Identification of Biased Beliefs in Games of Incomplete Information Using Experimental Data

Victor Aguirregabiria[*]
University of Toronto and CEPR

Erhao Xie[*]
University of Toronto

April 28, 2017

Abstract

This paper studies the identification of players' preferences and beliefs in empirical applications of discrete choice games using experimental data. The experiment comprises a set of games with similar features (e.g., two-player coordination games) where each game has different values for the players' monetary payoffs. Each game can be interpreted as an experimental treatment group. The researcher assigns randomly subjects to play these games and observes the outcome of the game as described by the vector of players' actions. The researcher is interested in the nonparametric identification of players' preferences (utility function of money) and players' beliefs about the expected behavior of other players, without imposing restrictions such as unbiased or rational beliefs or a particular functional form for the utility of money. We show that, given a particular design of the matrices of payoffs in the treatments of the experiment, the hypothesis of unbiased/rational beliefs is testable. We propose a nonparametric test of this null hypothesis. We apply our method to two sets of experiments conducted by Goeree and Holt (2001) and Heinemann, Nagel and Ockenfels (2009). Our empirical results suggest that in the matching pennies game, a player is able to correctly predict other player's behavior. In the public good coordination game, our test can reject the null hypothesis of unbiased beliefs when the payoff of the non-cooperative action is relatively low.

Keywords: Empirical Games; Randomized experiments; Laboratory experiments; Testing for biased beliefs; Identification; Multiple equilibria; 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. Department of Economics, University of Toronto. 150 St. George Street, Toronto, Ontario, M5S 3G7.
Email: erhao.xie@mail.utoronto.ca

[*]We would like to thank Martin Burda, Evan Calford, Robert Gazzale, Tanjim Hossain, Arvind Magesan, Eduardo Souza-Rodrigues, Colin Stewart, Junichi Suzuki, and especially Ryan Webb for multiple comments and helpful discussions. We also want to thank comments from participants at the 2016 annual conference of the Economics Science Association (ESA). 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. Knowing both terms is essential to understand player's behavior and make counterfactual prediction. However, both terms are unobserved to researchers and it is an empirical challenge to distinguish them. To fill this gap, we study the joint identification of preferences and beliefs using experimental data from the outcomes of multiple realizations of games.

In the experimental economics literature, researchers design laboratory experiments and generate experimental data to study player's behavior in games. From the point of view of identification, there are clear advantages of having data from a controlled experiment. In particular, the design of the experiment determines players' monetary payoffs such that these payoffs are perfectly known to the researcher. Most of the experimental games literature has exploited this advantage using three alternative approaches. A first approach imposes the restriction that the utility function is equal to the monetary payment (plus a mean-zero private information variable, henceforth, linear utility assumption) and then identifies beliefs using choice data. Examples of this approach include Cheung and Friedman (1997) who estimate a belief-learning process in a repeated game, and Nyarko and Schotter (2002) who compare beliefs estimated in this way with elicited beliefs. A second approach assumes that players form equilibrium/unbiased beliefs. Under this assumption, a player's belief is identified by the probability distribution of other players' actions given the information available and the utility function of money can be identified thereafter using 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 another stream, experimental researchers measure subjects' beliefs in games by elicitation process. Recently, Karni (2009), Offerman et al. (2009) and Hossain and Okui (2013) developed different mechanisms that can correctly elicit player's belief regardless of risk preference.[1] See Schotter and Trevino (2014) and Schlag et al. (2015) for recent reviews of prominent elicitation methods and their practical issues.

This paper proposes an alternative approach to identify preferences and beliefs in discrete games of incomplete information using data from a controlled experiment. Our approach relaxes the assumption of unbiased or equilibrium beliefs and it does not impose any parametric restriction

[1] Karni's (2009) mechanism involves randomized payment and such idea was first proposed by Smith (1961). In addition, a similar design is conducted by Grether (1992).

on the functional form of the utility function nor needs information of elicited beliefs. Relaxing these restrictions is important in different empirical applications.

First, there are multiple reasons why players may have biased beliefs. Playing a Bayesian Nash Equilibrium strategy requires player to determine other players' equilibrium strategy and to be able to integrate it over the other player's private information. Such calculation is burdensome, and human cognition limits may preclude the equilibrium behavior, particularly in one-shot experimental games. Even in the absence of cognition limits, in games with multiple equilibria players may have uncertainty about which equilibrium strategy will be chosen by other players. A player may believe that the selected equilibrium is A, while other player may think that it is B. This type of 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.

Second, the linear utility assumption places strong restrictions on subjects' preferences that are at odds with important empirical findings in the experimental literature. Kahneman and Tversky (1979) note that individuals may respond to loss more sensitively than to gains. Harrison and Rutström (2008) show that risk aversion is prevalent even for the payoff scale typically found in experimental data. The linear utility assumption also rules out social preferences and heterogeneity across players in their marginal utility of money.[2] Our framework treats a player's utility as an unknown unrestricted function of her monetary payoff and is able to capture both risk preference and loss aversion.

Third, even though there is substantial empirical evidence showing that elicited beliefs are consistent with individuals' actions in corresponding tasks,[3] there exist some practical issues. Perhaps, the most serious problem is that the process for the elicitation of beliefs can affect players' behavior in games. Schotter and Trevino (2014) denote this problem as *Heisenberg problem* using as 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).[4] In contrast, our method does not require elicitation data and consequently, avoid these potential issues.

To avoid the estimation biases and the misleading results associated with the failure of these

[2] For examples of social preferences such as fairness, see Güth, Schmittberger and Schwartz (1982), Kahneman, Knetsch and Thaler (1986) and Fehr and Schmidt (1999), among others.

[3] To 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

[4] 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 hedging problem and the complexity of the methods. The empirical evidence on the importance of hedging is also mixed. See section 3 in Schotter and Trevino (2014) or section 4 in Schlag et al. (2015) for more details.

assumptions, we treat both utilities and beliefs as unrestricted (nonparametric) functions to be estimated. In this paper, we focus on the identification and estimation of *population average utilities and beliefs*. Obtaining precise estimates of individual specific beliefs/preferences requires an experimental design where each subject makes plays the same game multiple times without information feedback. Such experimental design could be costly and is rarely conducted.[5] Instead, our identification approach can be applied to experimental data where each subject is randomly assigned to only one treatment (game). As a consequence, our focus on population average beliefs/preferences can avoid the undesired effects of order and repetition, and potential hedging problems. Most importantly, population average utilities and beliefs are useful parameters for researchers and policy makers. They can be used to obtain population average effects of factual and counterfactual policy changes. Researchers can be also interested in testing rational/unbiased beliefs, or risk/loss aversion in preferences, at the average population level (e.g., Goeree and Holt, 2001).

Our identification results and tests are based on an exclusion restriction that can be easily generated by the researcher in the design of the experiment. Suppose that each individual in the sample is randomly assigned to play one of K different two-player games. The researcher designs the monetary payoff matrices in these K games such that the payoffs of the column player vary across the games but the payoffs of the row player are the same in the K games.[6] This variation across games in the payoff matrix is what we describe as our exclusion restriction in the sense that it does not affect the payoff function of the row player but it can affect the beliefs of this player about the behavior of the column player. Under this exclusion restriction, the variation across the K games in the empirical distribution of the actions of the row player provides information about this player's beliefs in these games. Without any assumptions on players' beliefs in any game and following an argument similar to Aguirregabiria and Magesan (2015), we show that this exclusion restriction identifies a function of beliefs. This identification results can be used to test different assumptions on beliefs such as: (a) population average unbiased (equilibrium) beliefs; (b) the validity of elicited beliefs, if these are available; and (c) monotonicity of the beliefs function with respect to monetary payoff of the other player(s). Our framework allows for an optimization error in players' "best responses", as in Quantal Response Equilibrium (QRE) proposed by McKelvey and Palfrey (1995).

[5] Note that if such data exists, then our approach provides indetification (i.e., consistent estimation) of utilities and beliefs at the individual level.

[6] We can find this experimental design of monetary payoffs in important studies such as Goeree and Holt (2001), Ochs (1995), McKelvey, Palfrey and Weber (2000) and Heinemann, Nagel and Ockenfels (2009), among others. None of these papers focuses on the joint nonparametric identification of subjects' preferences and beliefs.

The complete identification of utility and beliefs functions requires some additional restrictions. These restrictions are weaker than linear utility assumptions and equilibrium assumptions. In a two-player binary choice game, the researcher needs to impose two restrictions on the beliefs or payoffs. We discuss different forms that these restrictions can take and how the choice of these restrictions can be informed by our tests on beliefs. For instance, the researcher may assume that elicited beliefs are valid or that beliefs are unbiased at two of the K games. How to choose these two games is also an important decision for the researcher, and in this paper, we discuss how our population average rationality test can provide informative guidance on researcher's decision.

Our nonparametric specification of risk preferences has potential advantages over other methods in the experimental literature. In an influential study, Roth and Malouf (1979) propose linearizing the utility function by assigning the payoff as the probability of winning a fixed reward. This mechanism has been applied by Ochs (1995) and Feltovich (2000), among others. Selten et al. (1999) have raised skepticism about the validity of this mechanism. Goeree, Holt and Palfrey (2003), using experimental data from Ochs (1995), show that this mechanism fails to linearize the utility function. The general validity of Roth and Malouf's approach seems unknown in the literature. Another common method consists in eliciting players' risk preference using a lottery choice with a known objective probability distribution. Such a method is used in Heinemann, Nagel, and Ockenfels (2009), among others. Elicitation introduces an additional cost in the implementation of the experiment and, as mentioned above, there may be different sources of bias in the elicitation of preferences and beliefs. A third approach involves estimating a common parametric function for the utility of money, e.g., a CRRA utility function. This is the approach used by Goeree, Holt and Palfrey (2003). As usual with parametric specification, the mis-specification of utility function can generate bias in estimates of beliefs such that, for instance, the researcher may spuriously conclude that players' beliefs are biased (not in equilibrium).

We apply our approach to estimate two types of games that have received substantial 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).[7] In the matching pennies game, our estimation results cannot reject that players, on average, can correctly predict other players' behavior. In coordination games, we find that subjects have biased beliefs when the monetary payoff for the safe action is low. As such payoff increases, this bias in beliefs declines and becomes

[7] As mentioned above, the design of monetary payoff matrices in Ochs (1995) and McKelvey, Palfrey and Weber (2001) also fits the requirement of our approach. In these two experiments, subjects play the same game for multiple rounds and we choose not to use their data to avoid modeling learning behavior.

not statistically different to zero. Our estimated payoff function is convex when the monetary payoff is low and becomes concave as the monetary payoff increases. This finding suggests that the commonly imposed globally concave utility functions, such as CRRA or logarithmic functions, are not able to capture subject's preference, and a non-parametric specification of the payoff function is more appropriate in this application.

The remainder of this paper is organized as follows. Section 2 describes the model and the experimental design that generates the exclusion restriction. Section 3 presents our identification results. Section 4 describes the two experimental data sets that we use in our empirical analysis and presents the estimation procedure and our empirical results. We summarize and conclude in section 5.

2 Model

2.1 Basic model

For simplicity, we present here a model with two players and binary choice. In section 4.2.1, we show that our approach can be extended to games with multiple players. There are two roles for players in the game: the "row" player (R) and the "column" player (C). We index player roles by $i, j \in \{R, C\}$. Let $a_R \in \{0,1\}$ and $a_C \in \{0,1\}$ be the actions and choice sets for the "row" player and for the "column" player, respectively. Players take their actions simultaneously to maximize their respective expected payoffs. The payoff function of player i is:

$$\Pi_i(a_i, a_j) = \pi\left(m_i(a_i, a_j)\right) + \varepsilon_i(a_i) \tag{1}$$

$m_i(a_i, a_j)$ is the monetary payoff of player i when players take actions (a_i, a_j). $\pi(\cdot)$ is a real-valued function that represents the population average utility of money. The matrix of monetary payoffs and the utility function $\pi(\cdot)$ are common knowledge to all the players. $\varepsilon_i(a_i)$ represents player i's deviation from the average utility, and it is idiosyncratic for each individual player and is private information of the individual; furthermore, it is independently distributed across subjects with a probability distribution that is public information for all the players. As we explain below, these private information variables can be also interpreted as *optimization errors*, along the line of the *Quantal Response Equilibrium* concept proposed by Mckelvey and Palfrey (1995, 1998).

