Stochastic Choice and Noisy Beliefs in Games: an Experiment*

Evan Friedman† and Jeremy Ward

November 18, 2019

JOB MARKET PAPER #1

Click here for LATEST VERSION

Abstract

We study an equilibrium model in which players make stochastic choices given their beliefs *and* there is noise in the beliefs themselves. The model primitives are an action-map, which determines a distribution of actions given beliefs, and a belief-map, which determines a distribution of beliefs given opponents' behavior. These are restricted to satisfy axioms that are stochastic generalizations of "best response" and "correct beliefs", respectively. In our laboratory experiment, we collect actions data and elicit beliefs for each game within a family of asymmetric 2-player games. These games have systematically varied payoffs, allowing us to "trace out" both the action- and belief-maps. We find that, while both "noise in actions" and "noise in beliefs" are important in explaining observed behaviors, there are systematic violations of the axioms. In particular, although all subjects observe and play the same games, subjects in different roles have qualitatively different belief biases. To explain this, we argue that the player role itself induces a higher degree of strategic sophistication in the player who faces more asymmetric payoffs. This is confirmed by structural estimates.

Keywords: beliefs; quantal response equilibrium; noisy belief equilibrium
JEL Classification: C72, C92, D84

*We thank Alessandra Casella, Yeon-Koo Che, Mark Dean, Ambuj Dewan, Duarte Goncalves, Navin Kartik, RC Lim, Suanna Oh, Tom Palfrey, Jacopo Perego, and Mike Woodford for helpful comments. This material is based upon work supported by the National Science Foundation under Grant Number 1730399.
†Email: ekf2119@columbia.edu

1 Introduction

Nash equilibrium is the central concept of game theory. It describes a situation of stability in which *(i)* players best respond to their beliefs over opponents' behavior and *(ii)* these beliefs are correct. However, both of these deterministic assumptions are unrealistic in many contexts.

The aim of this paper is to understand the ways in which beliefs and actions deviate from the assumptions of Nash equilibrium. Since the deterministic assumptions of Nash will be trivially rejected,[1] we first characterize as a benchmark model a generalization that allows for both stochastic choice given beliefs and randomness in the beliefs themselves. This model is based on four natural axioms which represent stochastic generalizations of "best response" and "correct beliefs". Next, we collect experimental data–actions and elicited beliefs–in order to test these axioms. While we find evidence for stochasticity in both actions and beliefs, there are systematic violations of the axioms. We show that these failures are qualitatively consistent with non-linearities in the utility function and an effect of the player role itself on subjects' strategic sophistication. This is confirmed by estimates of a unified structural model applied to actions and belief statements jointly.

Existing equilibrium models that incorporate stochastic elements have had success in explaining deviations from Nash. Most notably, quantal response equilibrium (QRE) (McKelvey and Palfrey [1995]), which allows for "noise in actions" while maintaining correct beliefs, has become a standard tool for analyzing experimental data. More recently, Friedman [2019] introduced noisy belief equilibrium (NBE), which is shown to explain several of the same phenomena as QRE by injecting "noise in beliefs" while maintaining best response.[2] Both models, however, make some unrealistic predictions that are directly related to the fact of having noise in only one of actions or beliefs.[3] Despite this, the equilibrium effects of allowing both sources of noise have not been examined.

Our first contribution is to introduce a model that allows for noise in both actions and beliefs, which will serve as our benchmark. The model, which we call QNBE, nests QRE and

[1] Even one failure to best respond or any variance in beliefs is inconsistent with the model's assumptions.

[2] NBE shares much of the same structure as random belief equilibrium (Friedman and Mezzetti [2005]), but differs in that the belief distributions are restricted to satisfy behavioral axioms which gives rise to these predictions.

[3] For instance, Friedman [2019] shows that QRE cannot be invariant to both scaling and translating payoffs, and, in order to explain observed deviations from Nash equilibrium within individual games (e.g. as documented in McKelvey et al. [2000]), QRE implies an oversensitivity to affine transformations. On the other hand, NBE is invariant to affine transformations but implies that non-rationalizable actions are played with probability zero, which is rejected in many datasets.

