) = \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{d \in \mathcal{D}} 1\{d_{it} = d\} \left(a_{it} \ \log\left[P_{i,r_d}(\mathbf{m}^d)\right] + (1 - a_{it}) \ \log\left[1 - P_{i,r_d}(\mathbf{m}^d)\right]\right) \tag{30}$$

where $\mathbf{P}$ is the vector with all the conditional choice probabilities (CCPs): $\mathbf{P} = \{P_{i,r_d}(\mathbf{m}^d) : d \in \mathcal{D}; i \in \mathcal{I}\}$. This is the log-likelihood function of a nonparametric multinomial model. As it is well-known, the unconstrained maximum likelihood estimator (MLE) of $\mathbf{P}$ is the frequency estimator: $\widehat{P}_{i,r_d}(\mathbf{m}^d) = [\sum_{t=1}^{T} 1\{d_{it} = d\} a_{it}] / [\sum_{t=1}^{T} 1\{d_{it} = d\}]$. Let $\widehat{\mathbf{P}}_U$ be this unconstrained MLE.

Proposition 1 establishes that – under assumptions 1 to 5 and the condition of unbiased beliefs for subject i as row player – the vector of CCPs $\mathbf{P}$ should satisfy the following restrictions. For every treatment $d \in \mathcal{D}(d^1, d^2)$:

$$\frac{F^{-1}[P_{i,R}(\mathbf{m}^d)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]}{F^{-1}[P_{i,R}(\mathbf{m}^2)] - F^{-1}[P_{i,R}(\mathbf{m}^1)]} - \frac{\sum_{j \neq i} P_{j,C}(\mathbf{m}^d) - P_{j,C}(\mathbf{m}^1)}{\sum_{j \neq i} P_{j,C}(\mathbf{m}^2) - P_{j,C}(\mathbf{m}^1)} = 0, \tag{31}$$

where $\mathcal{D}(d^1, d^2)$ is the subset of treatments where the row player has the same payoff matrix as in treatments d^1 and d^2. We can represent these restrictions using the system of equation $\mathbf{f}(\mathbf{P}) = \mathbf{0}$. With some abuse of notation, this system can represent the restrictions of unbiased beliefs for only one subject, or for a subset of subjects, or for all the subjects in the experiment. We can define a constrained MLE of CCPs that imposes the restrictions $\mathbf{f}(\mathbf{P}) = \mathbf{0}$. This constrained MLE is:

$$\widehat{\mathbf{P}}_{const} \quad = \quad \arg\max_{\mathbf{P}} \ \ell(\mathbf{P}) \quad \text{subject to:} \ \ \mathbf{f}(\mathbf{P}) = 0 \tag{32}$$

Using this nonparametric maximum likelihood setting, we can test the null hypothesis $\mathbf{f}(\mathbf{P}) = 0$ using a likelihood ratio test, a Wald test, or a Lagrange multiplier test. For instance, the likelihood ratio test is based on the statistic:

$$2\left[\ell\left(\widehat{\mathbf{P}}_U\right) - \ell\left(\widehat{\mathbf{P}}_{const}\right)\right] \ \text{that under null} \rightarrow_d \chi^2_q \tag{33}$$

where q is the number of restrictions in the vector $\mathbf{f}(\mathbf{P})$. The Wald test is based on the statistic:

$$\mathbf{f}\left(\widehat{\mathbf{P}}_U\right)' \left[\frac{\partial \mathbf{f}(\widehat{\mathbf{P}}_U)}{\partial \mathbf{P}'} Var\left(\widehat{\mathbf{P}}_U\right) \frac{\partial \mathbf{f}(\widehat{\mathbf{P}}_U)'}{\partial \mathbf{P}}\right]^{-1} \mathbf{f}\left(\widehat{\mathbf{P}}_U\right) \ \text{that under null} \rightarrow_d \chi^2_q \tag{34}$$

The MLEs and tests described above consider only estimation of CCPs and tests of unbiased beliefs . The researcher may be interested in the estimation of the full model. Let $\mathbf{B}$, α, and β be the vectors with beliefs and preference parameters $\alpha_{ir}(\mathbf{m}_r)$ and $\beta_{ir}(\mathbf{m}_r)$, respectively. We can distinguish two sets of restrictions: the best response equations, $F^{-1}[P_{i,r}(\mathbf{m})] - \alpha_{i,r}(\mathbf{m}_r) - \beta_{i,r}(\mathbf{m}_r)$ $B_{i,r}(\mathbf{m}) = 0$; and additional restrictions that we need to just identify or to over-identify the full model. We represent all the restrictions as $\mathbf{c}(\mathbf{P}, \mathbf{B}, \alpha, \beta) = 0$. The constrained MLE is:

$$\left(\widehat{\mathbf{P}}_{const},\ \widehat{\mathbf{B}},\ \widehat{\alpha},\ \widehat{\beta}\right) \quad = \quad \arg \max_{\mathbf{P}, \mathbf{B}, \alpha, \beta} \ \ell\left(\mathbf{P}\right) \qquad \text{subject to:} \ \ \mathbf{c}(\mathbf{P}, \mathbf{B}, \alpha, \beta) = \mathbf{0} \tag{35}$$

We can test the overidentifying restrictions involved in this constrained MLE using likelihood ratio test, Wald test, or Lagrange multiplier test.

6 Empirical applications

In this section, we illustrate our identification results and tests using data from two laboratory experiments. These experiments incorporate the exclusion restriction in Assumption 5. Section 6.1 presents the matching pennies game studied by Goeree and Holt (2001). Section 6.2 deals with the public good game from Heinemann, Nagel and Ockenfels (2008).

6.1 Matching pennies

6.1.1 Experiment

Two players simultaneously choose between two possible actions: 0 or 1. Table 1 presents the monetary payoff matrices. The pair of numbers between brackets, $[m_R,\, m_C]$, represents the monetary payoffs of the row and the column player, respectively, measured in cents. The experiment contains three games or monetary payoff matrices. The only difference across games is in the monetary payoff of the column player under action profile $(a_R, a_C) = (0, 0)$. It is clear that this experimental design satisfies the exclusion restriction in Assumption 5.

Table 1: Matching Pennies Experiment (Goeree and Holt, 2001)

Monetary Payoff Matrix $\mathbf{m}^1$

		Player C: $a_C = 0$	Player C: $a_C = 1$
Player R	$a_R = 0$	[40 , 320]	[80 , 40]
	$a_R = 1$	[80 , 40]	[40 , 80]

Monetary Payoff Matrix $\mathbf{m}^2$

		Player C: $a_C = 0$	Player C: $a_C = 1$
Player R	$a_R = 0$	[40 , 80]	[80 , 40]
	$a_R = 1$	[80 , 40]	[40 , 80]

Monetary Payoff Matrix $\mathbf{m}^3$

		Player C: $a_C = 0$	Player C: $a_C = 1$
Player R	$a_R = 0$	[40 , 44]	[80 , 40]
	$a_R = 1$	[80 , 40]	[40 , 80]

The experiment includes 50 subjects ($N = 50$) who – at the moment of the experiment – were undergraduates in an economic class at the University of Virginia. Subjects were randomly matched and assigned as row or column player. Every subject played the three games once, and his/her role as either row or column player was fixed. The order of the three games – the three payoff matrices — was also randomly alternated.[11]

In this experiment, we have that $T = 3$ such that each subject is observed only once in each treatment. Therefore, it is impossible to estimate utility and beliefs at the individual level. In this application, we impose the restriction that subjects have the same utility function of money. Under this restriction, the data can be seen as coming from a single subject making 25 independent decisions for each treatment $(\mathbf{m}, r)$.