Player i does not know the values of the variables $\varepsilon_j(0)$ and $\varepsilon_j(1)$ which are private information of player j. 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 represent the subjective belief of player i about the probability that the other player chooses action $a_j = 1$. Then, player i's expected payoff of action a_i is:

$$\Pi_i^e(a_i, B_i) = [1 - B_i] \;\; \pi\left(m_i(a_i, 0)\right) + B_i \;\pi\left(m_i(a_i, 1)\right) + \varepsilon_i(a_i) \tag{2}$$

Given utility function and beliefs, players maximize their expected payoffs. The best response of player i is alternative $a_i = 1$ if

$$\begin{aligned} &[1 - B_i]\, \pi\left(m_i(1, 0)\right) + B_i \;\pi\left(m_i(1, 1)\right) + \varepsilon_i(1) \\ &\geq [1 - B_i]\, \pi\left(m_i(0, 0)\right) + B_i \;\pi\left(m_i(0, 1)\right) + \varepsilon_i(0) \end{aligned} \tag{3}$$

Let $Q_i\left(\mathbf{m}_i, B_i\right)$ be the probability that player i chooses action $a_i = 1$ given beliefs B_i and matrix of payoffs $\mathbf{m}_i \equiv \{m_i(0, 0),\, m_i(0, 1),\, m_i(1, 0),\, m_i(1, 1)\}$, and integrated over her private information variables. Integrating the best response condition (3) over the private information, we obtain player i's best response probability function:

$$Q_i\left(\mathbf{m}_i, B_i\right) = F_{\widetilde{\varepsilon}}\left(\alpha_\pi(\mathbf{m}_i) + \beta_\pi(\mathbf{m}_i)\;B_i\right) \tag{4}$$

where $F_{\widetilde{\varepsilon}}$ is the CDF of $\widetilde{\varepsilon}_i \equiv \varepsilon_i(0) - \varepsilon_i(1)$, $\alpha_\pi\left(\mathbf{m}_i\right) \equiv \pi\left(m_i(1, 0)\right) - \pi\left(m_i(0, 0)\right)$, and $\beta_\pi(\mathbf{m}_i) \equiv [\pi(m_i(1, 1)) - \pi(m_i(0, 1))] - [\pi(m_i(1, 0)) - \pi(m_i(0, 0))]$. The payoff matrix and the utility function are such that $\beta_\pi(\mathbf{m}_i) \neq 0$, i.e., the model is a game and not a single-agent decision problem.

This model includes the Bayesian Nash Equilibrium as a particular case.

Definition. *The model is consistent with Bayesian Nash Equilibrium (BNE) if: (a) players' choice probabilities are equal to their best responses, i.e.,* $Q_i\left(\mathbf{m}_i, B_i\right) = F_{\widetilde{\varepsilon}}\left(\alpha_\pi(\mathbf{m}_i) + \beta_\pi(\mathbf{m}_i)\;B_i\right)$*, and* $Q_j\left(\mathbf{m}_i, B_j\right) = F_{\widetilde{\varepsilon}}\left(\alpha_\pi(\mathbf{m}_j) + \beta_\pi(\mathbf{m}_j)\;B_j\right)$*; and (b) players' beliefs about other players' actions are equal to these players' best response probabilities:* $B_i = Q_j\left(\mathbf{m}_j, B_j\right)$ *and* $B_j = Q_i\left(\mathbf{m}_i, B_i\right)$*.* ■

This Bayesian Nash Equilibrium is also consistent with the *Quantal Response Equilibrium (QRE).* In particular, private information variables $\varepsilon_i(a_i)$ can be interpreted as optimization errors. Importantly, regardless we interpret ε's as optimization errors or heterogeneous preferences, or a combination of both, we assume that they are private information only known by the own player such that we have a game of incomplete information.

Our framework relaxes two assumptions in existing empirical applications of QRE and BNE models. First, subjects are not restricted to have correct/equilibrium beliefs when they play the game. They can have biased beliefs, and the bias can be heterogeneous across subjects. Second, the payoff function π can be different than the monetary payoff, and we treat this function as an unknown to be estimated from the data.[8]

[8] See Mckelvey and Palfrey (1995), Mckelvey, Palfrey and Weber (2000) and Goeree, Holt and Palfrey (2003) among

2.2 Experimental design and subject heterogeneity

Given the game described above, the experimental researcher chooses T different matrices of monetary payoffs that are indexed by $t \in \{1, 2, ..., T\}$. Let $\mathbf{m}_t = (\mathbf{m}_{Rt}, \mathbf{m}_{Ct})$ be the $t-th$ matrix of monetary payoffs, where $\mathbf{m}_{Rt}$ and $\mathbf{m}_{Ct}$ represent the matrices of payoffs for the row and the column player, respectively.

There is a sample of N subjects indexed by $n \in \{1, 2, ..., N\}$. Subjects are randomly assigned to $2T$ possible treatments. A treatment in this experiment is defined as a pair (i, t), where t is the index of the payoff matrix in the game the subject has to play, and $i \in \{R, C\}$ represents the player role of the subject in that game (i.e., either row or column player). The random allocation of players to treatments is anonymous such that each subject does not have any information about who is the other subject she is playing against. This design where each subject is assigned only to one treatment can avoid the order effect, repetition effect, and potential hedging problem.

Once subjects have been allocated to treatments, they play their respective games. We use the categorical variable $d_n \in \{R, C\} \times \{1, 2, ..., T\}$ to represent the treatment (i, t) received by subject n, and the binary variable $a_n \in \{0, 1\}$ is used to represent the subject's actual choice in the game. Therefore, the data from this randomized experiment can be described in terms of the observations $\{d_n, a_n : n = 1, 2, ..., N\}$.

Subjects can be heterogeneous in preferences and beliefs. Variables $\varepsilon_{nit}(0)$ and $\varepsilon_{nit}(1)$ represent the idiosyncratic components of the payoff function for subject n if she is assigned to treatment (i, t). Similarly, the probability B_{nit} represents the subjective belief of subject n when assigned to treatment (i, t). We make the following assumption of additive separability and independence on subjects' heterogeneity in preferences and beliefs.

Assumption 1. *(A) All the heterogeneity in preferences across subjects is captured by the private information variables $\varepsilon_{nit}(0)$ and $\varepsilon_{nit}(1)$ that have zero mean, are independently distributed across subjects n and payoff matrices $\mathbf{m}_{jt}$. (B) Subject n's beliefs in treatment (i, t) are $B_{nit} = \overline{B}_{it} + \xi_{nit}$, where $\overline{B}_{it}$ represents the average beliefs in the population of subjects conditional on treatment (i, t), and ξ_{nit} is subject n's idiosyncratic component in beliefs that is private information of this subject, has zero mean, and it is independently distributed across subjects n and payoff matrices $\mathbf{m}_{jt}$.* ■

In Assumption 1, the condition that random variables $\varepsilon_{nit}(\cdot)$ and ξ_{nit} are independent across players is an implication of the randomized experiment and the anonymity of the assignment of

others. Goeree, Holt and Palfrey (2003) relax the assumption that the utility function π is equal to the monetary payoff and estimate a parametric model for this function. In this paper, we do not impose any functional form for the utility, other than being an increasing function.

subjects to treatments. Also, by construction, variables $\varepsilon_{nit}(\cdot)$ and ξ_{nit} have zero mean. Assumption 1(A) restricts preference heterogeneity $\varepsilon_{nit}(\cdot)$ to be independent of $\mathbf{m}_{jt}$. This holds true for any self-regarding preferences, as player j's monetary reward $\mathbf{m}_{jt}$ has no effect on player i's utility and consequently does not affect $\varepsilon_{nit}(\cdot)$.[9] The main restriction in Assumption 1 is the condition that beliefs heterogeneity ξ_{nit} is independent of monetary payoffs $\mathbf{m}_{jt}$. Note that Assumption 1(B) allows average beliefs $\overline{B}_{it}$ to be completely unrestricted, but restricts the distribution of ξ_{nit} to be the same across treatments t. This is not an implication of the randomized experiment.

Assumption 2. *Define the random variable* $\omega_{nit} \equiv \varepsilon_{nit}(0) - \varepsilon_{nit}(1) - \xi_{nit}\ \beta_{\pi}(\mathbf{m}_{it})$, *where* $\beta_{\pi}(\mathbf{m}_{it})$ *has been defined above. Conditional on* $\mathbf{m}_{it}$, *the random variable* ω_{nit} *has a probability distribution* $F_{\omega|\mathbf{m}_{it}}$ *that is known to the researcher and it is strictly monotonic in* $\mathbb{R}$. ■

ω_{nit} is a composite error that consists of both utility and belief heterogeneity. Under assumption 1, ω_{nit} is independent of $\mathbf{m}_{jt}$ and its distribution only depends on $\mathbf{m}_{it}$. Assumption 2 requires researchers to know such distribution while it may be unknown in practice. In our empirical applications in section 5, we try several distributional assumptions as robustness checks.

Consider subject n that has been assigned to treatment t as player i. Given that this subject has beliefs B_{nit}, her best response probability is $Q_{nit} \equiv Q_i(\mathbf{m}_{it}, B_{nit})$, and using the definition in equation (4) this best response probability is equal to $F_{\widetilde{\varepsilon}}(\alpha_{\pi}(\mathbf{m}_{it}) + \beta_{\pi}(\mathbf{m}_{it})\ B_{nit})$. This best response probability depends on the idiosyncratic beliefs of subject n, B_{nit}. Under Assumptions 1-2, we can integrate this best response probability function over the idiosyncratic component of beliefs, ξ_{nit}. We obtain the (average) *Conditional Choice Probability (CCP)* function:

$$P_{it} \equiv P_i\left(\mathbf{m}_{it}, \overline{B}_{it}\right) = F_{\omega|\mathbf{m}_{it}}\left(\alpha_{\pi}(\mathbf{m}_{it}) + \beta_{\pi}(\mathbf{m}_{it})\ \overline{B}_{it}\right) \tag{5}$$

Equation (5) is the key restriction of the model that we use to identify and estimate the (average) utility and beliefs functions.

There is a substantial empirical literature in behavioral and experimental economics that studies players' non-equilibrium behavior and heterogeneous beliefs. Level-k models by Nagel (1995) and Stahl and Wilson (1994, 1995) and the cognitive hierarchy model by Camerer, Ho, and Chong (2004) are some important contributions in this literature.[10] Our model relaxes some restrictions in these previous studies. We do not impose BNE, QRE, or level-k rationalizability, and ξ_{nit} captures heterogeneity in beliefs across players and across subjects in the same role. We also consider a nonparametric specification of the utility of money.

[9] This assumption may be restrictive when players have social preferences.

[10] For a survey of papers in this field, see Crawford, Costa-Gomes and Iriberri (2013).

3 Identification

The dataset consists of N observations $\{d_n, a_n\}$, one for each subject, where d_n represents the treatment received by subject n, and a_n is her action in the game. Each subject n is randomly assigned to one of the $2T$ treatments such that d_n is independent of the unobservables in ω_n.

Let π_i be the vector of payoff parameters for player i in the experiment, $\pi_i \equiv \{\pi(m_{it}(a_R, a_C)) : (a_R, a_C) \in \{0,1\}^2$ and $t = 1, 2, ..., T\}$. Similarly, let $\overline{\mathbf{B}}_i \equiv \{\overline{B}_{it} : t = 1, 2, ..., T\}$ be the vector of average belief parameters for player i in the experiment. The researcher is interested in using the experimental data to estimate (average) preferences and beliefs parameters π_R, π_C, $\overline{\mathbf{B}}_R$, and $\overline{\mathbf{B}}_C$.

Let $\mathcal{M}_T \equiv \{\mathbf{m}_t = (\mathbf{m}_{Rt}, \mathbf{m}_{Ct}) : t = 1, 2, ..., T\}$ be the set of payoff matrices in the T treatments of the randomized experiment. Assumption 3 establishes a condition on the set $\mathcal{M}_T$ that plays a fundamental role in our identification results.

Assumption 3. *The set $\mathcal{M}_T$ of payoff matrices in the randomized experiment is such that there are at least two treatments, say t_1 and t_2, such that: (A) player i has the same payoffs in the two treatments but the payoffs of player $j \neq i$ are different, i.e., $\mathbf{m}_{it_1} = \mathbf{m}_{it_2}$ and $\mathbf{m}_{jt_1} \neq \mathbf{m}_{jt_2}$; (B) player i's conditional choice probabilities vary across the two treatments, i.e., $P_{it_1} \neq P_{it_2}$.* ■

Assumption 3(A) establishes that the experimental design generates a particular variation in monetary payoffs across treatments: the payoff matrix of player j varies while the payoff matrix of player i remains constant. We show below that this condition provides an *exclusion restriction* that can be used to identify player i's beliefs from this player's observed behavior. Assumption 3(B) is a "Relevance condition" that is necessary for identification. Since the conditional choice probabilities P_{it_1} and P_{it_2} are identified from the data under mild conditions (see section 3.1 below), Assumption 3(B) is testable from the data. This assumption can be also interpreted as an implication of *Rationalizability*, i.e., player i knows that player j maximizes expected payoff given beliefs. Since player j's payoff matrix varies across treatments t_1 and t_2, player i's beliefs about player j's behavior also varies between the two treatments, and given that her own payoff matrix did not change, her actual behavior should be different as long as her behavior depends on beliefs.