NBE. It is defined by an *action-map* which determines the mixed actions taken by players given their beliefs, and a *belief-map* that determines the distribution of players' beliefs (i.e. a distribution over opponents' mixed actions) as a function of the opponents' mixed actions. The action-map is restricted to satisfy the axioms of regular QRE (Goeree et al. [2005]), requiring that, for any given belief, higher payoff actions are played with higher probability (*monotonicity*) and that an all-else-equal increase in the payoff to some action increases the probability that action is played (*responsiveness*). The belief-map is restricted to satisfy the axioms of NBE, requiring that belief distributions are unbiased (*unbiasedness*) and shift (in the sense of stochastic dominance) in the same direction as changes in the opponents' behavior (*belief-responsiveness*).

As we illustrate through examples, QNBE does impose testable restrictions in standard actions data, but it is fairly permissive in games for which optimal actions depend on beliefs. Hence, using actions data alone, the test of the model would be weak. Moreover, even if we did find a rejection in actions data, this would not pin down which axiom is violated. To resolve this, we elicit beliefs directly, which allows us to identify both the action- and belief-maps without strong auxiliary assumptions. Using this augmented data, we test the axioms. Now, even if the actions data can be rationalized as QNBE outcomes, we may still reject the model and name the offending axiom.

In our second contribution, we run a laboratory experiment in which subjects choose actions and we directly elicit beliefs for a series of games with systematically varied payoffs. This allows us to observe multiple points on (or "trace out") the empirical action-map and belief-map. Using these maps, we *(i)* test the axioms, *(ii)* offer explanations to the extent that the axioms fail, and *(iii)* quantify the relative importance of action- and belief-noise in explaining features of the data.

Central to our design are the 2×2 asymmetric matching pennies games whose payoffs are in Table 1.[4] Indexed by different values of player 1's payoff parameter $X > 0$, these *X-games* have unique mixed strategy Nash equilibria. By varying X, QNBE predicts variation in actions and beliefs for both players so that we may observe multiple points on the empirical action- and belief-maps. This is important because some of the axioms cannot be falsified otherwise, and violations of axioms may be local to particular regions of the domain.

In addition to the X-games, which are our focus, we also include some dominance solvable games.[5] We use these to derive a subject-level measure of strategic sophistication that helps

[4]Similar games were played in the lab for the first time in Ochs [1995].

[5]We also include a small number of additional games whose data we do not analyze in this paper.

		Player 2	
		L	**R**
Player 1	**U**	0 X	20 0
	D	20 0	0 20

Table 1: *Game X*. Player 1's payoff parameter X controls the asymmetry of payoffs. We assume $X > 0$ which ensures a unique, fully mixed Nash equilibrium.

to rationalize our findings on beliefs.

At the beginning of our experiment, subjects are sorted into player roles (row or column), which they maintain throughout. Subjects state beliefs and take actions for games that appear in random order. These include the X-games for six different values of X. At no point do subjects receive feedback, and each game appears several times so that we may observe multiple elicitations per subject.

In testing the axioms, we find that comparative statics (*responsiveness* and *belief-responsiveness*) hold, but restrictions on levels (*monotonicity* and *unbiasedness*) do not. For the axioms that are rejected, our findings differ across player roles.

Consistent with *responsiveness*, we find that an increase in the expected payoff to an action (through variations in beliefs for a given game) increases the probability the action is played. This is true for both players, all games, and for all regions of expected payoffs.

In testing *monotonicity*, we find systematic failures for player 1 only: for each game, there is an interval of beliefs to which subjects fail to best respond more often than not. These intervals involve beliefs for which the action that has a higher expected payoff is also more likely to result in a zero payoff.

Consistent with *belief-responsiveness*, we find that player i's belief distributions tend to be ordered by stochastic dominance across games in the same direction as differences in player j's action frequencies. Beliefs tend to overreact in the sense that small differences in action frequencies are associated with large differences in average beliefs, but this is entirely consistent with the axiom.

In testing *unbiasedness*, we find that player 1 is marginally biased, tending to form slightly conservative beliefs that are closer to the uniform distribution than player 2's actual frequency of play. Player 2, on the other hand, forms very biased, extreme beliefs: whereas player 1's behavior is relatively close to uniform across all X-games, player 2 tends to think that player 1 will overwhelmingly choose U when X is large and similarly choose D when

X is small. Whereas conservative beliefs have been found in games played without feedback (e.g. Huck and Weizsacker [2002]), we believe this asymmetric pattern of bias is novel.