6.1.2 Empirical results

(i) Unconstrained estimation of CCPs.

[11] For more details on this experiment, visit http://www.people.virginia.edu/~cah2k/trdatatr.pdf. In addition to this matching pennies game, subjects played other nine types of games which are not the focus of our paper. Every subject was paid $6 for showing up. The average earnings of a subject over all the 10 games — during a two-hour session — was about $35 ranging between $15 and $60.

Table 2 presents frequency estimates of players' choice probabilities from this experiment and the corresponding standard errors. These estimates correspond to the unconstrained MLE defined in equation (30). The behaviors of both players vary across the three treatments. Though the monetary payoff matrix of the row player is the same in the three treatments, the behavior of this player varies substantially. According to the model, the change in the row player's behavior should be attributed to the change in this player's beliefs about the column player's choices. We exploit this source of variation in the experiment to test unbiased beliefs of the row player and to identify beliefs and utilities for this player. Since the experiment does not provide the same source of variation for the column player, we cannot identify beliefs and preferences for this player.

Table 2: Matching Pennies Game Experiment
Empirical Choice Probabilities: $N = 50$
(Standard errors in parentheses)

	Player R – $\widehat{P}_R(\mathbf{m})$		*Player C* – $\widehat{P}_C(\mathbf{m})$	
Payoff matrix $\mathbf{m}^1$	0.84	(0.073)	0.04	(0.039)
Payoff matrix $\mathbf{m}^2$	0.52	(0.100)	0.52	(0.100)
Payoff matrix $\mathbf{m}^3$	0.20	(0.080)	0.92	(0.054)

Note: For player r, $\widehat{P}_r(\mathbf{m}) = [\sum_{i=1}^{N} \sum_{t=1}^{T} a_{it,r} 1\{d_{it} = (\mathbf{m}, r)\}]/[\sum_{i=1}^{N} \sum_{t=1}^{T} 1\{d_{it} = (\mathbf{m}, r)\}]$

The monetary payment for player C in outcome $(0, 0)$ declines monotonically when we go from payoff matrix $\mathbf{m}^1$ to $\mathbf{m}^2$, and from $\mathbf{m}^2$ to $\mathbf{m}^3$. Therefore, alternative $a_C = 0$ becomes less attractive to the column player when we go down along the sequence of games $\mathbf{m}^1 \to \mathbf{m}^2 \to \mathbf{m}^3$. This shows in the estimated choice probability of the column player. The probability of choosing alternative 1 increases: $\widehat{P}_C(\mathbf{m}^1)[= 0.04] < \widehat{P}_C(\mathbf{m}^2)[= 0.52] < \widehat{P}_C(\mathbf{m}^3)[= 0.92]$, and these inequalities are statistically significant.

If player R had rational beliefs, she would predict that player C chooses $a_C = 0$ with higher probability when we move through the sequence of games $\mathbf{m}^1 \to \mathbf{m}^2 \to \mathbf{m}^3$. The best response to such a belief is to choose $a_R = 1$ less frequently. The estimated choice probabilities in Table 2 are consistent with this argument: $\widehat{P}_R(\mathbf{m}^1)[= 0.84] > \widehat{P}_R(\mathbf{m}^2)[= 0.52] > \widehat{P}_R(\mathbf{m}^3)[= 0.20]$, and these inequalities are statistically significant.

However, this argument is qualitative and fails to consider the effect of preferences. Specifically, the above inequalities only imply that, when moving from $\mathbf{m}^1$ to $\mathbf{m}^3$, player R correctly predicts that player C will choose action 1 less frequently. However, these inequalities are not informative about the magnitude of player R's beliefs. This player may still have a biased belief on the magnitude of the change of player C's choices even though she correctly predicts the direction of the change.

Furthermore, without knowing how much does player R prefer 80 cents than 40 cents, it is impossible to test whether she has an unbiased belief on the magnitude. The method studied in this paper takes preferences into account and is suitable to test unbiased belief.

(ii) Wald test of unbiased beliefs of row players.

Table 3 presents Wald tests of the null hypothesis of unbiased beliefs as described in equation (34) above. In this experiment, the matrices of monetary payoffs are Toeplitz. Therefore, the restrictions of unbiased beliefs have the form in Proposition 2 and equation (18). The vector of restrictions $\mathbf{f}(\mathbf{P}) = \mathbf{0}$ includes two restrictions: $f_{1d}(\mathbf{P}) = 0$ for $d = 2, 3$, where:

$$f_{1d}(\mathbf{P}) \equiv F^{-1}[P_R(\mathbf{m}^d)] \ \left(1 - 2\ P_C(\mathbf{m}^1)\right) - F^{-1}[P_R(\mathbf{m}^1)] \ \left(1 - 2\ P_C(\mathbf{m}^d)\right) \tag{36}$$

Table 3: Wald Tests of Unbiased Beliefs
Matching Pennies

	Probit	*Logit*	*Exponential*	*Double Exp.*
$H_0 : f_{12}(\mathbf{P}) = 0$				
$f_{12}(\widehat{\mathbf{P}})$ [s.e]	-0.0859 [0.3242]	-0.1400 [0.5402]	0.2114 [0.3041]	0.0935 [0.4033]
(p-value)	(0.7932)	(0.7926)	(0.4338)	(0.8146)
$H_0 : f_{13}(\mathbf{P}) = 0$				
$f_{13}(\widehat{\mathbf{P}})$ [s.e]	0.0758 [0.2886]	0.1277 [0.4760]	-0.1924 [0.2442]	-0.0859 [0.3416]
(p-value)	(0.7876)	(0.7868)	(0.4142)	(0.8014)
$H_0 : f_{12}(\mathbf{P}) = f_{13}(\mathbf{P}) = 0$				
Chi-square	0.1392	0.1336	0.6621	0.0633
(p-value)	(0.9328)	(0.9354)	(0.7327)	(0.9688)

Table 3 presents Wald tests for three different null hypotheses: (1) $f_{12}(\mathbf{P}) = 0$; (2) $f_{13}(\mathbf{P}) = 0$; and (3) the joint restriction, $f_{12}(\mathbf{P}) = 0$ and $f_{13}(\mathbf{P}) = 0$. We report our tests under four specifications for the distribution of the unobservables: Probit, Logit, Exponential and Double Exponential. All the tests are consistent with the hypothesis that the row player has unbiased beliefs in the three treatments. All the p-values are greater than 0.4 and highly insignificant.

(iii) Full estimation of preferences and beliefs.

Table 4 presents estimates of utility and belief parameters using the constrained maximum likelihood estimator defined in equation (35). In addition to the restrictions from best response equations, we further impose the restriction that the row player has unbiased beliefs in one of the three games. Given the Toeplitz matrix of monetary payoff we have that $\beta_R(\mathbf{m}_R) = -2\alpha_R(\mathbf{m}_R)$.

And given that the payoff matrix of the row player is constant over the three treatments, the likelihood function depends only on one utility parameter: $\alpha_R(\mathbf{m}_R) = \pi(80) - \pi(40)$. We present only estimates for the Logit model, but the results are qualitatively the same under the other distributions.