We show below that under Assumptions 1 to 3 we can test for the null hypothesis of unbiased (equilibrium) beliefs without parametric assumptions on the utility function and average beliefs. Then, we present additional conditions for the full nonparametric identification of the model.

3.1 Tests of unbiased beliefs

Under the conditions in Assumption 1 the choice probabilities P_{it} are identified for every player-role and treatment (i,t). In particular, the probability P_{it} is equal to $\mathbb{E}[a_n \mid d_n = (i,t)]$ and we can use the following frequency estimator $\widehat{P}_{it}$ to consistently estimate P_{it}:

$$\widehat{P}_{it} = \frac{\sum_{n=1}^{N} a_n \ 1\{d_n = (i,t)\}}{\sum_{n=1}^{N} 1\{d_n = (i,t)\}} \tag{6}$$

where $1\{.\}$ is the indicator function. $\widehat{P}_{it}$ is the fraction of subjects who choose alternative $a = 1$ among all subjects who are assigned as player i in treatment t. As choice probabilities P_{it} are consistently estimated by $\widehat{P}_{it}$, we treat P_{it} as known in the proofs of our identification results.

Let $F^{-1}_{\omega|\mathbf{m}_{it}}(.)$ be the inverse function of the CDF of $F_{\omega|\mathbf{m}_{it}}$. This inverse function exits because the strict monotonicity of the CDF. Under Assumption 2, the inverse function $F^{-1}_{\omega|\mathbf{m}_{it}}(P_{it})$ can be consistently estimated for every treatment (i,t) by $F^{-1}_{\omega|\mathbf{m}_{it}}(\widehat{P}_{it})$. For notational simplicity, we use the variable S_{it} to represent $F^{-1}_{\omega|\mathbf{m}_{it}}(P_{it})$. The model implies that:

$$S_{it} = \alpha_{\pi}\left(\mathbf{m}_{it}\right) + \beta_{\pi}\left(\mathbf{m}_{it}\right) \ \overline{B}_{it} \tag{7}$$

Let t_1 and t_2 be the two treatments in Assumption 3. Let $\mathcal{T}_{i,t_1}$ be the subset of treatments in the experiment where player i has the same monetary payoffs as in treatment t_1 i.e., $\mathcal{T}_{i,t_1} \equiv \{t : \mathbf{m}_{it} = \mathbf{m}_{it_1}\}$. For any treatment $t \in \mathcal{T}_{i,t_1}$, we have that

$$S_{it} - S_{it_1} \quad = \quad \beta_{\pi}\left(\mathbf{m}_{it_1}\right) \ \left[\overline{B}_{it} - \overline{B}_{it_1}\right] \tag{8}$$

Assumption 3(B) and the strict monotonicity of the CDF $F_{\omega|\mathbf{m}_{it}}$ imply that $S_{it_2} - S_{it_1} \neq 0$. Therefore, given that $\beta_{\pi}(\mathbf{m}_{it_1}) \neq 0$, equation (8) implies that $\overline{B}_{it_2} - \overline{B}_{it_1} \neq 0$. Taking this into account, we have that for any treatment $t \in \mathcal{T}_{i,t_1}$,

$$\frac{S_{it} - S_{it_1}}{S_{it_2} - S_{it_1}} \quad = \quad \frac{\overline{B}_{it} - \overline{B}_{it_1}}{\overline{B}_{it_2} - \overline{B}_{it_1}} \tag{9}$$

This expression shows that, under assumptions 1-3, the observed behavior of subjects as player i identifies the beliefs ratio $(\overline{B}_{it} - \overline{B}_{it_1})/(\overline{B}_{it_2} - \overline{B}_{it_1})$ for any treatment $t \in \mathcal{T}_{i,t_1}$. That is, observed behavior can identify an object that depends only on beliefs and not on preferences. This result implies that the assumption of unbiased or equilibrium (average) beliefs is testable.

Under the restriction of equilibrium beliefs, the ratio $(\overline{B}_{it} - \overline{B}_{it_1})/(\overline{B}_{it_2} - \overline{B}_{it_1})$ should be equal to the ratio of the choice probabilities of the other player (subjects as player j), i.e., $(P_{jt} - P_{jt_1})/(P_{jt_2} - P_{jt_1})$. This provides a testable restriction.

Proposition 1. *Under Assumptions 1 to 3, for any treatment $t \in \mathcal{T}_{i,t_1}$, the hypothesis of equilibrium (unbiased) beliefs implies the restriction:*

$$\frac{S_{it} - S_{it_1}}{S_{it_2} - S_{it_1}} \quad = \quad \frac{P_{jt} - P_{jt_1}}{P_{jt_2} - P_{jt_1}} \tag{10}$$

with $S_{it} \equiv F^{-1}_{\omega|\mathbf{m}_{it}}(P_{it})$. Given that the choice probabilities P_{it} and P_{jt} are identified, this restriction is testable when the number of treatments in the set $\mathcal{T}_{i,t_1}$ is at least three. $\blacksquare$

For experiments where the payoff matrix of player i has a particular structure, it is possible to construct a test of unbiased beliefs that requires only two treatments in the set $\mathcal{T}_{i,t_1}$. Suppose that the matrix of monetary payoffs of player i is symmetric and diagonal-constant (Toeplitz matrix) such that $m_i(0,0) = m_i(1,1)$ and $m_i(0,1) = m_i(1,0)$. For instance, this is form of the payoff matrix in a matching pennies game. Under this condition, we have that $\beta_\pi(\mathbf{m}_{it}) = -2\ \alpha_\pi(\mathbf{m}_{it})$ and equation (7) becomes $S_{it} = \alpha_\pi(\mathbf{m}_{it})\ [1 - 2\ \overline{B}_{it}]$. Therefore, under Assumption 3, for treatments t_1 and t_2 we have that

$$\frac{S_{it_2}}{S_{it_1}} \quad = \quad \frac{1 - 2\ \overline{B}_{it_2}}{1 - 2\ \overline{B}_{it_1}} \tag{11}$$

This condition provides a different test for the null hypothesis of unbiased beliefs.

Proposition 1'. *Under Assumptions 1 to 3 and the condition that the matrix of monetary payoffs of player i is symmetric and diagonal-constant (Toeplitz matrix), the hypothesis of equilibrium (unbiased) beliefs implies the testable restriction:*

$$\frac{S_{it_2}}{S_{it_1}} \quad = \quad \frac{1 - 2\ P_{jt_2}}{1 - 2\ P_{jt_1}}\ . \qquad \blacksquare \tag{12}$$

3.2 Identification of utility and beliefs

We now consider the identification of utility parameters $\alpha_\pi(\mathbf{m}_{it})$ and $\beta_\pi(\mathbf{m}_{it})$ and belief parameters $\overline{B}_{it}$ for any treatment t in the set of treatments $\mathcal{T}_{i,t_1}$. Later we discuss the identification of the utility function from the functions $\alpha_\pi(\mathbf{m}_{it})$ and $\beta_\pi(\mathbf{m}_{it})$.

Equations (7) and (9) imply that, for any treatment $t \in \mathcal{T}_{i,t_1}$, preferences and beliefs of player i are identified up to two constants. To see this, define the constant parameters μ and λ as $\mu \equiv \overline{B}_{it_1}$ and $\lambda \equiv \overline{B}_{it_2} - \overline{B}_{it_1}$. And for any treatment $t \in \mathcal{T}_{i,t_1}$, define the ratio $R_{it} \equiv (S_{it} - S_{it_1})/(S_{it_2} - S_{it_1})$ that is identified from the data. Note that by definition $R_{it_1} = 0$ and $R_{it_2} = 1$. Then, we can write equation (9), that describes the model restrictions on beliefs, as:

$$\overline{B}_{it} = \mu + \lambda\ R_{it} \tag{13}$$

Similarly, for any treatment $t \in \mathcal{T}_{i,t_1}$ we can write equation (7) as:

$$S_{it} = \alpha_\pi (\mathbf{m}_{it_1}) + \beta_\pi (\mathbf{m}_{it_1}) \ \ [\mu + \lambda \ R_{it}] \tag{14}$$

Operating in this equation we can obtain the following expressions for the preference parameters in terms of identified objects and the unknown constants μ and λ. For any $t \in \mathcal{T}_{i,t_1}$,

$$\beta_\pi (\mathbf{m}_{it}) = \beta_\pi (\mathbf{m}_{it_1}) \ \ = \ \ \frac{1}{\lambda} \left(S_{it_2} - S_{it_1}\right) \tag{15}$$

$$\alpha_\pi (\mathbf{m}_{it}) = \alpha_\pi (\mathbf{m}_{it_1}) \ \ = \ \ S_{it_1} - \frac{\mu}{\lambda} \left(S_{it_2} - S_{it_1}\right) \tag{16}$$

Equations (13), (15) and (16) show that the vector of parameters $\theta_{t_1} \equiv \{\overline{B}_{it}, \, \alpha_\pi (\mathbf{m}_{it}), \, \beta_\pi (\mathbf{m}_{it}) : t \in \mathcal{T}_{i,t_1}\}$ is identified up to the two constants μ and λ.

The model implies an additional restriction on the sign of $\alpha_\pi (\mathbf{m}_{it_1})$. Remember that $\alpha_\pi (\mathbf{m}_i) \equiv \pi (m_i(1,0)) - \pi (m_i(0,0))$. Since the utility of money is an increasing function, we have that the sign of $\alpha_\pi (\mathbf{m}_{it_1})$ is equal to the sign of the money difference $m_{it_1}(1,0) - m_{it_1}(0,0)$, such that:

$$sign \left\{ S_{it_1} - \frac{\mu}{\lambda} \left(S_{it_2} - S_{it_1}\right) \right\} = sign \left\{ \ m_{it_1}(1,0) - m_{it_1}(0,0) \ \right\} \tag{17}$$

Suppose that $m_{it_1}(1,0) - m_{it_1}(0,0) \geq 0$ and that $S_{it_2} - S_{it_1} > 0$. This is without loss of generality because we can always label the two choice alternatives such that $m_{it_1}(1,0) - m_{it_1}(0,0) \geq 0$, and we can label treatments t_1 and t_2 such that $S_{it_2} - S_{it_1} > 0$. The sign restriction in (17) implies the following inequality constraint for μ/λ:

$$\frac{\overline{B}_{it_1}}{\overline{B}_{it_2} - \overline{B}_{it_1}} \equiv \frac{\mu}{\lambda} \leq \frac{S_{it_1}}{S_{it_2} - S_{it_1}} \tag{18}$$

Note that this inequality also provides a testable restriction for the null hypothesis of unbiased (equilibrium) beliefs: under this null hypothesis, we should have that $P_{jt_1}/(P_{jt_2} - P_{jt_1}) \leq S_{it_1}/(S_{it_2} - S_{it_1})$.

Proposition 2. *Under Assumptions 1 to 3 and monotonicity of the payoff function, the hypothesis of equilibrium (unbiased) beliefs implies the inequality restriction:*

$$\frac{P_{jt_1}}{P_{jt_2} - P_{jt_1}} \ \ \leq \ \ \frac{S_{it_1}}{S_{it_2} - S_{it_1}} \tag{19}$$

Given that the choice probabilities P_{it} and P_{jt} are identified, this restriction is testable as long as the set $\mathcal{T}_{i,t_1}$ contains at least two treatments. ■

Suppose that we have an empirical application where the number of treatments in the set $\mathcal{T}_{i,t_1}$ is greater than two. Suppose that for any treatments t different than t_1 and t_2 we reject the null hypothesis in Proposition 1, but that for treatments t_1 and t_2 we cannot reject the null hypothesis in Proposition 2. Therefore, we cannot reject the null hypothesis that player i has unbiased beliefs at treatments t_1 and t_2 but has biased beliefs at other treatments in the set $\mathcal{T}_{i,t_1}$. Given this condition, the whole vector of structural parameters $\theta_{t_1} \equiv \{\overline{B}_{it},\ \alpha_\pi(\mathbf{m}_{it})\ \beta_\pi(\mathbf{m}_{it}) : t \in \mathcal{T}_{i,t_1}\}$ is point identified.

Proposition 3. *Under Assumptions 1 to 3 and the condition that player i has unbiased beliefs in treatments t_1 and t_2, the vector of structural parameters $\theta_{t_1} \equiv \{\overline{B}_{it},\ \alpha_\pi(\mathbf{m}_{it}),\ \beta_\pi(\mathbf{m}_{it}) : t \in \mathcal{T}_{i,t_1}\}$ is point identified.* ■

Proof: If beliefs at treatments t_1 and t_2 are unbiased, we have that $\mu \equiv \overline{B}_{it_1} = P_{jt_1}$ and $\lambda \equiv \overline{B}_{it_2} - \overline{B}_{it_1} = P_{jt_2} - P_{jt_1}$ such that constants μ and λ are identified. Then, equations (13), (15) and (16) imply that the parameters $\overline{B}_{it}$, $\alpha_\pi(\mathbf{m}_{it})$, and $\beta_\pi(\mathbf{m}_{it})$ are identified for any $t \in \mathcal{T}_{i,t_1}$. ■

Note that the selection of the baseline treatments t_1 and t_2 in the set $\mathcal{T}_{i,t_1}$ should be based on the test in Proposition 2.