This gives two puzzles with respect to the benchmark QNBE model. These are a failure of *monotonicity* for player 1-subjects who fail to best respond more often than not given some intervals of beliefs and a failure of *unbiasedness*, with the nature of bias depending on player role. We provide explanations and a fitted model that can capture these features of the data.

To explain the failure of *monotonicity* for player 1, we show that, given stated beliefs, concavity in the utility function over payoffs qualitatively predicts precisely the violations we observe (payoffs are in probability points of earning a prize, so this is distinct from risk aversion). This is backed by structural estimates, which suggest that most subjects individually have concavity and that a reasonable calibration can accommodate most of the violations.[6]

To explain the failure of *unbiasedness*, our first clue is that the belief-bias is qualitatively different for the two players. This leads us to conjecture that one player, by merit of her role in the game, is induced to think about her opponent more deeply or with greater "strategic sophistication". This could generate the bias as player 2 believes that player 1 tends to take the low-level action (the best response to random behavior: U when X is large, D when X is small) whereas player 1 anticipates this and acts accordingly.

This sophistication hypothesis cannot be tested within the X-games directly, but can be studied with the help of the dominance solvable games. All subjects face the same action and belief choices in these games, so we can use the belief statements to derive a subject-level measure of strategic sophistication that is collected identically for all subjects. We formally justify this measure through a level k-type framework.[7] Using this measure, we find that player 1-subjects have much higher levels of sophistication than player 2-subjects. All subjects see exactly the same games throughout the experiment, were randomly assigned to their roles, and played a number of X-games before playing a dominance solvable game. Hence, we conclude that experience in the player 1-role of the X-games causally induces greater sophistication and this somehow spills over to the dominance solvable games.

Based on these sophistication results, we consider generalized level k-type models to rationalize the beliefs data. The simplest model that explains the large majority of individual

[6]Previous studies eliciting beliefs in games (Costa-Gomes and Weizsacker [2008] and Rey-Biel [2009]) find little evidence of risk aversion though there are exceptions (Ivanov [2011]).

[7]The level k literature was started by Nagel [1995] and Stahl and Wilson [1995], and is reviewed in Crawford et al. [2013].

subjects' belief patterns is a parametric, subjective cognitive hierarchy model (Camerer [2003]) that embeds payoff sensitivity (as in QRE) into players' models of others. We fit the model to individual subjects' belief data from the entire set of X-games. We interpret one of the fitted parameters as strategic sophistication, and find that this model also captures the sophistication gap between the players. The model predicts that this inferred sophistication should be correlated at the subject level with the measure derived from the dominance solvable games. We find very strong correlations, which we take as further evidence that player 1-subjects' beliefs in the X-games indicate higher levels of sophistication.

Next, we conduct a counterfactual exercise to determine the relative importance of action-noise and belief-noise for explaining the data. Specifically, we consider *(i)* the action frequencies that we *would have* observed if subjects best responded to all of their stated beliefs and *(ii)* the action frequencies we *would have* observed if subjects had beliefs equal to their opponents' empirical action frequencies (and their actions were determined by a best-fit random utility model). These correspond to "turning off" action-noise and belief-noise, respectively. Both counterfactuals deviate considerably from the empirical action frequencies, indicating that both sources of noise are important. Comparing the performance of the counterfactuals, we find that the latter is more accurate for player 1 (i.e. action-noise is more important) and the former is more accurate for player 2 (i.e. belief-noise is more important). Hence, ignoring any one source of noise may lead to misspecification, and which source of noise is more important depends on the context.

Our analysis throughout the paper implicitly assumes that stated beliefs equal the underlying "true" beliefs that subjects hold in their minds and guide their actions. This is, of course, a hypothesis that cannot be directly tested (see Rutstrom and Wilcox [2009] for a discussion). Since we take "noise in beliefs" seriously, we consider the possibility that stated beliefs are simply noisy signals of true beliefs. Assuming this were the case, can we reject the axioms with respect to true beliefs? Could we say that true beliefs are noisy *at all*? We show that, under mild assumptions, the answer to both questions is yes.