Table 4: Full ML Estimation of Preferences and Beliefs[(a)]
Matching Pennies

Parameter	*Under unbiased beliefs in game 1*	*Under unbiased beliefs in game 3*	*Under unbiased beliefs in games 1, 2, 3*
$\pi(80) - \pi(40)$	1.8024*** (0.6129)	1.6504*** (0.6325)	1.7309*** (0.4413)
$B_R(\mathbf{m}^1)$	0.04 (0.039) $[con]^{(b)}$	-0.0024 (0.2538)$^{(c)}$	0.04 (0.039) $[con]^{(b)}$
$B_R(\mathbf{m}^2)$	0.4778*** (0.1113)	0.4757*** (0.1216)	0.52 (0.100) $[con]^{(b)}$
$B_R(\mathbf{m}^3)$	0.8846*** (0.1907)	0.92 (0.054) $[con]^{(b)}$	0.92 (0.054) $[con]^{(b)}$
log-likelihood	-69.2870	-69.2870	-69.3417
LR test - unbiased beliefs (p-value)			0.1093 (0.9468)

Note (a): Standard errors are in parenthesis; ***, **, * represent significance at 10%, 5% and 1% respectively.
Note (b): The estimated belief parameter under the unbiased restriction is the MLE of $P_C(m^1)$.
Note (c): The estimate of $B_R(m^1)$is negative because we do not impose the restriction $0 \leq B_R(m^1) \leq 1$. Restricting $B_R(m^1)$ to be zero implies an almost identical value for the log-likelihood.

Table 4 presents our estimation results. The first two columns show the estimates when we impose player R has unbiased belief in games 1 and 3, respectively.[12] As shown by Proposition 6, each of these two restrictions just identifies the model. Therefore, by construction, they yield the same value for the log-likelihood function. The last column shows the estimates of preferences when we impose the restriction of unbiased beliefs in the three games. In all these specifications, the estimates of the utility parameter $\alpha_R = \pi(80) - \pi(40)$ are significantly positive and of very similar magnitude. The last row shows the likelihood ratio test of unbiased belief hypothesis in all treatments. The test statistic has a high p-value of 0.9468. This is also reflected by the closeness between the unconstrained and constrained estimates of beliefs.

[12] We omit the specification that assumes player R has unbiased belief in game 2. This is because player C's choice probability is very close to 50%, with Toeplitz monetary payoff matrix, player R's expected utility of two actions will be very close, regardless of the utility function she has. Therefore, the utility function is imprecisely estimated.

Goeree and Holt (2001) conclude that player R tends to correctly predict player C's behaviors, based on the observations of choice probabilities. In this paper, we confirm their qualitative observation in a game with incomplete information and a nonparametric specification of the utility function.

6.2 Coordination game

Heinemann, Nagel, and Ockenfels (2009, henceforth HNO) study and measure the strategic uncertainty that appears in games with multiple equilibria when players have non-coordinated beliefs about the selected equilibrium. To study this phenomenon, they design and implement a randomized experiment using a set of coordination games with different group sizes, monetary payoffs and coordination difficulty. Unlike the matching pennies in previous subsection, HNO experiment contains: (a) many experiments for each subject – large T – such that we can allow preferences and beliefs to vary freely over subjects; and (b) a rich variation of monetary payoffs such that we can estimate a non-linear utility function of money.

6.2.1 Experiment

Heinemann, Nagel and Ockenfels (2009) study a coordination game of G players with binary choice. The game is a particular case of the public good game that we have presented in section 4.2.2 under the restrictions $c = m_s$ – the contribution to the public good is equal to the safe amount of money – and $g = 15 - m_s$. Consequently, the identification results in Propositions 4 and 5 apply.

In this coordination game, the matrix of players' monetary payoffs is determined by $\mathbf{m} \equiv (m_s, \phi, G)$. In addition, $B_i(\mathbf{m})$ represents individual i's belief on other players' aggregate behavior, instead of beliefs about the choice of a single player. It is the belief probability that at least a fraction ϕ of the other $(G-1)$ players will choose action 1. We use $P^{co}_{-i}(\mathbf{m})$ to denote the actual probability of this event, i.e. $P^{co}_{-i}(\mathbf{m}) = Prob[\sum_{j=1,j\neq i}^{G} a_j \geq \phi(G-1)]$. Therefore, subject i has unbiased beliefs in game $\mathbf{m}$ if $B_i(\mathbf{m}) = P^{co}_{-i}(\mathbf{m})$. As shown in section 4.2.2, the coordination difficulty ϕ and the group size G can affect a subject's beliefs but not her own monetary payoff, such that they can be used as exclusion restrictions to identify beliefs.

The experiment was conducted in different locations: Frankfurt, Barcelona, Bonn and Cologne. Heinemann, Nagel and Ockenfels (2009) report that there are differences among subject pools. Our empirical analysis mainly focuses on the observations from Frankfurt since this pool contains most of subjects. We also report results from Bonn where data on elicited beliefs is available such that we can test the validity of elicitation process.[13]

The experiment includes 90 treatments or games according to all the possible values of the parameters $\mathbf{m} = (m_s, \phi, G)$ with $G \in \{4, 7, 10\}$, $\phi \in \{1/3,\ 2/3,\ 1\}$, and 10 values of m_s between 1.5 Euros and 15 Euros with an incremental unit of 1.5 Euros. As this coordination game has symmetric monetary payoffs, the role of player is indistinguishable and it is not a treatment variable in this section. Subjects were randomly assigned into a group G. Then, given the selection of group size G, a subject participates in all the treatments for every value of ϕ and m_s. Therefore, each subject participates in 30 treatments. There are 64, 42, 40 subjects assigned to group size $G = 4, 7, 10$, respectively.

In addition to this coordination game, each subject also makes 10 binary lottery choices. Action 0 is a safe choice that pays m_s Euros (same set of 10 values of m_s as coordination game) and action 1 is risky that rewards 15 Euros with probability 2/3. From the point of view of expected utility maximization, in absence of altruism and/or envy, this lottery choice is equivalent to a coordination problem such that the coordination probability is exogenously and objectively fixed at 2/3. We include these 10 lottery choices in the estimation as it is natural to impose unbiased belief assumption and facilitate the estimation of individual's preferences.

To prevent learning, Heinemann, Nagel and Ockenfels (2009) give no feedback across treatments. At the end of a session, only one treatment is randomly selected to determine subject's earning. This avoids potential hedging and each decision situation can be treated as independent. The duration of a session is about 40-60 minutes with an average earning of 16.68 Euros per subject.

6.2.2 Empirical results

(i) Pooled estimates of CCPs and of the true probability if successful coordination

Figure 1 plots the empirical choice probabilities pooling all the subjects in the same treatment. Each of the three panels corresponds to a value for group size G. The horizontal axis represents the safe action's monetary reward m_s. We report the estimated CCPs for $\phi = 1/3$ (i.e. blue line) and $\phi = 1$ (i.e. red line). It is clear that subjects are more reluctant to choose the risky action when they have a better safe choice. Coordination difficulty ϕ significantly affects subjects' choice probabilities, especially when the safe action's monetary payoff is in the medium range.

[13] The experiment in Frankfurt was done at a computer laboratory in the Economics Department of the University of Frankfurt between May and July 2003. Most of subjects were undergraduates in business and economics. For details about this experiment, see section 3 in Heinemann, Nagel and Ockenfels (2009). The experimental instructions are available on the supplements page of the *Review of Economic Studies* website at http://www.restud.org.