When the matrix of monetary payoffs of player i is symmetric and diagonal-constant, we can construct a different version of the inequality test in Proposition 2 and of the identification result of beliefs and payoffs in Proposition 3. Under this structure of the payoff matrix, there is only one unknown constant to determine beliefs and payoff parameters. Taking into account that $S_{it} = \alpha_\pi(\mathbf{m}_{it})\ [1 - 2\ \overline{B}_{it}]$ and $\mu \equiv \overline{B}_{it_1}$, it is straightforward to show that for any treatment $t \in \mathcal{T}_{i,t_1}$,

$$\alpha_\pi(\mathbf{m}_i) \quad = \quad \frac{S_{it_1}}{1 - 2\ \mu} \tag{20}$$

and

$$\overline{B}_{it} = \frac{1}{2}\left[1 - (1 - 2\ \mu)\,\frac{S_{it}}{S_{it_1}}\right] \tag{21}$$

Given these conditions, we have versions of Propositions 2 and 3 for games with a symmetric and diagonal-constant matrix of monetary payoffs.

Proposition 2'. *Under Assumptions 1 and 2, monotonicity of the payoff function, and a Toeplitz matrix of monetary payoffs, the hypothesis of equilibrium (unbiased) beliefs implies the testable inequality restrictions:*

$$\frac{S_{it}}{1 - 2P_{jt}} \quad \geq \quad 0 \tag{22}$$

for any $t \in \mathcal{T}_{i,t_1}$. ■

Proposition 3'. *Under Assumptions 1 to 3, a Toeplitz matrix of monetary payoffs, and the condition that player i has unbiased beliefs in one of the treatments in set T_{i,t_1}, the vector of structural parameters $\theta_{t_1} \equiv \{\overline{B}_{it},\ \alpha_\pi(\mathbf{m}_{it}),\ \beta_\pi(\mathbf{m}_{it}) : t \in \mathcal{T}_{i,t_1}\}$ is point identified.* $\blacksquare$

Proof: Suppose (without loss of generality) that the treatment with unbiased beliefs is t_1. Then, we have that $\mu \equiv \overline{B}_{it_1} = P_{jt_1}$ such that constants μ is identified. Then, equations (20) and (21) imply that the parameters $\overline{B}_{it}$, $\alpha_\pi(\mathbf{m}_{it})$, and $\beta_\pi(\mathbf{m}_{it})$ are identified for any $t \in \mathcal{T}_{i,t_1}$. $\blacksquare$

Based on the identification results in Propositions 1 to 3 (or in Propositions 1' to 3' for games with a Toeplitz matrix of monetary payoffs) we propose the following sequential procedure for the empirical analysis of these games.

Step 1. Test the restriction of Proposition 1 for all treatments. If such test is not rejected, then subjects' observed behavior is consistent with BNE (or QRE) and imposing such equilibrium consistently estimates subjects' utility function. Otherwise, if the test is rejected for some treatments, we proceed to step 2.

Step 2. We apply the test in Proposition 2 to every pair of treatments t_1 and t_2. If we identify a pair of treatments that satisfy the inequality restriction in this test, then we proceed to step 3. Otherwise, we conclude that subjects have unbiased beliefs in at most one treatment and we are not able to identify beliefs and preferences in this case.

Step 3. We impose the restriction of unbiased beliefs in a pair of treatments t_1 and t_2 that pass the test in Proposition 2. Then, we estimate utility functions and beliefs using the restrictions in Proposition 3.

4 Empirical applications

In this section, we apply the model and the identification results in section 2 and 3 to datasets from two laboratory experiments that incorporate the exclusion restriction in Assumption 3. Section 4.1 presents an application to a matching pennies game using experimental data from Goeree and Holt (2001). Section 4.2 deals with a coordination game from Heinemann, Nagel and Ockenfels (2008).

In the two applications we present tests and estimation results under four different parametric specifications for the distribution of the unobserved variable ω_{nit}: (a) standard normal (Probit); (b) standard logistic (Logit); (c) exponential with zero mean; and (d) double exponential with zero mean. The form of the inverse function $S_{it} \equiv F^{-1}_{\omega|\mathbf{m}_{it}}(P_{it})$ for these four distributions is: (a) for the Probit model, $S_{it} = \Phi^{-1}(P_{it})$, where Φ^{-1} is the inverse CDF of the standard normal; (b) for the

Logit model, $S_{it} = \ln(P_{it}) - \ln(1 - P_{it})$; (c) for the exponential model, $S_{it} = -\ln([1 - P_{it}]) - 1$; and (d) for double exponential model, $S_{it} = -\ln(-\ln(P_{it})) - \gamma$ where γ is Euler constant. Note that Probit and Logit model assume unobservable ω_{nit} is symmetrically distributed while it is asymmetric under exponential and double exponential models.

4.1 Matching pennies

Goeree and Holt (2001) conducted an experiment of ten types of games. For each type of game, the find that there exists a particular monetary payoff matrix (i.e. they refer to it as the *treasure* matrix) such that subjects' behaviors are consistent with Nash Equilibrium. However, when monetary payoffs depart from the *treasure* matrix, subjects' behaviors become not consistent with equilibrium. Their empirical results are based on the assumption that monetary payoffs are subjects' true payoff. Our analysis is able to detect whether subjects have unbiased beliefs under a nonparametric specification of the utility function.

4.1.1 Experiment

Table 1 presents the payoff matrices in the experiment by Goeree and Holt (2001, henceforth GH). Each player simultaneously chooses between two possible actions, 0 or 1. The pairs of numbers between brackets, $[m_R, m_C]$, represent the monetary payoffs of row player and the column player, respectively, measured in cents. The experiment contains three games or treatments. The only difference across treatments is in the monetary payoff of the row player under action profile $(a_R, a_C) = (0, 0)$. It is clear that this experimental design satisfies the exclusion restriction in Assumption 3(A). Furthermore, note that the payoff matrix of the column player is symmetric and diagonal-constant. Therefore, for this game we can apply the test and identification result in Propositions 1', 2', and 3'.

The experiment includes 50 subjects ($N = 50$): five cohorts of ten subjects who were undergraduates in an economic class from University of Virginia. They were randomly matched and assigned as row or column player. In addition, the ordering of treatments is alternated for different sessions. Each subject records his/her decision of the game described by table 1 in an instruction sheet. In addition to this matching pennies game, subjects are also asked to play other nine different games which are not the focus of this paper. In this experiment each subject is paid \$6 for showing up. The average earnings for a two-hour session is about \$35 with range from \$15 to \$60 for all 10 games.

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

Treatment 1

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

Treatment 2

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

Treatment 3

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

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

	Player R $[\widehat{P}_{R,t}]$		*Player C* $[\widehat{P}_{C,t}]$	
Treatment 1	0.52	(0.100)	0.52	(0.100)
Treatment 2	0.04	(0.039)	0.84	(0.073)
Treatment 3	0.92	(0.054)	0.20	(0.080)

Note: For player-type i, $\widehat{P}_{it} = \left[\sum_{n=1}^{N} a_n \; 1\{d_n = (i,t)\}\right] / \left[\sum_{n=1}^{N} 1\{d_n = (i,t)\}\right]$

Half of the subjects are randomly selected as row players and the remaining subjects are column players. Each subject plays all three treatments once, and his/her role as either row or column player is fixed across treatments.[11] Table 2 presents the frequencies or players' choice probabilities from this experiment and the corresponding standard errors. The behavior of both players varies across treatments. In particular, though the payoff matrix of the column player is the same in

[11] For detailed instruction of this experiment, visit http://www.people.virginia.edu/~cah2k/trdatatr.pdf.

the three treatments the behavior of this player varies considerably. According to the model, the change in the behavior of the column player should be attributed to the change in this player's beliefs on the behavior of the row player. This evidence is consistent with the "relevance" restriction in Assumption 3(B) that establishes that player C's behavior varies across treatments. We will exploit this source of variation in this experiment to test for unbiased beliefs of the column player and to identify beliefs and utilities for this player. Since the experiment does not provide the same source of variation for the row player, we cannot identify beliefs and preferences for this player.

The monetary payment for player R in outcome $(0,0)$ is higher in treatment 2 compared to treatment 1. Therefore, alternative $a_R = 0$ becomes more attractive to R in treatment 2. If player C has rational beliefs, she would predict that player R will choose $a_R = 0$ with higher probability in treatment 2 than in treatment 1. The best response to such belief is to choose $a_C = 1$ more frequently. A similar argument applies to the comparison of treatments 1 and 3. The estimated choice probabilities in Table 2 are consistent with this argument: P_{C2} $[= 0.84] >$ P_{C1} $[= 0.52] >$ P_{C3} $[= 0.20]$, and these inequalities are statistically significant. However, this argument is not a formal and rigorous test of unbiased beliefs. Without taking into account players' preferences and their degree of risk aversion/loving, we do not know whether or not these changes in the choice probability are consistent with unbiased beliefs. Here we implement formal tests of unbiased beliefs that takes into account these considerations.

4.1.2 Testing procedures and estimation method

Let $\widehat{P}_{Rt}$ and $\widehat{P}_{Ct}$ be the estimated choice probabilities in table 2 for $t = 1, 2, 3$, and let $\widehat{S}_{Ct}$ be $F_{\omega}^{-1}(\widehat{P}_{Ct})$. Based on Proposition 1, we could construct the statistic $\frac{\widehat{S}_{C3}-\widehat{S}_{C1}}{\widehat{S}_{C2}-\widehat{S}_{C1}} - \frac{\widehat{P}_{R3}-\widehat{P}_{R1}}{\widehat{P}_{R2}-\widehat{P}_{R1}}$ and its standard error to test for the null hypothesis of unbiased beliefs using a t-test. This test is asymptotically valid. However, this test does not have good small-sample properties when one, or both, of the values in the denominator, $\widehat{S}_{C2} - \widehat{S}_{C1}$ and $\widehat{P}_{R2} - \widehat{P}_{R1}$, are close to zero. To deal with this issue, we use instead the following statistic:[12]

$$\widehat{\delta} = \left(\widehat{S}_{C3} - \widehat{S}_{C1}\right)\left(\widehat{P}_{R2} - \widehat{P}_{R1}\right) - \left(\widehat{S}_{C2} - \widehat{S}_{C1}\right)\left(\widehat{P}_{R3} - \widehat{P}_{R1}\right) \tag{23}$$

Under the null hypothesis that the column player has unbiased beliefs, $\widehat{\delta}$ is normally distributed with mean zero. We use the bootstrap method to calculate the standard error $se(\widehat{\delta})$ (Horowitz,

[12] Another way to deal with small denominator problems is to introduce a truncation parameter h_N in which the test based on proposition 1 turns $\frac{\widehat{S}_{C3}-\widehat{S}_{C1}}{\max(h_N, \widehat{S}_{C2}-\widehat{S}_{C1})} - \frac{\widehat{P}_{R3}-\widehat{P}_{R1}}{\max(h_N, \widehat{P}_{R2}-\widehat{P}_{R1})}$ and h_N converges to zero as sample size N goes to infinity. However, the empirical result of such approach crucially depends on the choice of h_N. Though a cross-valiation method can be used to choose the value of this parameter, in this paper we have preferred to use a simpler method.

2001). Since in this experiment the matrix of monetary payoffs is Toeplitz, we can also apply the test of unbiased beliefs in Proposition 1'. We can construct the statistics[13]

$$\begin{aligned} \widehat{\delta}_{12} &= \widehat{S}_{C2}\ (1-2\widehat{P}_{R1}) - \widehat{S}_{C1}(1-2\widehat{P}_{R2}) \\ \widehat{\delta}_{13} &= \widehat{S}_{C3}\ (1-2\widehat{P}_{R1}) - \widehat{S}_{C1}\ (1-2\widehat{P}_{R3}) \end{aligned} \tag{24}$$

$\widehat{\delta}_{12}$ is the statistic for the unbiased belief in treatments 1 and 2, and $\widehat{\delta}_{13}$ is the same type of test statistic but for treatments 1 and 3. Define the vector $\hat{\delta}_1 = (\hat{\delta}_{12}, \hat{\delta}_{13})'$. Under the null hypothesis of unbiased beliefs in treatments 1, 2, and 3, the statistic $\hat{\delta}_1' \cdot \widehat{Var}(\hat{\delta}_1) \cdot \hat{\delta}_1$ has a Chi-square distribution with two degrees of freedom where $\widehat{Var}(\hat{\delta}_1)$ is an estimate of the variance-covariance matrix of $\hat{\delta}_1$.