This paper contributes to a large literature on behavioral game theory (Camerer et al. [2004a]) that focuses on bounded rationality. More narrowly, we contribute to the theory and empirical study of equilibrium models that inject stochasticity into actions and beliefs (especially McKelvey and Palfrey [1995], Chen et al. [1997], Friedman and Mezzetti [2005], and Friedman [2019]). Our central tool is belief elicitation, so we engage with the growing methodological literature on belief elicitation (see Schotter and Trevino [2014] and Schlag et al. [2015] for review articles) and benchmark our findings against those from well-known

studies that elicited beliefs (e.g. Nyarko and Schotter [2002], Costa-Gomes and Weizsacker [2008] and Rey-Biel [2009]). Our key innovation is to collect multiple elicitations per subject without feedback for each game within a family of closely related games. This allows us to study noise in beliefs and examine how beliefs vary across games. We refer to Section 9 for a detailed discussion of the literature.

The paper is organized as follows. Section 2 presents the theory, Section 3 gives the experimental design, and Section 4 provides an overview of the data. Section 5 presents the results from testing the axioms, Section 6 offers explanations for the axioms' failure, and Section 7 studies the relative importance of action- and belief-noise. Section 8 discusses issues of belief elicitation and the interpretation of stated beliefs–and how these may affect the interpretation of our results. Section 9 discusses the relationship of this paper to the existing literature, and Section 10 concludes.

2 Theory

The deterministic assumptions of "best response" and "correct beliefs" implicit in Nash equilibrium will be trivially rejected, so we introduce a new benchmark model that replaces these deterministic assumptions with stochastic generalizations.

The model we study is a hybrid, defined by an action-map satisfying the axioms of regular QRE (Goeree et al. [2005]) and a belief-map satisfying the axioms of NBE (Friedman [2019]). Anticipating the experiment, we present the case of 2×2 games in which there are two players with two actions each, but as QRE and NBE are defined very generally, the model generalizes to all finite, normal form games.

A game is defined by $\Gamma^{2\times2} = \{N, A, u\}$ where $N = \{1, 2\}$ is the set of players, $A = A_1 \times A_2 = \{U, D\} \times \{L, R\}$ is the action space, and $u = (u_1, u_2)$ is a vector of payoff functions with $u_i : A \to \mathbb{R}$. In other words, this is any game in which player 1 can move up (U) or down (D) and player 2 can move left (L) or right (R).

We use i to refer to a player and j for her opponent. We reserve k and l for action indices. Since each player has only two actions, we write player i's mixed action as $\sigma_i \in [0, 1]$. In an abuse of notation, we use $\sigma_1 = \sigma_U$ and $\sigma_2 = \sigma_L$ to indicate the probabilities with which player 1 takes U and player 2 takes L, respectively.

2.1 Action-map

Let $\sigma_j^{'} \in [0,1]$ be an arbitrary belief that player i holds over player j's action. Given this belief, player i's vector of expected payoffs is $\bar{u}_i(\sigma_j^{'}) = (\bar{u}_{i1}(\sigma_j^{'}), \bar{u}_{i2}(\sigma_j^{'})) \in \mathbb{R}^2$, where $\bar{u}_{ik}(\sigma_j^{'})$ is the expected payoff to action k. We use $v_i = (v_{i1}, v_{i2}) \in \mathbb{R}^2$ as shorthand for an arbitrary such vector. That is, v_i is understood to satisfy $v_i = \bar{u}_i(\sigma_j^{'})$ for some $\sigma_j^{'}$.

As in QRE, the action-map is induced by a quantal response function $Q_i : \mathbb{R}^2 \to [0,1]$. This maps any vector of expected payoffs (given beliefs) to a mixed action, and it is assumed to satisfy the following regularity axioms (Goeree et al. [2005]):

(A1) **Interiority**: $Q_{ik}(v_i) \in (0,1)$ for all $k \in 1,2$ and for all $v_i \in \mathbb{R}^2$.

(A2) **Continuity**: $Q_{ik}(v_i)$ is a continuous and differentiable function for all $v_i \in \mathbb{R}^2$.

(A3) **Responsiveness**: $\frac{\partial Q_{ik}(v_i)}{\partial v_{ik}} > 0$ for all $k \in 1,2$ and $v_i \in \mathbb{R}^{J(i)}$.

(A4) **Monotonicity**: $v_{ik} > v_{il} \implies Q_{ik}(v_i) > Q_{il}(v_i)$ and $v_{ik} = v_{il} \implies Q_{ik}(v_i) = \frac{1}{2}$.