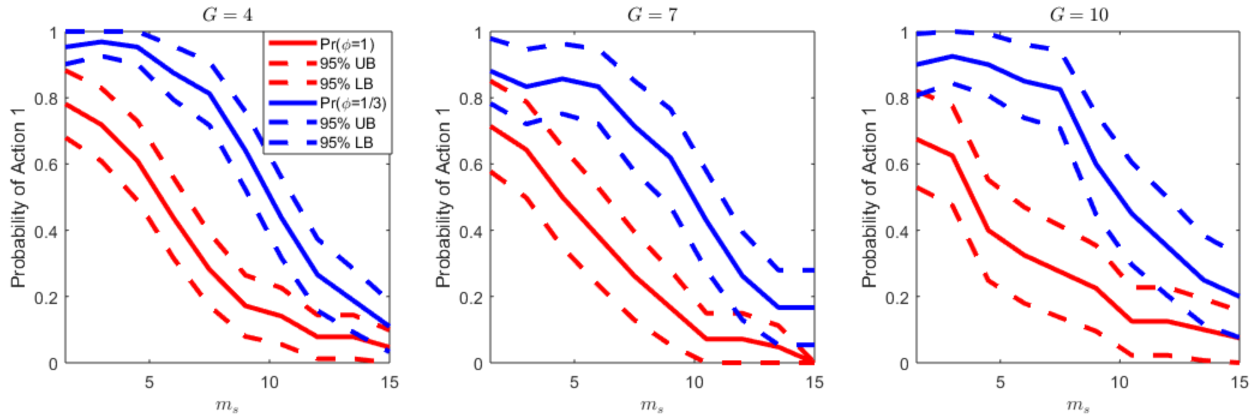

Figure 1: Empirical Choice Probabilities

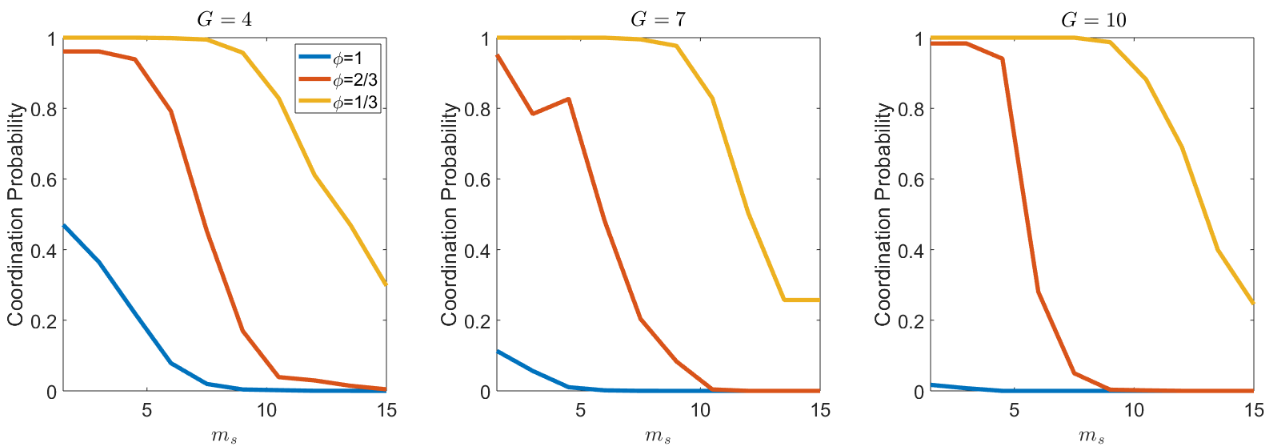

Figure 2: Actual Coordination Probabilities

Figure 2 plots the empirical probability of successful coordination pooled over all the subjects in the same treatment: $P^{co}_{-i}(\mathbf{m}) = \Pr[\sum_{j=1, j \neq i}^{G} a_j \geq \phi(G-1)]$. We calculate this probability by drawing all possible combinations of $(G-1)$ individuals from the N subjects in the experiment. Since N is relative large, each subject i faces the same actual coordination probability: $P^{co}_{-i}(\mathbf{m}) = P^{co}(\mathbf{m})$. Moreover, since the number of combinations is very large, the sampling error in the estimation of $P^{co}(\mathbf{m})$ is negligible. We see that $P^{co}(\mathbf{m})$ declines with m_s and with ϕ. In particular, $P^{co}(\mathbf{m})$ is practically equal to one for treatments with $\phi = 1/3$ and $m_s \leq 10$, and it is equal to zero when $\phi = 1$ and $m_s > 10$.

(ii) Specification assumptions

In this experiment, we aim to estimate preferences and test unbiased belief assumption at the individual level. Since a subject only participates in a single group size G, we rely on the variation of ϕ to identify the model.

We apply the constrained MLE as described in equation (35). To fully identify the preferences and beliefs of an individual, we impose the restriction that she has unbiased beliefs when $\phi = 1/3$. Together with the data on lottery choice, this restriction just identifies the preference and subject i's belief when $\phi = 2/3$ and when or $\phi = 1$.

We apply the likelihood ratio test of unbiased beliefs separately for each subject. In the implementation of this test, the constrained MLE imposes the restriction that individual i has unbiased beliefs at every treatment.

We select $\phi = 1/3$ as treatment with unbiased beliefs because coordination is relatively easy to maintain in this game. As shown in figure 2, successful coordination is achieved with probability very close to 1 for most values of m_s. This observation suggests that strategic uncertainty is small when $\phi = 1/3$, and it is relatively easy for subjects to form unbiased beliefs. In contrast, when $\phi = 2/3$, the probability of successful coordination varies substantially as m_s increases. Such a drastic change in $P^{co}(\mathbf{m})$ may imply substantial strategic uncertainty and prevent correct beliefs.

There are $T = 40$ observations per subject: 10 lottery choices and 30 coordination games. Using a fully nonparametric specification of preferences and beliefs, the model has 30 unknown parameters per subject: 10 utility parameters $\pi_i(m_s) - \pi_i(0)$ for 10 values of m_s, and 20 belief parameters. The estimates of these 30 parameters with 40 observations is quite imprecise. Therefore, we impose some parametric restrictions.

We assume a CRRA utility function: $\pi_i(m) = m^{\rho_i}$, where ρ_i is a subject-specific parameter and $1 - \rho_i$ represents the subjects' relative risk aversion. That is, $0 < \rho_i < 1$ implies risk aversion, and $\rho_i > 1$ implies risk loving preferences.

For the beliefs function, we consider the following specification:

$$B_i(\mathbf{m}) = \frac{\exp\left\{\ln\left(P^{co}(\mathbf{m})\right) - \ln\left(1 - P^{co}(\mathbf{m})\right) + \gamma_i(\mathbf{m})\right\}}{1 + \exp\left\{\ln\left(P^{co}(\mathbf{m})\right) - \ln\left(1 - P^{co}(\mathbf{m})\right) + \gamma_i(\mathbf{m})\right\}} \tag{37}$$

Function $\gamma_i(\mathbf{m})$ captures the bias in the beliefs of subject i in treatment $\mathbf{m}$ such that $\gamma_i(\mathbf{m}) = [\ln\left(B_i(\mathbf{m})\right) - \ln\left(1 - B_i(\mathbf{m})\right)] - [\ln\left(P^{co}(\mathbf{m})\right) - \ln\left(1 - P^{co}(\mathbf{m})\right)]$. If function $\gamma_i(\mathbf{m})$ were unrestricted, our specification of beliefs would be completely nonparametric. We impose the restriction that $\gamma_i(\mathbf{m}) = \gamma_i(\phi)$ for any $\mathbf{m} = (m_s, G, \phi)$. That is, we restrict γ_i to vary only with ϕ. This restriction is motivated by the need of a parsimonious specification due to the relatively small T. Importantly, we allow all parameters to vary across subjects.

Our identification results require that the exclusion restriction exogenously affects players' choice probabilities; therefore, we exclude the subjects who always make the same decision when ϕ or m_s vary. We further eliminate 5 subjects whose behaviors are anomalies.[14] This reduces our sample from 146 subjects to 95 individuals.