For the estimation of payoffs and beliefs, we can exploit the symmetry of the payoff matrix of the column player to identify payoffs and beliefs parameters with only one restriction of unbiased beliefs (Proposition 3'). Suppose that we impose the restriction of unbiased beliefs in treatment $t = 1$. This implies that we can estimate beliefs of the column player at treatments $t = 2, 3$ using the estimator:

$$\widehat{\widehat{B}}_{Ct} = \frac{1}{2}\left[1 - \left(1 - 2\ \widehat{P}_{R1}\right)\frac{\widehat{S}_{Ct}}{\widehat{S}_{C1}}\right] \tag{25}$$

And we can estimate the payoff parameter of the column player using the estimator:

$$\widehat{\alpha}_{\pi}\left(m_C\right) = \widehat{\pi\left(80\right) - \pi}\left(40\right) \quad = \quad \frac{\hat{S}_{C1}}{1 - 2\hat{P}_{R1}} \tag{26}$$

4.1.3 Empirical results

Table 3 presents results for the tests of unbiased beliefs in the GH experiment. We report results from four different tests: the test from Proposition 1 for $\delta = 0$, and three different tests from Proposition 1', for $\delta_{12} = 0$ and $\delta_{13} = 0$ separately, and for the joint restriction. For each test, we report our tests under four specifications for the distribution of the unobservables (Probit, Logit, Exponential and Double Exponential). Standard errors are calculated by bootstrap with 5,000 bootstrap samples. All the tests are consistent with the hypothesis that the column player has unbiased beliefs in the three treatments. All the p-values are greater than 0.4 and highly insignificant.

Goeree and Holt (2001) conclude that the column player tends to predict correctly the behavior of the row player. In this paper, we verify their qualitative observation in a framework with incomplete information and a nonparametric specification of the utility function. They also conclude

[13] Note that the restrictions $\delta_{12} = 0$ and $\delta_{13} = 0$ imply the restriction $\delta_{23} = 0$, and therefore this third restriction is redundant. Also, note that the restrictions $\delta_{12} = 0$ and $\delta_{13} = 0$ imply δ defined in equation (23) is also zero.

that the row player seems to simply respond to monetary payoff without considering column player's response. Unfortunately, we cannot verify this point using our method because this experiment does not include treatments with different monetary payoffs for the column player. Taking treatment 2 as an example, the row player's high probability of choosing action 0 can be explained by either this player over-predicting the probability that the column player chooses action 0, or by a utility function that values 320 cents far more than 80 cents and 40 cents. Without an exclusion restriction, we cannot distinguish between these two interpretations.

Table 3: Tests of Unbiased Beliefs
Matching Pennies

	Probit	*Logit*	*Exponential*	*Double Exp.*
Treatments 1,2, & 3				
$H_0 : \delta = 0$: $\widehat{\delta}$ [s.e]	0.0503 [0.3753]	0.0726 [0.6479]	-0.1942 [0.3683]	-0.0965 [0.4738]
(p-value)	(0.8818)	(0.9032)	(0.5426)	(0.8226)
Treatments 1 & 2				
$H_0 : \delta_{12} = 0$: $\widehat{\delta}_{12}$ [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)
Treatments 1 & 3				
$H_0 : \delta_{13} = 0$: $\widehat{\delta}_{13}$ [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)
Treatments 1, 2, & 3				
$H_0 : \delta_{12} = \delta_{13} = 0$: Chi-square	0.1392	0.1336	0.6621	0.0633
(p-value)	(0.9328)	(0.9354)	(0.7327)	(0.9688)

Note: Standard error is calculated using bootstrap with 5,000 replications.

Table 4 presents estimates of the preference parameter $\alpha_\pi = \pi(80) - \pi(40)$. We report estimates under the condition of unbiased beliefs at each of the treatments.[14] Except for the exponential model with unbiased belief in treatment 2, all specifications yield significantly positive estimates which implies strict monotonicity of the payoff function. The four different assumptions for the distribution of the unobservables provide very similar estimates of α_π, after adjusting their scales.[15]

[14] In Table 4, we do not include the estimate of $\pi(80) - \pi(40)$ under the restriction of unbiased beliefs in treatment 1. This is because, in treatment 1, both players' choice probabilities are close to 50%, the exclusion restriction has little power, and as a result the estimated preference is very imprecise.

[15] Remember that the standard deviation of the error is 1 for Probit model, $\sqrt{\frac{1}{3}}\pi$ for the Logit model 1 for the exponential model, and $\sqrt{\frac{1}{6}}\pi$ for the double exponential model.

In this experiment, player C receives only two possible monetary payoffs, and therefore, we cannot study possible departures from the restriction of linear utility function.

Table 4: Estimation of Payoff Parameters
Matching Pennies

Parameter	*Probit*	*Logit*	*Exponential*	*Double Exp.*
Unbiased beliefs at $t = 2$				
$\pi(80) - \pi(40)$	1.0809***	1.8024**	0.9050	1.2711*
(s.e.)	(0.3899)	(0.7464)	(0.6418)	(0.6900)
Unbiased beliefs at $t = 3$				
$\pi(80) - \pi(40)$	1.0019**	1.6504*	0.9248***	1.2537***
(s.e.)	(0.4321)	(0.7849)	(0.1940)	(0.3922)

***, **, * indicate significance at 1%, 5%, and 10% levels.

4.2 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 rich variation of player's monetary payoffs such that our method can be applied to these data to relax the restriction of a linear utility function of money. The game in this application has multiple players. In the next subsection, we show that the symmetry of the game implies that we can represent it as a game between one player and the rest of the players, and we can apply our identification results in section 3.

4.2.1 Experiment

Table 3 presents the payoff matrices in the different treatments of the experiment. There are G players in the game. Players simultaneously choose between action 0 and 1. Action $a = 0$ is a safe action that gives the player m_0 Euros regardless of other players' decisions. Action $a = 1$ is a risky action that yields 15 Euros if at least a fraction λ of other players also choose action $a = 1$, but it yields zero monetary payoff otherwise. In this game, m_0 is a measure of the opportunity cost of coordination, and λ measures coordination difficulty. We expect that coordination behavior becomes more unlikely to be maintained as m_0 and λ increase.

Table 5: Coordination Game Experiment (Heinemann, Nagel and Ockenfels, 2009)

		Other players	
	$q =$	fraction of other players choosing $a = 1$	
		$q < \lambda$	$q \geq \lambda$
Player R	$a_R = 0$	m_0	m_0
	$a_R = 1$	0	15 Euros

Treatments

$T = 81$ treatments. Set of treatment consists of every combination (G, m_0, λ) with:

$$\begin{cases} G \equiv \#\ players \in \{4, 7, 10\} \\ m_0 \in \{j * 1.5 \text{ Euros, with } j = 1, 2, ..., 9\} \\ \lambda \in \{1/3, 2/3, 1\} \end{cases}$$

The experiment was conducted in different locations: Barcelona, Bonn, Cologne, and Frankfurt. Heinemann, Nagel and Ockenfels (2009) report that there are substantial differences among subject pools. For instance, subjects in Frankfurt are more risk averse than students from other locations. Therefore, it is reasonable to believe that those subjects from different locations are from different populations. Our analysis focuses on Frankfurt as it contains most of the subjects and treatments.[16]

The experiment was run 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. There are 90 treatments or games according to all the possible values of the parameters G, m_0, and λ with $G \in \{4, 7, 10\}$, $\lambda \in \{1/3,\ 2/3,\ 1\}$, and any value of m_0 between 1.5 Euros and 15 Euros with an incremental unit of 1.5 Euros.[17] Subjects were randomly assigned into a group G, where G is 4, 7 or 10. Then, given the selection of group size G, a subject participates in all the treatments/games for every value of λ and m_0. Therefore, each subject participates in 30 treatments. To prevent learning, Heinemann, Nagel and Ockenfels (2009) do not give feedback between blocks. At the end of a session, only 1 of 40 situations 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.

[16] 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.

[17] We did not use treatments with $m_0 = 15$ in our analysis because subjects' choice probabilities are very imprecisely estimated for these treatments.

We now show that this game can be interpreted as a two-player game. As subjects are randomly and anonymously matched, it is reasonable to assume that player i believes each of the other players chooses alternative 1 with the same probability. Let b_i be such belief probability, and let $BIN(n; N, P)$ be the CDF of a Binomial distribution with parameters (N, P). By definition, we have that $B_i \equiv 1 - BIN(\lambda[G-1], G-1, b_i)$ is player i's belief probability that at least a proportion λ of the other $(G-1)$ players choose alternative 1. Therefore, player i's expected payoff of action a_i can be written as:

$$\Pi_i^e(a_i, B_i) = [1 - B_i] \;\; \pi\left(m_i(a_i, q < \lambda)\right) + B_i \;\pi\left(m_i(a_i, q \geq \lambda)\right) + \varepsilon_i(a_i) \tag{27}$$

As in table 3, $m_i(a_i, q < \lambda)$ represents player i's monetary payoff when she chooses a_i and less than a fraction λ of other players choose alternative 1, and $m_i(a_i, q \geq \lambda)$ is defined similarly. Note that equation (27) has the same structure as equation (2) in section 2.1. The model and the identification results in section 2 and 3 trivially extend to this game. Intuitively, as table 3 shows, we can view this game as player i coordinates with an aggregate player who can choose either $q < \lambda$ or $q \geq \lambda$. To play this game, player i only needs to form a belief about the probability that $q \geq \lambda$, that we denote by B_i. Through this section, we refer to B_i as player i's belief of *successful coordination*.

Unlike the GH experiment, the HNO coordination game does not have a variable that shifts one player's monetary payoff while has no impact on other players' utility. Note that the parameter m_0 shifts all players' monetary payoff and cannot be an exclusion restriction. However, changes in the coordination difficulty parameter λ and group size G play the same role as our exclusion restriction. In particular, a change in λ or G does not shift the payoff matrix of any player but it affects player's belief of successful coordination.[18] With exogenous (randomized) variation in λ and G, all the identification results in previous section hold in HNO coordination game.

The number of subjects N in this experiment is 64, 42, or 40 depending on the treatment. Note that for each value of the safe monetary payoff m_0, there are nine treatments (i.e. 9 combinations of (G, λ)). For illustration purpose, we index these nine treatments by t in HNO experiment.

Figure 1 and table 6 present players' empirical choice probabilities and their corresponding standard errors for each of the 81 treatments. Each graph in figure 1 corresponds to a value of (G, λ) and plots the choice probability of the risky action as a function of m_0, and 95% confidence intervals. Each column in the panel corresponds to a value of λ, and each row to a value of G. For every value of (G, λ), the probability of choosing the risky action declines monotonically with the safe amount of money m_0. For every value of G, the choice probability function shifts

[18] It is easy to see through formula $B_i = 1 - BIN(\lambda[G-1], G-1, b_i)$ as both λ and G affect B_i.

upwards when λ (the *coordination difficulty*) decreases. This effect is statistically significant and illustrates a change in players' behavior due to a change in beliefs. This implies that λ is a *relevant* instrument because it affects players' beliefs without affecting their own payoff matrix, i.e., it satisfies Assumption 3. Comparing graphs across rows in the panel shows that the number of players G has a small and a non-significant effect on players' behavior.