(A1) and (A2) are non-falsifiable technical axioms. Taken together, (A3) and (A4) are a stochastic generalization of "best response", requiring than an all-else-equal increase in the payoff to an action increases the probability it is played and that, given any belief, the best response is taken more often than not.[8]

2.2 Belief-map

Player i's belief about j's mixed action is drawn from a distribution that depends on j's mixed action. In other words, player i's beliefs are a random variable $\sigma_j^*(\sigma_j)$ whose distribution depends on σ_j and is supported on $[0,1]$. This family of random variables, or belief-map, is described by a family of CDFs: for any potential belief $\bar{\sigma}_j \in [0,1]$, $F_i(\bar{\sigma}_j|\sigma_j)$ is the probability of realizing a belief less than or equal to $\bar{\sigma}_j$ given that player j is playing σ_j. Following Friedman [2019], the belief-map is assumed to satisfy the following axioms:

(B1) **Interior full support**: For any $\sigma_j \in (0,1)$, $F_i(\bar{\sigma}_j|\sigma_j)$ is strictly increasing and continuous in $\bar{\sigma}_j \in [0,1]$.

(B2) **Continuity**: For any $\bar{\sigma}_j \in (0,1)$, $F_i(\bar{\sigma}_j|\sigma_j)$ is continuous in $\sigma_j \in [0,1]$.

(B3) **Belief-responsiveness**: For all $\sigma_j < \sigma_j^{'} \in [0,1]$, $F_i(\bar{\sigma}_j|\sigma_j^{'}) < F_i(\bar{\sigma}_j|\sigma_j)$ for $\bar{\sigma}_j \in (0,1)$.

[8] Requiring that $v_{ik} = v_{il} \implies Q_{ik}(v_i) = \frac{1}{2}$ in (A4) is unnecessary since it is implied by $v_{ik} > v_{il} \implies Q_{ik}(v_i) > Q_{il}(v_i)$ and (A1). We added this condition to (A4) in order to have a clean division between technical and behavioral axioms.

(B4) **Unbiasedness**: $F_i(\sigma_j|\sigma_j) = \frac{1}{2}$ for $\sigma_j \in (0,1)$. $\sigma_j^*(0) = 0$ and $\sigma_j^*(1) = 1$ with prob. 1.

(B1) and (B2) are non-falsifiable technical axioms. (B1) requires that belief distributions have full support and no atoms when the opponent's action is interior, and (B2) requires that the belief distributions vary continuously in the opponent's behavior except possibly as the opponent plays a pure action with a probability that approaches one. Taken together, (B3) and (B4) are a stochastic generalization of "correct beliefs". (B3) requires that, when the opponent's action increases, beliefs shift up in a strict sense of stochastic dominance.[9] (B4) imposes that belief distributions are correct *on median*. Both median- and mean-unbiasedness can be microfounded via a model of sampling (Friedman [2019]). The technical axioms allow for either or both types of unbiasedness. We use median-unbiasedness to derive theoretical results because it turns out to be much simpler in our setting, but we test for both types of unbiasedness in our data.

2.3 Equilibrium

In equilibrium, player i quantal responds to belief realizations where the beliefs are drawn from a distribution that depends on j's mixed action–and j's mixed action is her expected quantal response similarly induced by quantal responding to belief realizations.

Given player j's mixed action $\sigma_j \in [0,1]$, player i's beliefs are drawn according to $F_i(\cdot|\sigma_j)$. For each belief realization $\sigma_j^{'} \in [0,1]$, player i's mixed action is given by quantal response to expected payoffs $Q_i(\bar{u}_i(\sigma_j^{'})) \in [0,1]$. Player i's expected quantal response as a function of σ_j, which we call the *reaction function*, simply integrates over belief realizations: $\Psi_i(\sigma_j; Q_i, \sigma_j^*) \equiv \int_{[0,1]} Q_i(\bar{u}_i(\sigma_j^{'}))dF_i(\sigma_j^{'}|\sigma_j) \in [0,1]$. Since $Q_i : \mathbb{R}^2 \to [0,1]$ is single-valued, Ψ_i is also single-valued, i.e. a function as opposed to a correspondence.

A given profile of quantal response functions $Q = (Q_1, Q_2)$ and belief-maps $\sigma^* = (\sigma_1^*, \sigma_2^*)$ induce the reaction function $\Psi = (\Psi_1, \Psi_2) : [0,1]^2 \to [0,1]^2$. Equilibrium is defined as a mixed action profile that is a fixed point along with the supporting belief distributions.