(iii) LR test of unbiased beliefs

We have implemented the LR test of unbiased beliefs separately for each of the 95 subjects, and obtained the corresponding p-values. Figure 3 plots the inverse of the empirical CDF of these p-values. The horizontal axis indexes the subject rank, rather than the cumulative probability. The blue line represents our benchmark model with utility function $\pi_i(m) = m^{\rho_i}$ and the black line plots the result under the restriction of linear utility, $\pi_i(m) = m$. At the individual level, the hypothesis of unbiased beliefs is rejected at 5% significance level for a majority of the subjects: 75 out of 96. Cautiously, since we perform unbiased belief test separately for each individual, this result should not be interpreted as claiming that 75 subjects have biased beliefs, because that null hypothesis involves multiple testing. To shed light on beliefs at aggregate level, we further test the null hypothesis that all individuals have unbiased beliefs using Bonferroni's correction. More specifically, to obtain a significance level α for the joint test for N individuals, we apply a significance level α/N for each of the individual tests. With this Bonferroni correction, the null hypothesis that all individuals have unbiased beliefs is rejected at 1% significance level.

A comparison between the blue and black lines shows an interesting result. Intuitively, if a subject is not risk neutral, imposing the restriction of a linear utility function biases the estimate of the belief functions. Therefore, the estimated beliefs may spuriously generate over-rejection of the null hypothesis of unbiased beliefs. Our results confirm this intuition, as reflected by a black curve that is below the blue curve. On average, the p-value under the linear utility function is 0.035 lower than the one under the CRRA utility.

[14] For instance, they may randomly choose each action or choose coordination with higher probability when safe action pays more monetary reward. These anomalies yield a utility function that is decreasing in monetary payoff.

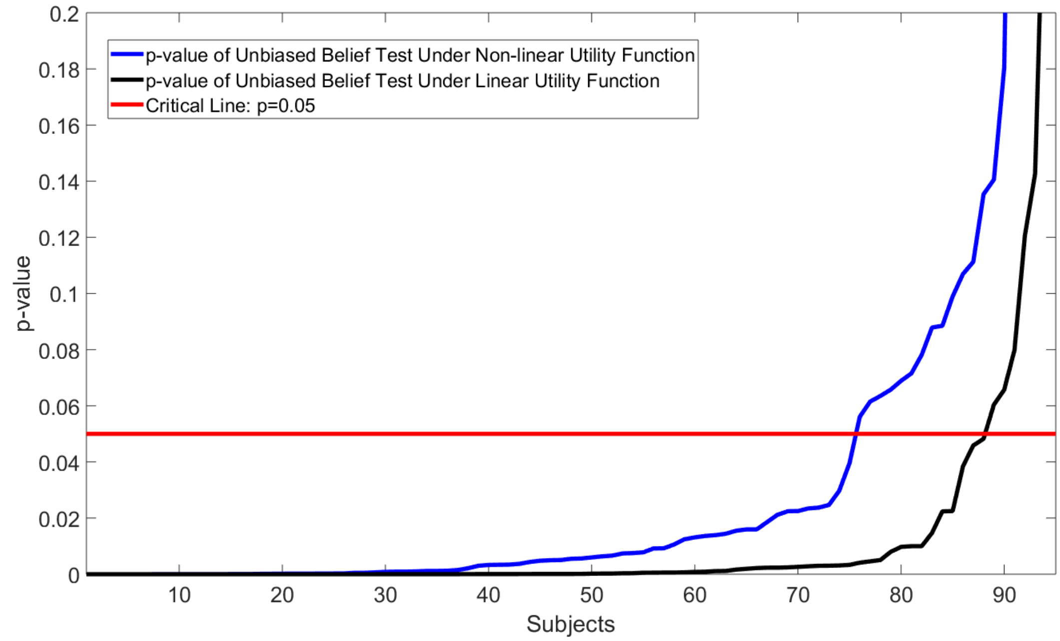

Figure 3: Inverse CDF of p-value of Unbiased Belief Test

(iv) Estimation of beliefs $\gamma_i(\phi)$

In the experiment, most of individuals correctly predict that the actual coordination probability decreases with the coordination difficulty ϕ: they choose the risky action less frequently as ϕ increases. However, our test suggests that even though most subjects respond in the right direction, they fail to respond in the right magnitude. As shown in Figure 2, the actual probability of successful coordination decreases very sharply when ϕ increases. In contrast, figure 1 shows that players' choice probability of risk action only decreases moderately as ϕ increases. This moderate change in players' actions imply that subjects underestimate the drop of coordination probability when coordination is more difficult. More specifically, 64 out of 95 subjects have positive estimates of both $\gamma_i(\phi = 2/3)$ and $\gamma_i(\phi = 1)$. These individuals over-predict the probability of coordination when $\phi = 2/3$ or 1, and therefore, they under-estimate the decrease in the actual probability of coordination.

(v) Estimation of risk preferences

Figure 4 plots the histogram of the estimated risk parameters ρ_i. This histogram shows that for approximately two-thirds of the individuals the estimate of $\hat{\sigma_i}$ implies risk aversion (i.e., $\hat{\sigma_i} < 1$). Interestingly, for around one-third of the subjects, the estimates of $\hat{\sigma_i}$ implies risk loving preferences

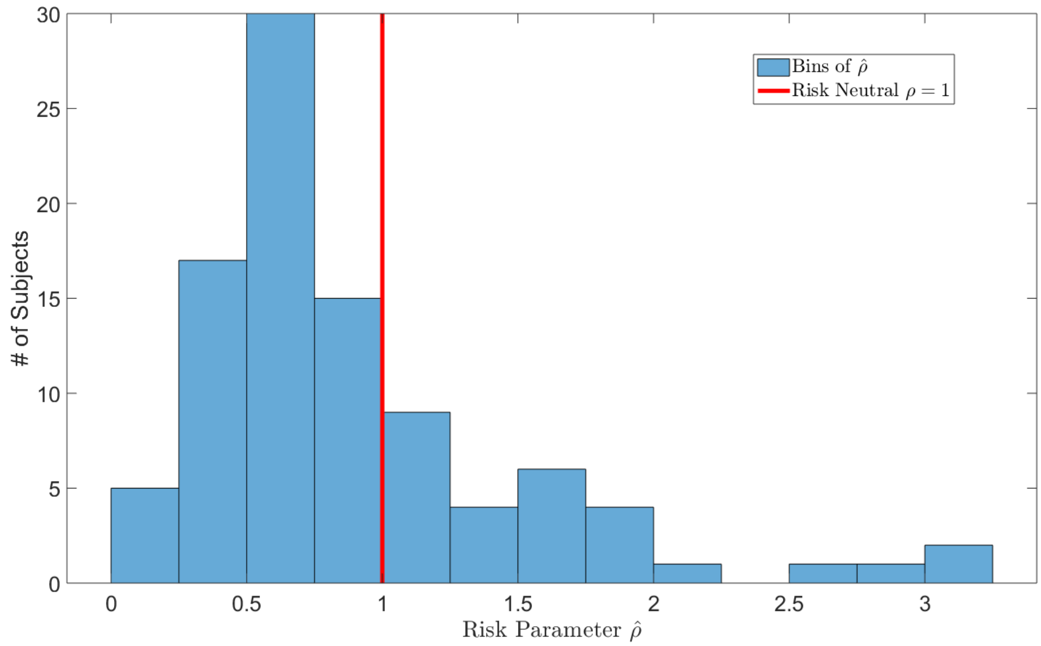

Figure 4: Histogram of Risk Parameter $\widehat{\rho}_i$

(i.e., $\hat{\sigma}_i > 1$). However, this figure does not provide any information for the statistical significance of the estimates.