Figure 1: Empirical Choice Probabilities

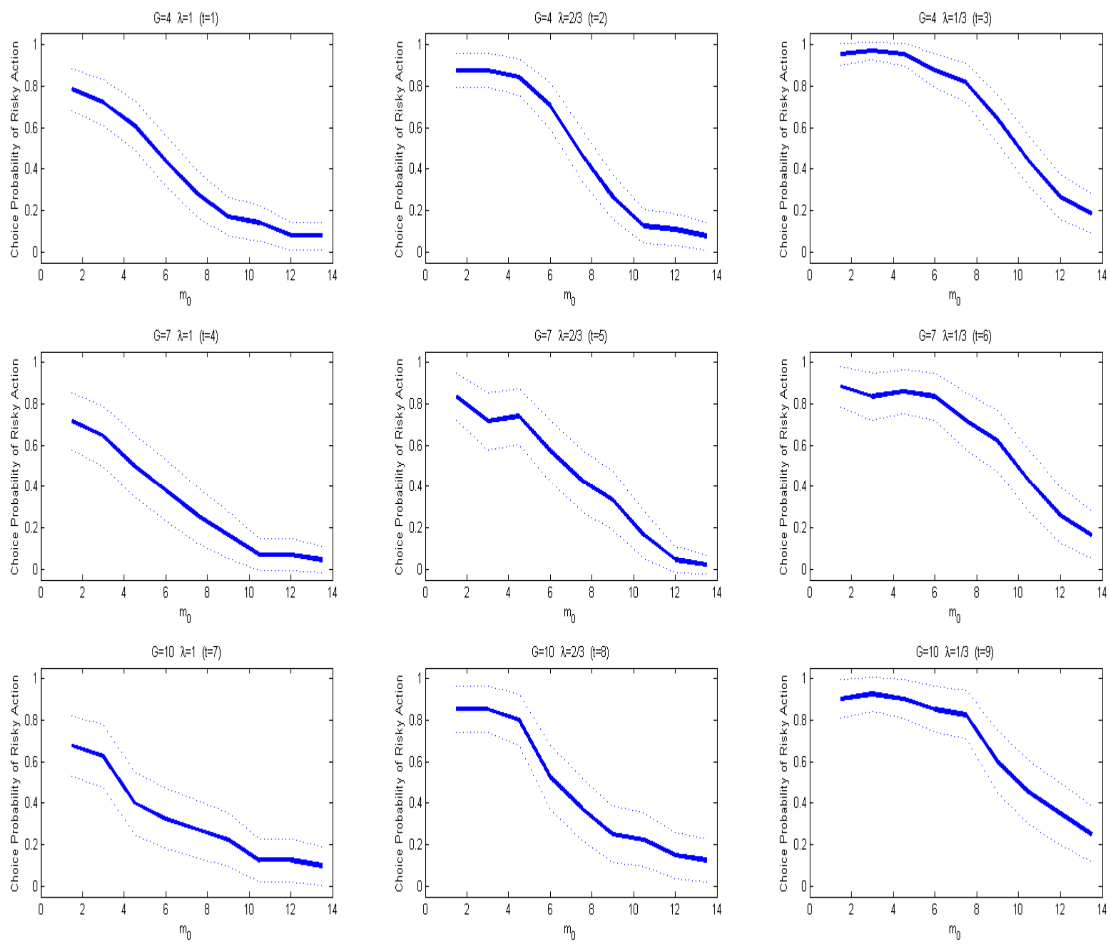

Table 6: Coordination Game Experiment
Empirical Choice Probabilities (Probability of choosing risky action)
(Standard errors in parentheses)

	G = 4			G = 7			G = 10		
	$\lambda=1$	$\lambda=\frac{2}{3}$	$\lambda=\frac{1}{3}$	$\lambda=1$	$\lambda=\frac{2}{3}$	$\lambda=\frac{1}{3}$	$\lambda=1$	$\lambda=\frac{2}{3}$	$\lambda=\frac{1}{3}$
m_0= 1.5	0.7813	0.8750	0.9531	0.7143	0.8333	0.8810	0.6750	0.8500	0.9000
	(0.0517)	(0.0413)	(0.0264)	(0.0697)	(0.0575)	(0.0500)	(0.0741)	(0.0565)	(0.0474)
m_0= 3.0	0.7188	0.8750	0.9688	0.6429	0.7143	0.8333	0.6250	0.8500	0.9250
	(0.0562)	(0.0413)	(0.0217)	(0.0739)	(0.0697)	(0.0575)	(0.0765)	(0.0565)	(0.0416)
m_0= 4.5	0.6094	0.8438	0.9531	0.5000	0.7381	0.8571	0.4000	0.8000	0.9000
	(0.0610)	(0.0454)	(0.0264)	(0.0772)	(0.0678)	(0.0540)	(0.0775)	(0.0632)	(0.0474)
m_0= 6.0	0.4375	0.7031	0.8750	0.3810	0.5714	0.8333	0.3250	0.5250	0.8500
	(0.0620)	(0.0571)	(0.0413)	(0.0749)	(0.0764)	(0.0575)	(0.0741)	(0.0790)	(0.0565)
m_0= 7.5	0.2813	0.4688	0.8125	0.2619	0.4286	0.7143	0.2750	0.3750	0.8250
	(0.0562)	(0.0624)	(0.0488)	(0.0678)	(0.0764)	(0.0697)	(0.0706)	(0.0765)	(0.0601)
m_0= 9.0	0.1719	0.2656	0.6406	0.1667	0.3333	0.6190	0.2250	0.2500	0.6000
	(0.0472)	(0.0552)	(0.0600)	(0.0575)	(0.0727)	(0.0749)	(0.0660)	(0.0685)	(0.0775)
m_0= 10.5	0.1406	0.1250	0.4375	0.0714	0.1667	0.4286	0.1250	0.2250	0.4500
	(0.0435)	(0.0413)	(0.0620)	(0.0397)	(0.0575)	(0.0764)	(0.0523)	(0.0660)	(0.0787)
m_0= 12.0	0.0781	0.1094	0.2656	0.0714	0.0476	0.2619	0.1250	0.1500	0.3500
	(0.0335)	(0.0390)	(0.0552)	(0.0397)	(0.0329)	(0.0678)	(0.0523)	(0.0565)	(0.0754)
m_0= 13.5	0.0781	0.0781	0.1875	0.0476	0.0238	0.1667	0.1000	0.1250	0.2500
	(0.0335)	(0.0335)	(0.0488)	(0.0329)	(0.0235)	(0.0575)	(0.0474)	(0.0523)	(0.0685)
Subjects per treatment		64			42			40	

In this game, $P_{m_0,t}$ represents the choice probability of the risky action when the treatment is (m_0, t) where t denotes a treatment index for (G, λ).[19] Given $G-1$ of the players (all except one), let $g_{m_0,t}$ be the number of these players who choose the risky action. According to the model, $g_{m_0,t}$ is a Binomial random variable with parameters $G-1$ and $P_{m_0,t}$. Therefore, the probability that at least a fraction λ of the other players choose the risky action is:

$$CP_{m_0,t} \equiv \Pr\left(g_{m_0,t} \geq \lambda[G-1]\right) = 1 - BIN\left(\lambda[G-1]\ ; G-1, P_{m_0,t}\right) \tag{28}$$

In this game, as $\overline{B}_{m_o,t}$ represents the average beliefs about the probability that at least a fraction λ of other players choose the risky action, the condition of unbiased beliefs is not $\overline{B}_{m_0,t} = P_{m_0,t}$ but instead $\overline{B}_{m_0,t} = CP_{m_0,t}$.

[19] Column 1 to 9 in table 4 represents $t = 1$ to $t = 9$.

4.2.2 Testing procedures and estimation method

We apply a similar test as for the GH matching pennies game but with some adjustments. For any m_0 and any three treatments, t_1, t_2 and t, we construct the following test-statistic:

$$\begin{aligned}\widehat{\delta}_{m_0,(t_1,t_2,t)} &= \left(\widehat{S}_{m_0,t_2} - \widehat{S}_{m_0,t_1}\right)\left(\widehat{CP}_{m_0,t} - \widehat{CP}_{m_0,t_1}\right) \\ &- \left(\widehat{S}_{m_0,t} - \widehat{S}_{m_0,t_1}\right)\left(\widehat{CP}_{m_0,t_2} - \widehat{CP}_{m_0,t_1}\right)\end{aligned} \tag{29}$$

$\widehat{\delta}_{m_0,(t_1,t_2,t)}$ is the unbiased beliefs test statistic for treatments t_1, t_2 and t. For given (m_0, t_1, t_2), there are seven other combinations of (G, λ) for treatment t, such that we can construct seven different statistics $\widehat{\delta}_{m_0,(t_1,t_2,t)}$. We use these seven $\widehat{\delta}$'s to construct a Chi-Square test for testing the hypothesis that players have unbiased beliefs in all treatments t for a given m_0.

As shown in Proposition 3, to estimate the utility function and the beliefs function we need to impose the restriction that subjects have unbiased beliefs in two treatments. Suppose subjects' beliefs are unbiased in treatments (m_0, t_1) and (m_0, t_2); let (m_0, t) denote another treatment in which subjects' beliefs are allowed to be biased. Following Proposition 3, beliefs in treatment (m_0, t) can be estimated using the expression:

$$\widehat{\widehat{B}}_{m_0,t} = \widehat{CP}_{m_0,t_1} + \left(\widehat{CP}_{m_0,t_2} - \widehat{CP}_{m_0,t_1}\right)\left(\frac{\widehat{S}_{m_0,t} - \widehat{S}_{m_0,t_1}}{\widehat{S}_{m_0,t_2} - \widehat{S}_{m_0,t_1}}\right) \tag{30}$$

Given estimated beliefs $\widehat{\widehat{B}}_{m_0,t}$ for every treatment (m_0, t), we apply OLS to the regression-like equation

$$\hat{S}_{m_0,t} = \frac{\pi(m_0)}{\sigma_{\omega,m_0}} + \frac{\pi(15)}{\sigma_{\omega,m_0}}\,\widehat{\widehat{B}}_{m_0,t} \tag{31}$$

to estimate the utility parameters $\pi(m_0)/\sigma_{\omega,m_0}$ and $\pi(15)/\sigma_{\omega,m_0}$, where σ^2_{ω,m_0} is the variance of the unobservable ω that we allow to be heteroskedastic with respect to m_0. We use a bootstrap re-sampling method to calculate standard errors that account for the two-step feature of the estimation method. Given estimates $\pi(m_0)/\sigma_{\omega,m_0}$ and $\pi(15)/\sigma_{\omega,m_0}$, we obtain the the normalized payoff $\widetilde{\pi}(m_0) = 15 * [\pi(m_0)/\sigma_{\omega,m_0}]/[\pi(15)/\sigma_{\omega,m_0}] = 15 * \pi(m_0)/\pi(15)$ such that $\widetilde{\pi}(m_0)$ does not depend on the variance of the unobservable and $\widetilde{\pi}(15)$ is normalized to 15 and $\widetilde{\pi}(0)$ is normalized to 0.

The estimates of utility and beliefs parameters can be quite sensitive to the choice of treatments t_1 and t_2 we where we impose the restriction of unbiased beliefs. As we explain at the end of section 3, we exploit the unbiased belief tests to guide the choice of treatments t_1 and t_2. More specifically, we apply the following procedure. First, for each value of m_0, we perform the Chi-square unbiased beliefs test described above. If the p-value of the test is high (i.e., $p \geq 0.5$), we assume that subjects have unbiased beliefs in all treatments for that value of m_0 and estimate utility function imposing the restriction of unbiased beliefs. If the p-value of the test is relatively low (i.e., $p < 0.5$), then we look separately at the t-tests of unbiased beliefs for each statistic $\widehat{\delta}_{m_0,(t_1,t_2,t)}$. For given m_0, we select the triple (t_1, t_2, t) with the highest p-value conditional on this p-value being greater than 0.1. We impose the restriction that beliefs are unbiased in these three treatments. For values of m_0 such that there is not any triple (t_1, t_2, t) with p-value greater than 0.1, we apply the inequality test in Proposition 2 and choose the two treatments with the highest p-value conditional on this p-value being greater than 0.1. Finally, we impose the restriction of unbiased beliefs at the selected treatments and estimate the rest of beliefs and utility parameters.

4.2.3 Empirical results

Table 7 presents the results of the unbiased belief tests for all the values of m_0.[20] We first draw some caution on the interpretation of our results for $m_0 = 1.5$ and $m_0 = 13.5$. As the monetary reward for safe action takes its lowest value at $m_0 = 1.5$ (highest value at $m_0 = 13.5$) the choice probability of the risky action remains very high (very low) regardless of the value of λ and G. For these values, choice probabilities have very little variation over the treatments $t = (\lambda,G)$, and this implies that our unbiased beliefs test has little power. In other words, for these extreme values of m_0, the high p-values of the test of unbiased beliefs can be spurious and not reflect evidence of unbiased beliefs but a low power of the test.[21]

Except for $m_0 = 1.5$ and $m_0 = 13.5$, the p-value of the unbiased belief test is increasing in m_0 under the four distributional assumptions for the unobservables.[22] It suggests that subjects tend to evaluate coordination probability incorrectly when coordination cost is relatively low; as such cost increases, subjects tend to have correct belief. In addition, our analysis suggests that the empirical results of unbiased belief tests are robust to distributional assumptions on unobservables as the results are qualitatively and quantitatively similar across different specifications.

[20] Heinemann, Nagel and Ockenfels (2009) also consider a Bayesian game which is different than the one in this paper. The Bayesian Nash Equilibrium in their framework predicts that the equilibrium probability of the risky action increases monotonically with m_0 and λ, and decreases with G. This prediction is clearly rejected by the empirical choice probabilities, and therefore, they conclude that BNE is not appropriate in this experiment. In contrast, the BNE in our framework does not predict their comparative statistics and it requires a formal test of unbiased beliefs.

[21] This conjecture is confirmed by our estimation results of utility function. If we impose the unbiased beliefs restrictions in all the treatments for $m_0 = 1.5$ and $m_0 = 13.5$, the estimated utility function violates monotonicity at those two points.

[22] The only exception happens at $m_0 = 6$.