Definition 1. Fix $\{\Gamma^{2\times 2}, Q, \sigma^*\}$. A quantal response-noisy belief equilibrium (QNBE) is a pair $\{\sigma, \sigma^*(\sigma)\}$ where $\sigma = \Psi(\sigma; Q, \sigma^*)$ is a mixed action profile and $\sigma^*(\sigma)$ is the supporting profile of belief distributions.

The mapping Ψ is continuous, so existence follows from Brouwer's fixed point theorem.

[9] This is stronger than standard stochastic dominance, which helps with comparative statics and in establishing uniqueness of equilibria, but the distributions can still be arbitrarily close, so it is only slightly stronger.

Proposition 1. *Fix* $\{\Gamma^{2\times 2}, Q, \sigma^*\}$. *A QNBE exists.*

Proof. See Appendix 11.2. □

2.4 QRE and NBE

QRE is defined as in QNBE except that beliefs are correct with probability one.

Definition 2. Fix $\{\Gamma^{2\times 2}, Q\}$. A quantal response equilibrium (QRE) is any mixed action profile σ such that $\sigma = Q(\bar{u}(\sigma))$.

Similarly, NBE is defined as in QNBE except that players best respond to all belief realizations.

Definition 3. Fix $\{\Gamma^{2\times 2}, \sigma^*\}$. A noisy belief equilibrium (NBE) is a pair $\{\sigma, \sigma^*(\sigma)\}$ where $\sigma \in \psi(\sigma; \sigma^*)$ and $\psi_i(\sigma_j; \sigma_j^*) \equiv \int_{[0,1]} BR_i(\bar{u}_i(\sigma_j^{'}))dF_i(\sigma_j^{'}|\sigma_j)$ defines the expected best response correspondence.[10]

In other words, the QRE belief-map is the identity map and the NBE action-map is the best response correspondence. For almost every game, the sets of attainable QRE and NBE mixed action profiles–that can be supported for some primitives–are nested in the set of attainable QNBE mixed action profiles.[11] For the games we analyze in this paper, we show this directly.

2.5 X-games

We specialize theory for the family of X-games whose payoffs are in Table 1. This serves to illustrate the QNBE model and provides our justification for using the X-games in the experiment.[12]

[10] BR_i is the standard best response correspondence: $BR_i(v_i) = 1$ if $v_{i1} > v_{i2}$, $BR_i(v_i) = 0$ if $v_{i1} < v_{i2}$ and $BR_i(v_i) = [0, 1]$ if $v_{i1} = v_{i2}$. ψ_i is the *expected* best response correspondence, where the expectation is over belief realizations whose distribution depends on the opponent's behavior. Friedman [2019] shows that ψ_i single-valued in generic games since the probability of indifference is zero by (B1).

[11] *A non-generic counterexample.* If players are indifferent between all of their actions, independent of their opponents' behavior, then any distribution of actions is an NBE for any belief-map by (B4). On the other hand, QNBE and QRE predict that all players uniformly mix, and this is true for all primitives by (A4).

[12] The results characterizing behavior within a game generalize to all 2×2 games with unique, mixed strategy Nash equilibria. Comparative static results generalize to any such game with respect to changes in any one payoff parameter.

The X-games have unique, mixed strategy Nash equilibria (NE). As is well-known, NE predicts each player must mix to make the other player indifferent, and so $\sigma_L^{NE,X} = \frac{20}{20+X}$ and $\sigma_U^{NE,X} = 0.5$ (i.e. constant for all X). Since we are only working within the X-game family and $\sigma_L^{NE,X}$ is a strictly decreasing function of X, we think of $\sigma_L^{NE,X}$ as a parameter of the game, and we freely go between X and $\sigma_L^{NE,X}$ as convenient.

First, we establish that, for any fixed primitives, the QNBE is unique.

Proposition 2. *Fix $\{X, Q, \sigma^*\}$. There is a unique QNBE.*

Proof. See Appendix 11.2. □

There is a unique QNBE for any fixed primitives, but since the primitives are only restricted to satisfy axioms, we characterize the set of equilibria that can be attained for *some* primitives. The next result characterizes the reaction functions consistent with the axioms and thus the set of mixed action profiles that can be supported as QNBE outcomes. The proof is by construction, and hence implicitly gives the equilibrium belief distributions as well, though we abstract from that here.