Figure 5 presents the scatter plot of the estimate of $\hat{\sigma}_i$ and the corresponding t-statistic for the null hypothesis of risk neutrality, i.e., $\rho_i = 1$. The horizontal red lines represent the 5% critical values of one-sided tests. Therefore, dots in the top right region correspond to subjects for which we reject the null hypothesis of risk neutrality in favor of the alternative hypothesis of risk loving preferences. Similarly, dots in the bottom left area correspond to subjects for which we reject the null hypothesis of risk neutrality in favor of the alternative hypothesis of risk aversion. This figure shows that 67 subjects have an estimate $\widehat{\rho}_i < 1$ and for 49 of them it is significantly lower than 1 at the 5% level. Also, we find that $\widehat{\rho}_i > 1$ for 28 individuals, for 5 of them we find that significant risk loving preferences at 5% significance level.

(vi) Relationship between risk aversion and beliefs

We investigate the relationship between risk preferences and beliefs. A potential source of biased beliefs is that a subject may over-estimate the similarity between his preferences towards risk and those of other subjects. For instance, a risk loving individual may believe other subjects are also

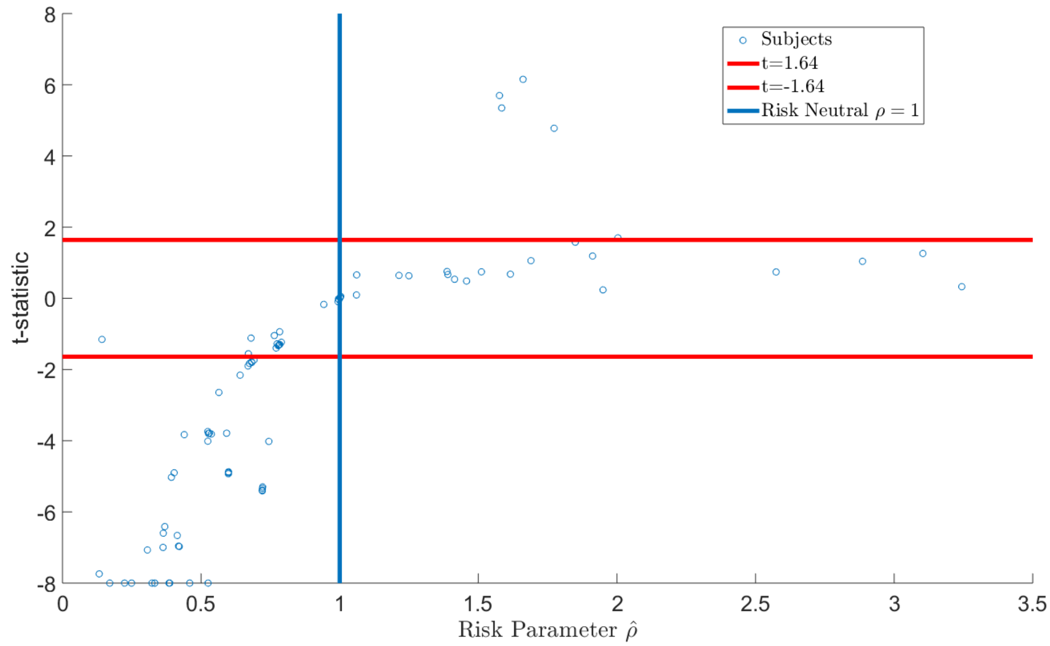

Figure 5: Point Estimates of $\widehat{\rho}$ and t-statistic

willing to take high risk, and this can generate an over-estimate of the coordination probability. To investigate this hypothesis, we first calculate the belief bias $\hat{B}_i(\mathbf{m}) - P^{co}(\mathbf{m})$ for each subject i and treatment $\mathbf{m}$, and then we average them at the subject level using the 30 treatments / observations per subject. This statistic $\overline{(\hat{B}_i - P^{co})}$ measures how much individual i's beliefs depart from unbiased beliefs. Figure 6 plots the scatter of these average belief biases and the estimated risk preference $\widehat{\rho}_i$. Most subjects tend to over-predict the probability of coordination, as reflected by a positive average belief bias. Interestingly, the magnitude of the over-prediction is positively correlated with the individual's risk preference. This correlation is equal to 0.5137 and is significant at the 1% significance level. This evidence is consistent with the interpretation that an individual who prefers risk tends to over-estimate other individuals' risk preference and consequently over-estimates the coordination probability.

(vii) Testing the validity of elicited beliefs

Heinemann, Nagel and Ockenfels (2009) also conducted a follow-up experiment with the same design but added an additional belief elicitation process that directly asked subjects to report their beliefs.[15] This follow-up experiment was done in Bonn and 40 subjects participated. We apply our

[15] The authors offered incentives to report true beliefs by rewarding subjects through a quadratic scoring rule. See

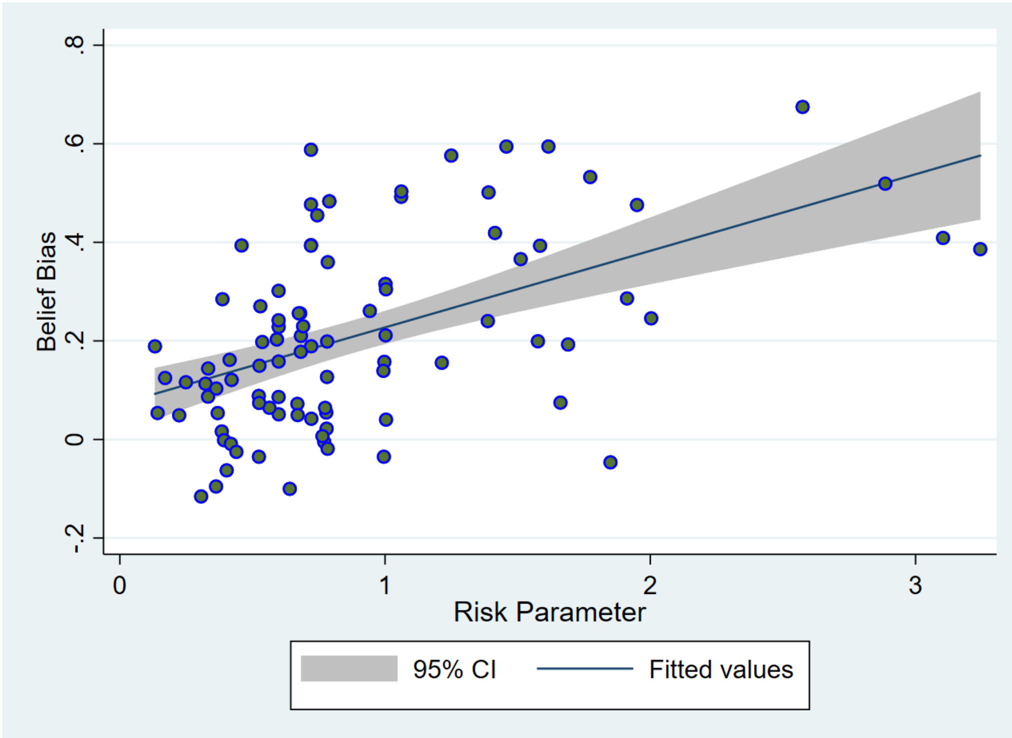

Figure 6: Average Belief Bias and Risk Parameter $\widehat{\rho}_i$

method to these data to test the validity of elicited beliefs: that is, to test whether subjects best respond to their reported beliefs.