Table 7: Tests of Unbiased Beliefs
Coordination Game (Chi-Square Test)
(p-value in parentheses)

	Probit Model	*Logit Model*	*Exponential Model*	*Double Exp. Model*
$m_0 = 1.5$	6.2229	5.0582	10.7770	7.2979
	(0.5140)	(0.6529)	(0.1486)	(0.3985)
$m_0 = 3.0$	27.7623***	21.6335***	30.7459***	32.7826***
	(0.0002)	(0.0029)	(0.0001)	(0.0000)
$m_0 = 4.5$	14.4177**	12.1147**	23.5703**	16.5377**
	(0.0442)	(0.0969)	(0.0014)	(0.0206)
$m_0 = 6.0$	7.0733	6.4311	4.1628	6.4730
	(0.4213)	(0.4904)	(0.7608)	(0.4857)
$m_0 = 7.5$	15.9080**	14.4899**	8.0014	11.5969
	(0.0260)	(0.0431)	(0.3325)	(0.1146)
$m_0 = 9.0$	11.5301	10.8009	10.3849	10.5683
	(0.1171)	(0.1475)	(0.1678)	(0.1586)
$m_0 = 10.5$	6.2696	5.3701	4.6734	5.0008
	(0.5086)	(0.6149)	(0.6998)	(0.6599)
$m_0 = 12.0$	3.8985	3.6526	3.5235	3.6314
	(0.7914)	(0.8188)	(0.8327)	(0.8211)
$m_0 = 13.5$	4.5249	4.1119	3.9706	4.0535
	(0.7177)	(0.7668)	(0.7832)	(0.7736)

***, **, * indicate significance at 1%, 5%, and 10% levels.

Table 8 and Figure 2 present the results for the estimation of the utility function. The column titled "t_1, t_2" reports the subset of treatments where we impose the restriction of unbiased beliefs and its corresponding p-value. Recall that we choose unbiased belief treatments through the procedure described in subsection 4.3.2 and normalize $\pi(15) = 15$ to account for heteroskedasticity of unobservable when m_0 varies. We first draw some caution on the reliability of the estimates of $\pi(m)$ for values $m < 4.5$ in the exponential and double exponential models. Though statistically insignificant, the point estimates are negative which implies a violation of the monotonicity of the utility function (i.e. $\pi(m) < \pi(0) = 0$ for $m > 0$). This result could be evidence of misspecification of the exponential and double exponential models. Except for these estimates, all specifications yield reasonable and qualitatively similar result. The utility function is increasing in the monetary payoff with only one violation at point $\pi(13.5)$.[23]

[23] As argued previously, this could be a result that data lacks variation of subjects' behaviors when m_0. Therefore, the unbiased belief tests have little power to provide informative guidance on the choice of unbiased belief treatments.

Table 8: Estimation of Payoff Parameters
Coordination Game
(standard error in parentheses)

	Probit		*Logit*		*Exponential*		*Double Exp.*	
	Payoffs	t_1,t_2	Payoffs	t_1,t_2	Payoffs	t_1,t_2	Payoffs	t_1,t_2
$\pi(1.5)$	0.3803	$1, 3$	1.4103	$1, 3$	-3.4960	$5, 8$	-4.0104	$1, 3$
	(4.7402)	p=n.a.	(4.2894)	p=n.a.	(7.6865)	p= n.a.	(6.9075)	p=n.a.
$\pi(3)$	1.3150	$1, 3$	2.0244*	$1, 3$	-3.2335	$5, 6$	-2.0775	$1, 3$
	(1.0919)	p=n.a.	(1.1733)	p=n.a.	(8.0327)	p= n.a.	(1.4991)	p=n.a.
$\pi(4.5)$	1.6543*	$1, 3, 7$	1.6358**	$1, 3, 7$	-1.4812	$5, 7, 9$	-0.0107	$1, 3, 7$
	(0.9241)	p=0.58	(0.8837)	p=0.78	(4.4394)	p=0.46	(1.4799)	p=0.32
$\pi(6)$	4.1466***	$3, 7, 8$	4.0637***	$3, 7, 8$	-0.3986	*All*	3.3847*	$5, 7, 9$
	(1.0131)	p=0.89	(1.0035)	p=0.96	(1.6953)	p=0.76	(1.7315)	p=0.98
$\pi(7.5)$	7.8417***	$2, 4, 6$	7.7179***	$1, 2, 6$	3.8047*	$2, 7, 9$	5.5005***	$2, 3, 4$
	(1.3819)	p=0.96	(1.2272)	p=0.95	(2.1053)	p=0.80	(1.2735)	p=0.79
$\pi(9)$	10.1793***	$1, 2, 3$	10.2377***	$1, 2, 3$	7.6988***	$1, 2, 3$	8.8530***	$1, 2, 3$
	(1.1710)	p=0.50	(1.1775)	p=0.46	(1.1723)	p=0.39	(1.0671)	p=0.37
$\pi(10.5)$	14.4220***	*All*	14.2903***	*All*	10.5850***	*All*	12.1592***	*All*
	(0.6984)	p=0.50	(0.6190)	p=0.61	(0.3916)	p=0.69	(0.4010)	p=0.65
$\pi(12)$	14.9357***	*All*	14.3011***	*All*	10.6007***	*All*	12.1546***	*All*
	(1.0404)	p=0.79	(1.0667)	p=0.81	(0.5842)	p=0.83	(0.7678)	p=0.82
$\pi(13.5)$	13.4269***	*All*	12.3466***	*All*	9.1063***	*All*	10.4623	*All*
	(1.6128)	p=0.71	(1.6070)	p=0.76	(1.0291)	p=0.78	(1.2633)	p=0.77
$\pi(15)$	15	n.a.	15	n.a.	15	n.a.	15	n.a.
	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.

***, **, * indicate significance at 1%, 5%, and 10% levels.
$t_1$and t_2 are the treatments where we imposed the restriction of unbiased beliefs.

Figure 2 presents the estimated utility function for the Probit specification of the unobservables (blue solid line) together with 95% confidence intervals (blue dashed line) and with a comparison to the linear utility assumption (black solid line).[24] The estimated utility function has an S-shape: it is convex for relatively low values of money and concave for large monetary payoffs. This feature holds also for the other three specifications of the unobservables. This result is significantly different to the standard specification that restricts utility to be equal to the monetary payoff (i.e. $\pi(m) = m$). Importantly, imposing the restriction that utility is equal to the monetary payoff can generate incorrect conclusions on beliefs. Furthermore, our estimates suggest that the specification of a globally concave utility function can generate also biases because this concavity does not hold for small monetary payoffs. This indicates that the conventional functional forms adopted for the utility function in many applications (i.e. linear, logarithmic, CRRA, or CARA functions) can be mis-specified. In contrast, this paper provides a method that estimates payoff functions

[24]This figure ignores the point at $\pi(13.5)$ as monotonicity of the utility function is violated.

without imposing any functional form assumption. Of course, this flexibility has the price of less precise estimates of the parameters. However, the nonparametric approach can be considered as an exploratory procedure to search for the correct parametric and parsimonious specification.

Figure 2: Estimated Payoff (Probit Model)

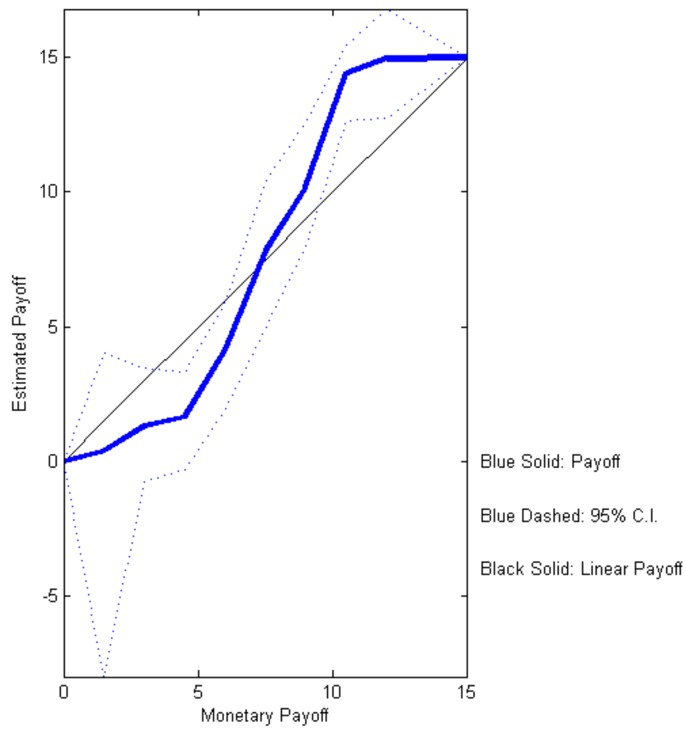

The S-shape of the estimated utility function is consistent with the prediction of Prospect Theory (i.e., Kahneman and Tversky, 1979). According to this interpretation, players of this game would have a reference point for the monetary payoff that they expect to obtain in the game. Let m_R be this reference point. For monetary payoffs above the reference point, preferences show risk aversion (i.e., the utility function is strictly concave), and for monetary payoffs below the reference point, preferences show loss aversion (i.e., the utility function is strictly convex). Therefore, the reference point m_R is the point of inflection where the utility function goes from convex to concave. As our framework non-parametrically specifies the utility function, it also provides an estimate of the reference point rather than imposing it *a priori*. According to our estimates under the Probit and Logit models, the interval estimate of the reference point is [$7.50, $10.50]. This interval includes the mean value of m_0 across all the different experiments of the game, $8.25.[25]

Figure 3: Estimated Beliefs and Choice Probabilities (Probit Model)

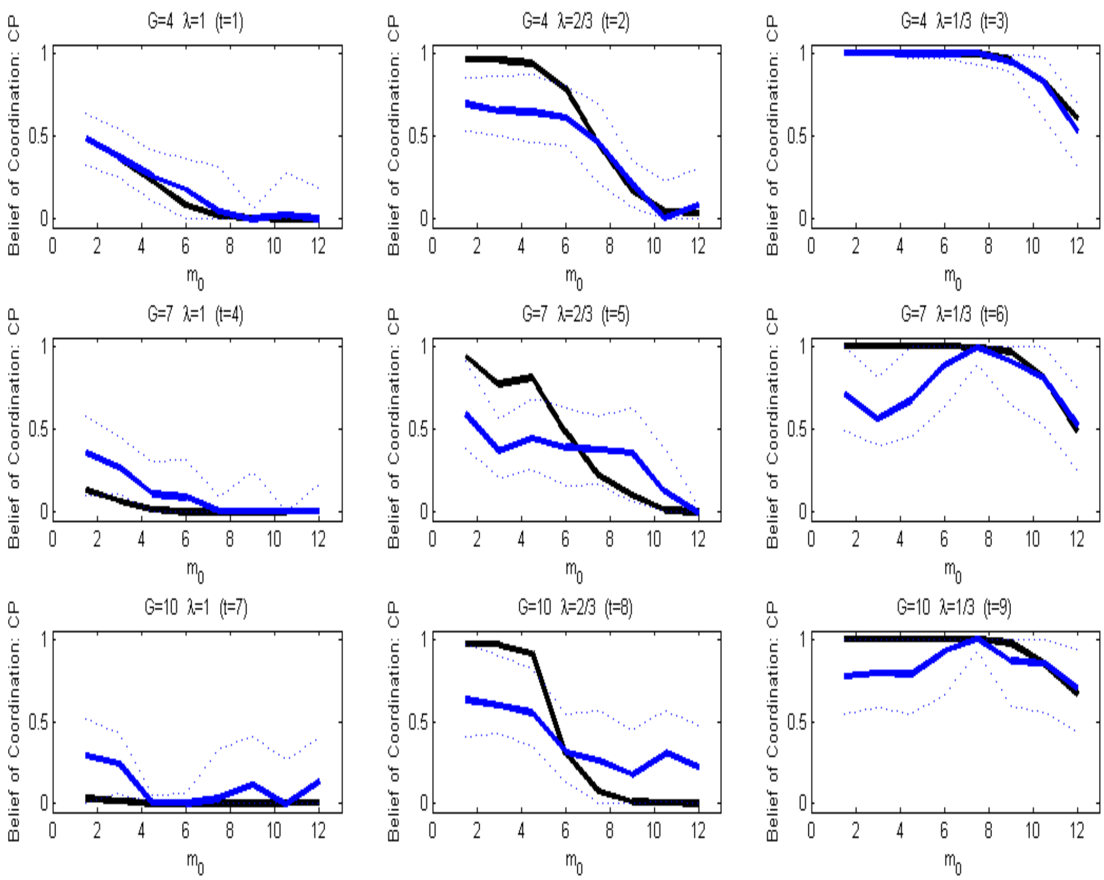

Tables 9 to 11 present our estimates of beliefs B_i, i.e., beliefs on the probability that at least a proportion λ of the other players choose the risky action. Each table corresponds to a particular value for the number of players G. Figure 3 presents a graphical representation of these estimates. In general, subject's belief about successful coordination decreases when either m_0, G, or λ increase. Furthermore, λ has the largest impact on subjects' beliefs. Specifically, when coordination is either easy ($\lambda = \frac{1}{3}$) or difficult ($\lambda = 1$), subjects tend to have unbiased beliefs. In contrast, when coordination difficulty is moderate ($\lambda = \frac{2}{3}$), subjects will significantly underestimate the coordination probabilities when coordination cost m_0 is relatively small and overestimate it when m_0 is large even, though the overestimation is not statistically significant.