Heinemann, Nagel and Ockenfels studied two sessions of elicitation methods with 20 subjects in each session. In the "global" session, each subject reports her/his subjective probability that the sum of the other players who coordinate is greater or equal to the corresponding threshold. In the "individual" session, a subject reports his/her subjective probability that a randomly chosen individual will choose the risky action. Since this reported belief on individual behavior does not correspond to the definition of $B_i(\mathbf{m})$, we need to calculate each subject's belief on aggregate behavior. We assume that each subject views other individuals as independent and identical draws from their reported belief. Under these conditions, $B_i(\mathbf{m})$ is equal to $1 - BIN[\phi(G-1) - 1, (G-1), 1 - b_i]$, where $BIN(\cdot)$ is the CDF of binomial distribution, and b_i is subject i's reported belief on individual behavior. As explained above, we exclude subjects whose actions are fixed when ϕ or m_s vary. This gives us 16 and 13 subjects in the "global" and the "individual" sessions.

For each subject, we apply a likelihood ratio test for the null hypothesis that the subject's belief

Heinemann, Nagel and Ockenfels (2009) for more details.

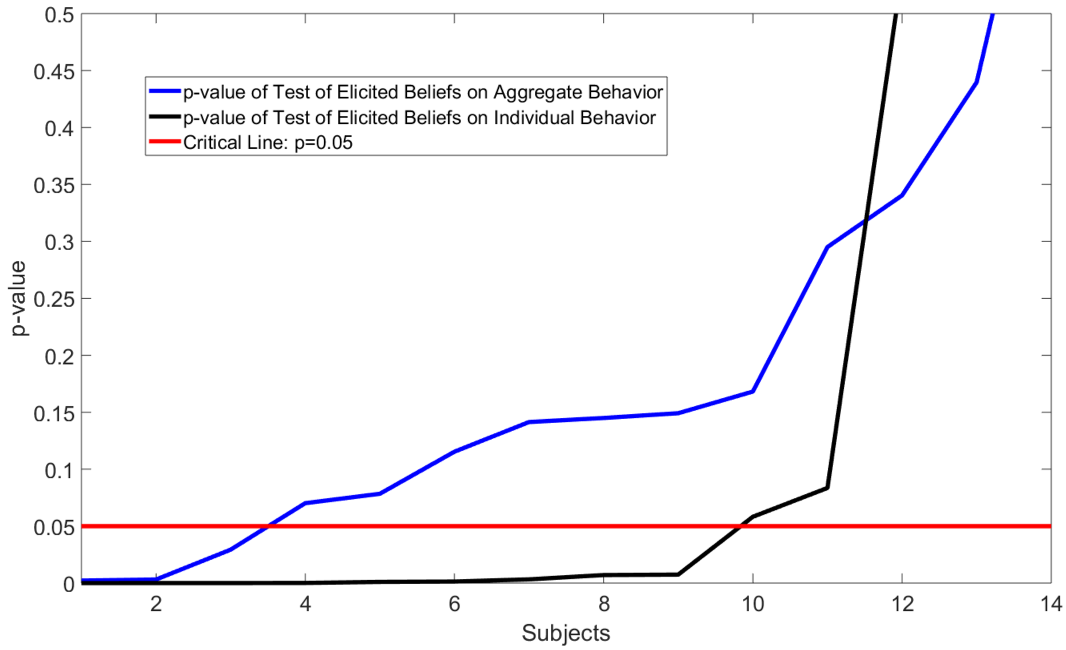

Figure 7: Inverse CDF of p-value of Belief Elicitation Test

is equal to the self-reported belief. Figure 7 plots the inverse CDF of the p-values of these individual tests. The blue and black lines show the results in "global" and "individual" sessions respectively. We find that most of subjects respond to their reported beliefs. In the "global" session, only 3 out of 16 subjects have a p-value less than 0.05. With Bonferroni correction, the hypothesis that all individuals best respond to elicited belief cannot be rejected at 1%. This is consistent with belief elicitation literature that shows elicited beliefs are consistent with individuals' actions. See Schotter and Trevino (2014) and Schlag et al. (2015).

In contrast, "individual" elicited beliefs explain the data substantially worse than "global" elicited beliefs. This is reflected by a lower black line than the blue one in Figure 7. In the "individual" session, 9 out of 13 subjects have a p-value less than 0.05. With Bonferroni correction, the hypothesis that all subjects best respond to elicited belief is rejected at 1% significance level.

There are at least two explanations for the rejection of elicited belief on individual behavior. First, subjects may believe that individuals are heterogeneous and/or their actions are correlated. If that were the case, our construction of $B_i(\mathbf{m})$ using the Binomial distribution would be incorrect. A second possible explanation is that subjects may make mistakes when calculating $B_i(\mathbf{m})$. This seems plausible as the Binomial distribution is cumbersome to evaluate, especially in an experimental environment where a calculator is not provided.

(viii) Other-regarding preferences

The existence of other-regarding preferences seems a very relevant concern in this application. An individual may get utility from her contribution to the public good because it has a positive effect on other individuals' utilities. Our identification result in Proposition 3 requires asymmetry in the matrix of monetary payoffs of the players, and therefore, it does not apply to the coordination game in this empirical application because monetary payoffs are symmetric in this game. Nevertheless, we think that other-regarding preferences is a potential issue in this empirical application and we have addressed it as follows.

The experiment in the coordination game includes treatments where subjects participate in a single-agent decision problem – a choice between two lotteries – with very similar structure as the coordination game – a safe choice and a risky choice – but where the probability distribution in the lottery for the risky choice – that corresponds to the probability of successful coordination in the game – is exogenous and known to the subjects. In these single-agent treatments, subjects' decisions do not depend on other-regarding preferences, simply because there are not "others" playing the game and receiving payoffs. Therefore, if a subject has a positive preference for the amount of money received by other players, then the probability that this subject chooses the risky alternative should be larger when playing the actual game than when playing the single-agent problem, provide they have the same lottery for the risky choice. Interestingly, for every subject in this experiment we find that either there is not significant difference between these choice probabilities or the difference goes in the opposite direction: subjects are more likely to take the risky alternative in the single-agent model than in the game even when the true probability of "successful coordination" in the game is higher than in the single-agent treatment. Note also that this is despite we find that most subjects have over-optimistic beliefs about the probability of successful coordination in the actual game

Based on this result, we think that our empirical results in the coordination game seem to be robust to potential other-regarding preferences.

7 Conclusion

This paper studies the non-parametric identification of players' preferences and beliefs using experimental data. With an exclusion restriction that affects a player's belief but has no impact on her preference, we show a function of beliefs is identified. This function can be used to test different belief restrictions, such as unbiased belief assumption and the validity of elicited beliefs. In addition, if the researcher imposes one of above belief restrictions in only two treatments, players' beliefs in all other treatments and preferences are non-parametrically identified. These identification results are obtained at individual level. The exclusion restriction required for identification can be either the variation of a single player's monetary payoff matrix or a change of payoff allocation rule in a coordination game.

These exclusion restrictions can be easily generated by experimental data. Therefore, our identification results provide a guidance on the design of the experiment if the researcher is interested in estimating subjects' preferences and beliefs. The experiment should include at least three treatments where the monetary payoffs of a player remain constant while the payoffs of other players vary.

We apply our test and identification results to experimental data from a matching pennies game (Goeree and Holt, 2001) and a coordination game (Heinemann, Nagel and Ockenfels, 2009). In the matching pennies game, subjects tend to correctly predict other players' behaviors when other players' monetary payoffs vary. In the coordination game, the null hypothesis of unbiased belief is rejected for a majority of subjects. Moreover, there is substantial individual heterogeneity among subjects' preferences. A majority of individuals are risk averse while some are risk loving. Finally, incorrectly imposing the linear utility assumption leads an over-rejection of the unbiased belief test.