[25] Remember that in this experiment every subject plays all the games for every value of G, λ, and m_0.

Table 9: Comparison Between Belief and Actual Coordination Probability
Probit Model, $G = 4$
(standard error in parentheses)

	$\lambda = 1$		$\lambda = \frac{2}{3}$		$\lambda = \frac{1}{3}$	
	Beliefs	True Choice Probability	Beliefs	True Choice Probability	Beliefs	True Choice Probability
$m_0 = 1.5$	0.4768***	0.4768***	0.6943***	0.9570***	0.9999***	0.9999***
s.e.	(0.0956)	(0.0956)	(0.1023)	(0.0283)	(0.0003)	(0.0003)
Equality Test	p-value=1		p-value=0.0090		p-value=1	
$m_0 = 3$	0.3713***	0.3713***	0.6511***	0.9570***	1.0000***	1.0000***
s.e.	(0.0880)	(0.0880)	(0.1131)	(0.0285)	(0.0002)	(0.0002)
Equality Test	p-value=1		p-value=0.0166		p-value=1	
$m_0 = 4.5$	0.2566***	0.2263***	0.6423***	0.9344***	0.9931***	0.9999***
s.e.	(0.0922)	(0.0693)	(0.1255)	(0.0369)	(0.0103)	(0.0004)
Equality Test	p-value=0.5664		p-value=0.0162		p-value=0.5734	
$m_0 = 6$	0.1782	0.0837**	0.6096***	0.7879***	0.9950***	0.9980***
s.e.	(0.1099)	(0.0379)	(0.1092)	(0.0722)	(0.0113)	(0.0026)
Equality Test	p-value=0.3000		p-value=0.0172		p-value=0.8374	
$m_0 = 7.5$	0.0457	0.0222	0.4582***	0.4532***	1.0000***	0.9934***
s.e.	(0.1121)	(0.0149)	(0.1454)	(0.0926)	(0.0336)	(0.006)
Equality Test	p-value=0.8926		p-value=0.9620		p-value=0.1734	
$m_0 = 9$	0.0000	0.0051	0.2124**	0.1742***	0.9468***	0.9536***
s.e.	(0.0236)	(0.0049)	(0.0873)	(0.0658)	(0.0282)	(0.0247)
Equality Test	p-value=0.2902		p-value=0.4926		p-value=0.4944	
$m_0 = 10.5$	0.0246	0.0028	0.0000	0.043	0.8247***	0.8220***
s.e.	(0.0966)	(0.0032)	(0.0788)	(0.0286)	(0.1092)	(0.0602)
Equality Test	p-value=0.8896		p-value=0.3282		p-value=0.9668	
$m_0 = 12$	0.0000	0.0005	0.0793	0.0333	0.5292***	0.6039***
s.e.	(0.0669)	(0.0009)	(0.1050)	(0.0243)	(0.1157)	(0.0899)
Equality Test	p-value=0.4256		p-value=0.7322		p-value=0.0408	
$m_0 = 13.5$	0.0389	0.0005	0.0389	0.0174	0.3594	0.4636
s.e.	(0.0831)	(0.0009)	(0.0837)	(0.0163)	(0.1129)	(0.0963)
Equality Test	p-value=0.7772		p-value=0.8520		p-value=0.0060	

***, **, * indicate significance at 1%, 5%, and 10% levels.

Table 10: Comparison Between Belief and Actual Coordination Probability
Probit Model, $G = 7$
(standard error in parentheses)

	$\lambda = 1$		$\lambda = \frac{2}{3}$		$\lambda = \frac{1}{3}$	
	Beliefs	True Choice Probability	Beliefs	True Choice Probability	Beliefs	True Choice Probability
$m_0 = 1.5$	0.3544**	0.1328	0.5879***	0.9377***	0.7114***	0.9999***
s.e.	(0.1420)	(0.0848)	(0.1576)	(0.0596)	(0.1576)	(0.0009)
Equality Test	p-value=0.0428		p-value=0.0352		p-value=0.0960	
$m_0 = 3$	0.2670**	0.0706	0.3649***	0.7703***	0.5615***	0.9993***
s.e.	(0.1068)	(0.0592)	(0.1096)	(0.1193)	(0.1285)	(0.0027)
Equality Test	p-value=0.0128		p-value=0.0004		p-value=0.0082	
$m_0 = 4.5$	0.1103	0.0156	0.4461***	0.8113***	0.6726***	0.9997***
s.e.	(0.1007)	(0.0206)	(0.1315)	(0.1087)	(0.1536)	(0.0016)
Equality Test	p-value=0.3586		p-value=0.001		p-value=0.0508	
$m_0 = 6$	0.0872	0.0031	0.3889***	0.4852***	0.8807***	0.9993***
s.e.	(0.1083)	(0.0064)	(0.1406)	(0.1470)	(0.1270)	(0.0024)
Equality Test	p-value=0.5518		p-value=0.1586		p-value=0.5166	
$m_0 = 7.5$	0	0.0003	0.3745***	0.2210*	0.9890***	0.9913***
s.e.	(0.0334)	(0.0012)	(0.1227)	(0.1159)	(0.0404)	(0.0143)
Equality Test	p-value=0.5124		p-value=0.062		p-value=0.9586	
$m_0 = 9$	0.0000	0.0000	0.3579**	0.1001	0.9042***	0.9671***
s.e.	(0.0901)	(0.0002)	(0.1651)	(0.0770)	(0.1254)	(0.0362)
Equality Test	p-value=0.4884		p-value=0.0322		p-value=0.6640	
$m_0 = 10.5$	0.0000	0.0000	0.1203	0.0087	0.8049***	0.8085***
s.e.	(0.0193)	(0)	(0.1300)	(0.0165)	(0.1420)	(0.1069)
Equality Test	p-value=0.342		p-value=0.4626		p-value=0.945	
$m_0 = 12$	0.0000	0.0000	0.0000	0.0001	0.5207***	0.4941***
s.e.	(0.0586)	(0.0000)	(0.0225)	(0.0008)	(0.1512)	(0.1495)
Equality Test	p-value=0.3958		p-value=0.3208		p-value=0.2462	
$m_0 = 13.5$	0.0000	0.0000	0.0000	0.0000	0.3109**	0.2632**
s.e.	(0.0451)	(0.0000)	(0.0106)	(0.0002)	(0.1371)	(0.1328)
Equality Test	p-value=0.3452		p-value=0.2612		p-value=0.1320	

***, **, * indicate significance at 1%, 5%, and 10% levels.

Table 11: Comparison Between Belief and Actual Coordination Probability
Probit Model, $G = 10$
(standard error in parentheses)

	$\lambda = 1$		$\lambda = \frac{2}{3}$		$\lambda = \frac{1}{3}$	
	Beliefs	True Choice Probability	Beliefs	True Choice Probability	Beliefs	True Choice Probability
$m_0 = 1.5$	0.2892**	0.0291	0.6280***	0.9661***	0.7706***	1.0000***
s.e.	(0.1442)	(0.0428)	(0.1615)	(0.0547)	(0.1572)	(0.0002)
Equality Test	p-value=0.064		p-value=0.0454		p-value=0.2086	
$m_0 = 3$	0.2437**	0.0146	0.5953***	0.9661***	0.7927***	1.0000***
s.e.	(0.1071)	(0.0270)	(0.1357)	(0.0539)	(0.1411)	(0.0001)
Equality Test	p-value=0.0196		p-value=0.0170		p-value=0.2234	
$m_0 = 4.5$	0.0000	0.0003	0.5536***	0.9144***	0.7853***	1.0000***
s.e.	(0.0204)	(0.0015)	(0.1416)	(0.0899)	(0.1517)	(0.0002)
Equality Test	p-value=0.3812		p-value=0.0102		p-value=0.2388	
$m_0 = 6$	0.0000	0.0000	0.3156**	0.3055*	0.9238***	1.0000***
s.e.	(0.0235)	(0.0004)	(0.1284)	(0.1616)	(0.1205)	(0.0008)
Equality Test	p-value=0.4954		p-value=0.8888		p-value=0.7010	
$m_0 = 7.5$	0.0303	0.0000	0.2603	0.0740	1.0000***	0.9999***
s.e.	(0.1234)	(0.0001)	(0.1765)	(0.0811)	(0.0384)	(0.0017)
Equality Test	p-value=0.8776		p-value=0.1772		p-value=0.4126	
$m_0 = 9$	0.1161	0.0000	0.1764	0.0100	0.8673***	0.9750***
s.e.	(0.1375)	(0.0001)	(0.1492)	(0.0238)	(0.1395)	(0.0436)
Equality Test	p-value=0.533		p-value=0.3092		p-value=0.4606	
$m_0 = 10.5$	0.0000	0.0000	0.3046*	0.0058	0.8522***	0.8505***
s.e.	(0.0929)	(0.0000)	(0.1601)	(0.0174)	(0.1401)	(0.1193)
Equality Test	p-value=0.4998		p-value=0.0488		p-value=0.9712	
$m_0 = 12$	0.1385	0.0000	0.2234	0.0006	0.7086***	0.6627***
s.e.	(0.1335)	(0.0000)	(0.1467)	(0.0046)	(0.1531)	(0.1708)
Equality Test	p-value=0.4134		p-value=0.1618		p-value=0.1458	
$m_0 = 13.5$	0.1212	0.0000	0.2004	0.0002	0.4878***	0.3993**
s.e.	(0.1177)	(0.0000)	(0.1326)	(0.0025)	(0.1429)	(0.1759)
Equality Test	p-value=0.4146		p-value=0.1622		p-value=0.0238	

***, **, * indicate significance at 1%, 5%, and 10% levels.

5 Conclusion

A common approach to study risk aversion and biased beliefs in experimental games is to directly elicit preferences and beliefs. An important concern in the experimental literature is that some elicitation processes may affect players' behavior in games. This paper complements the existing literature by treating utility and beliefs as unknowns and estimating them directly from choice data. Our approach requires an experimental design with multiple treatments where payoff matrices vary across treatments for some players but not others. This revealed preference/beliefs approach avoids the potential biases introduced by the elicitation process. We propose different tests for the null hypothesis of unbiased (equilibrium) beliefs and present identification results on beliefs and payoff function.

We apply our test and identification results to experimental data from a matching pennies game conducted by Goeree and Holt (2001) and a coordination game studied by Heinemann, Nagel, and Ockenfels (2009). Our empirical results show that in the matching pennies game, subjects tend to correctly predict other players' behavior when other players' monetary payoffs change. In the coordination game, the null hypothesis of unbiased belief is rejected when the monetary payoff to safe action is relatively low. When that payoff increases, subjects tend to adjust their beliefs to eliminate the bias. The estimated utility function of money has an S-shape, indicating that subjects are risk loving when receiving small monetary payoffs but they become risk averse when the payoff increases. Finally, our estimates of beliefs suggest that subjects underestimate the probability of coordination when the payoff of the safe action is low.

References

[1] Aguirregabiria, V., and A. Magesan (2015): "Identification and Estimation of Dynamic Games when Players' Beliefs Are not in Equilibrium," Working Paper, University of Toronto.

[2] Anderson, S., Fountain. J, and G. Ruström (2014): "Estimating Subjective Probabilities," *Journal of Risk Uncertainty*, 48, 207-229.

[3] Bajari., P and A. Hortacsu (2005): "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data," *Journal of Political Economy*, 113 (4), 703-741.

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

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

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

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

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

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

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

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

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

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

[14] Güth, W., Schmittberger, R., and B. Schwarze (1982): "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, 3, 367-388.

[15] Haile P., A. Hortacsu, and G. Kosenok (2008): "On the Empirical Content of Quantal Response Equilibrium," *American Economic Review*, 98, 180-200.

[16] 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.

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

[18] Horowitz, J. (2001): "The Bootstrap," In *Handbook of Econometrics*, vol.5, edited by James J. Heckman and Edward E. Leamer. Amsterdam: Elsevier Science.

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

[20] Kahneman, D., J.L. Knetsch, and R.H. Thaler (1986): "Fairness and the Assumptions of Economics," *The Journal of Business*, 59, 285-300.

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

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

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

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

[25] 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.

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

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

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

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

[30] 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.

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

[32] Rabin, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281-1292.

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

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

[35] 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.

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

[37] 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.

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

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