References

[1] Aguirregabiria, V., and C-Y. Ho (2012): "A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments," *Journal of Econometrics*, 168(1), 156-173.

[2] Aguirregabiria, V., and A. Magesan (2019): "Identification and Estimation of Dynamic Games when Players' Beliefs Are not in Equilibrium," *Review of Economic Studies.* Forthcoming.

[3] Anderson, S., J. Goeree, and C. Holt (2001): "Minimum-Effort Coordination Games: Stochastic Potential and Logit Equilibrium," *Games and Economic Behavior*, 34(2), 177-199.

[4] Bajari, P., H. Hong, J. Krainer, and D. Nekipelov (2010): "Estimating Static Models of Strategic Interactions," *Journal of Business & Economic Statistics*, 28(4), 469-482.

[5] Bajari, P., H. Hong, and D. Nekipelov (2013): "Game Theory and Econometrics: A Survey of Some Recent Research." *Advances in Economics and Econometrics*, 10th World Congress of the Econometric Society. Vol. 3.

[6] Berry, S. and E. Tamer (2006): "Identification in Models of Oligopoly Entry," in Blundell, R., Newey, W. and Persson, T. (eds) *Advances in Economics and Econometrics* (Cambridge: Cambridge University Press) 46-85.

[7] Brown, A., C. Camerer, and D. Lovallo (2013): "Estimating Structural Models of Equilibrium and Cognitive Hierarchy Thinking in the Field: The Case of Withheld Movie Critic Reviews," *Management Science*, 59, 733-747.

[8] Camerer, C., T. Ho, and J. Chong (2004): "A Cognitive Hierarchy Model of One-Shot Games," *Quarterly Journal of Economics*, 119, 861-898.

[9] Cheung, Y. W., and D. Friedman (1997): "Individual Learning in Normal Form Games: Some Laboratory Result." *Games and Economic Behavior*, 19 (1): 46-76.

[10] Chmura, T., S. Goerg, and R. Selten (2012)" "Learning in Experimental 2×2 Games," *Games and Economic Behavior*, 76, 44-73.

[11] Costa-Gomes, M., and V. Crawford (2006): "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96, 1737-1768.

[12] Costa-Gomes M., and G. Weizsacker (2008): "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, 75, 729-762.

[13] Crawford, V., M. Costa-Gomes, and N. Iriberri (2013): "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 51(1), 5-62.

[14] Crawford, V., and H. Haller (1990): "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, 58(3), 571-595.

[15] Crawford, V., and N. Iriberri (2007): "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?", *Econometrica*, 75, 1721-1770.

[16] Doraszelski, U., G. Lewis, and A. Pakes (2018): "Just Starting Out: Learning and Equilibrium in a New Market," *American Economic Review*, 108, 565-615.

[17] Ellickson, P. and S. Misra (2008): "Supermarket Pricing Strategies," *Marketing Science*, 27(5), 811-828.

[18] Fehr E. and K. Schmidt, (1999): "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, 114, 817-868.

[19] Feltovich, N. (2000): "Reinforcement-based vs. Belief-based Learning Models in Experimental Asymmetric-information Games," *Econometrica*, 68, 605-641.

[20] Goeree J. and C. Holt, (2001): "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91, 1402-1422.

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

[22] Goeree J., C. Holt, and T. Palfrey (2005): "Regular Quantal Response Equilibrium," *Experimental Economics*, 8, 347-367.

[23] Goldfarb, A., and M. Xiao (2011): "Who Thinks About the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets," *American Economic Review*, 101, 3130 3161.

[24] Grether, D. M. (1992): "Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence," *Journal of Economic Behavior & Organization*, 17 (1), 31-57.

[25] Guerra, G. and Zizzo, D. (2004): "Trust Responsiveness and Beliefs," *Journal of Economic Behavior & Organization*, 55 (1), 25-30.

[26] Harrison, G., and E. Rutström (2008): "Risk Aversion in the Laboratory," in James C. Cox, Glenn W. Harrison (ed.) *Risk Aversion in Experiments (Research in Experimental Economics, Volume 12)* Emerald Group Publishing Limited, pp.41 - 196.

[27] Heinemann F., R. Nagel and P. Ockenfels (2009): "Measuring Strategic Uncertainty in Coordination Games," *Review of Economic Studies*, 76, 181-221.

[28] Hossain T., and R. Okui (2013): "The Binarized Scoring Rule," *Review of Economic Studies*, 80, 984-1001.

[29] Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk." *Econometrica,* 47, 263-293.

[30] Karni, E. (2009): "A Mechanism for Eliciting Probabilities," *Econometrica*, 77, 603-606.

[31] Kline, B. (2018): "An Empirical Model of Non-Equilibrium Behavior in Games," *Quantitative Economics*, 9(1), 141-181.

[32] Mckelvey, R., and T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.

[33] Mckelvey, R., and T. Palfrey (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9-41.

[34] Mckelvey, R., Palfrey, T., and R. Weber (2000): "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2× 2 Games with Unique Mixed Strategy Equilibria," *Journal of Economic Behavior & Organization*, 42, 523-548.

[35] Melo, E., K. Pogorelskiy, and M. Shum (2019): "Testing the Quantal Response Hypothesis," *International Economic Review*, 60(1), 53-74.

[36] Morris, S., and H. Shin (2002): "Social Value of Public Information," *American Economic Review*, 92, 1521-1534.

[37] Morris, S., and H. Shin (2004): "Coordination Risk and the Price of Debt," *European Economic Review*, 48, 133-153.

[38] Nyarko Y., and A. Schotter (2002): "An Experimental Study of Belief Learning Using Elicited Beliefs," *Econometrica*, 70, 971-1005.

[39] Ochs, J. (1995): "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.

[40] Offerman, T., Sonnemans, J., Van De Kuilen, G., & Wakker, P.P. (2009): "A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes," *Review of Economic Studies*, 76 (4), 1461-1489.

[41] Palfrey, T., and S. Wang (2009): "On Eliciting Beliefs in Strategic Games," *Journal of Economic Behavior and Organization*, 71, 98-109.

[42] Rogers, B., T. Palfrey, and C. Camerer (2009): "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," *Journal of Economic Theory*, 144(4), 1440-1467.

[43] Roth, A., and M. Malouf (1979): "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 574-594.

[44] Rutström, E., and N. Wilcox (2009): "Stated Beliefs versus Inferred Beliefs: A Methodological Inquiry and Experimental Test," *Games and Economic Behavior*, 67, 616-632.

[45] Schlag, K. H., Tremewan, J. and Van Der Weele, J.J. (2015): "A Penny for Your Thoughts: A Survey of Methods for Eliciting Beliefs," *Experimental Economics*, 18 (3), 457-490.

[46] Schotter, A. and Trevino, I. (2014): "Belief Elicitation in the Laboratory," *Annual Review of Economics*, 6 (1), 103-128.

[47] Selten, R. and Chmura, T. (2008): "Stationary Concepts for Experimental 2×2 Games," *American Economic Review*, 98, 938-966.

[48] Seltern, R., Sadrieh, A., and Abbink, K. (1999): "Money Does Not Induce Risk Neutral Behavior, But Binary Lotteries do Even Worse," *Theory and Decision*, 46, 211-249.

[49] Smith, C. (1961): "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society. Series B,* 23 (1), 1-37.

[50] Sweeting, A. (2009): “The Strategic Timing Incentives of Commercial Radio Stations: An Empirical Analysis Using Multiple Equilibria,” *The RAND Journal of Economics*, 40(4), 710-742.

[51] Van Huyck,. R. Battalio and R. Beil (1990): “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80(1), 234-48.