Proposition 3. *Fix X. (i) Any reaction function $\Psi_U : [0,1] \to [0,1]$ that is continuous, strictly increasing, and satisfying the restrictions of (1) can be induced for some primitives $\{Q_U, \sigma_L^*\}$. (ii) Any reaction function $\Psi_L : [0,1] \to [0,1]$ that is continuous, strictly decreasing, and satisfying the restrictions of (2) can be induced for some primitives $\{Q_L, \sigma_U^*\}$. (iii) Any $\sigma = (\sigma_U, \sigma_L)$ satisfying $\sigma_U \in \Phi_U^X(\sigma_L)$ and $\sigma_L \in \Phi_L^X(\sigma_U)$ can be supported as QNBE outcomes for some primitives $\{Q, \sigma^*\}$.*

$$\Phi_U^X(\sigma_L) \in \begin{cases} (0, 3/4) & \sigma_L < \sigma_L^{NE,X} \\ (1/4, 3/4) & \sigma_L = \sigma_L^{NE,X} \\ (1/4, 1) & \sigma_L > \sigma_L^{NE,X} \end{cases} \tag{1}$$

$$\Phi_L^X(\sigma_U) \in \begin{cases} (1/4, 1) & \sigma_U < \frac{1}{2} \\ (1/4, 3/4) & \sigma_U = \frac{1}{2} \\ (0, 3/4) & \sigma_U > \frac{1}{2} \end{cases} \tag{2}$$

Proof. See Appendix 11.2. □

Figure 1 illustrates the proposition for $X = 80$, in which case $\sigma_L^{NE,X} = 1/5$. Here, we plot equilibrium mixed actions in the unit square of $\sigma_L - \sigma_U$ space. The first panel plots

$\Phi_U^X(\sigma_L)$ (1) and the second panel plots $\Phi_L^X(\sigma_U)$ (2). Where these two regions intersect (third panel) is the set of QNBE mixed action profiles that can be attained for some $\{Q, \sigma^*\}$ (part *(iii)* of the proposition).

As shown in Figure 1, the set of attainable QNBE mixed action profiles can be rather large. For $X = 80$, the Lebesgue measure is 51.25%, meaning just over half of all possible mixed action profiles can be supported as QNBE outcomes. However, QNBE makes predictions over actions and beliefs, so even if the actions data falls in this region, the axioms–and thus the model–may be falsified.

Friedman [2019] showed that for any 2×2 game with a unique, fully mixed NE–and hence for any X-game also–the sets of attainable QRE and NBE mixed action profiles coincide. In Figure 1, we plot this set as a cross-hatched rectangle, which has a measure of 15% (see Goeree et al. [2005] and Friedman [2019] for the derivation of such sets in similar games). Hence, allowing for just one of action-noise or belief-noise leads to the same measure of outcomes, but allowing for both increases the set of outcomes more than 3-fold.

Our next result is a comparative static.

Proposition 4. *Fix $\{Q, \sigma^*\}$. σ_U^{QNBE} is strictly decreasing in $\sigma_L^{NE,X}$ and σ_L^{QNBE} is strictly increasing in $\sigma_L^{NE,X}$.*

Proof. See Appendix 11.2. □

The proposition says that, under QNBE, varying X will cause systematic variation in mixed actions, and thus belief distributions also, for both players. This comparative static holds for QRE and NBE also (Goeree et al. [2005] and Friedman [2019]) and is essential in order to "trace out" the empirical action- and belief-maps.

We extend our results to a characterization of QNBE for any finite number of games. To this end, let $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ be a *dataset of mixed actions* from an arbitrary (finite) set of X-games. It is immediate that in order to support the dataset as QNBE outcomes for some primitives (held fixed across games), it is necessary for the restrictions of Proposition 3 to hold for each X and the comparative static of Proposition 4 to hold across any pair of Xs. As it turns out, this is also sufficient.

Proposition 5. *Let $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ be a dataset of mixed actions for any finite number of X-games. The data can be supported as QNBE outcomes for some primitives $\{\sigma^*, Q\}$ (held fixed across games) if and only if*

(i) $\hat{\sigma}_U^X \in \Phi_U^X(\hat{\sigma}_L^X)$ for all X, where Φ_U^X is defined as in (1),

(ii) $\hat{\sigma}_L^X \in \Phi_L^X(\hat{\sigma}_U^X)$ for all X, where Φ_L^X is defined as in (2),

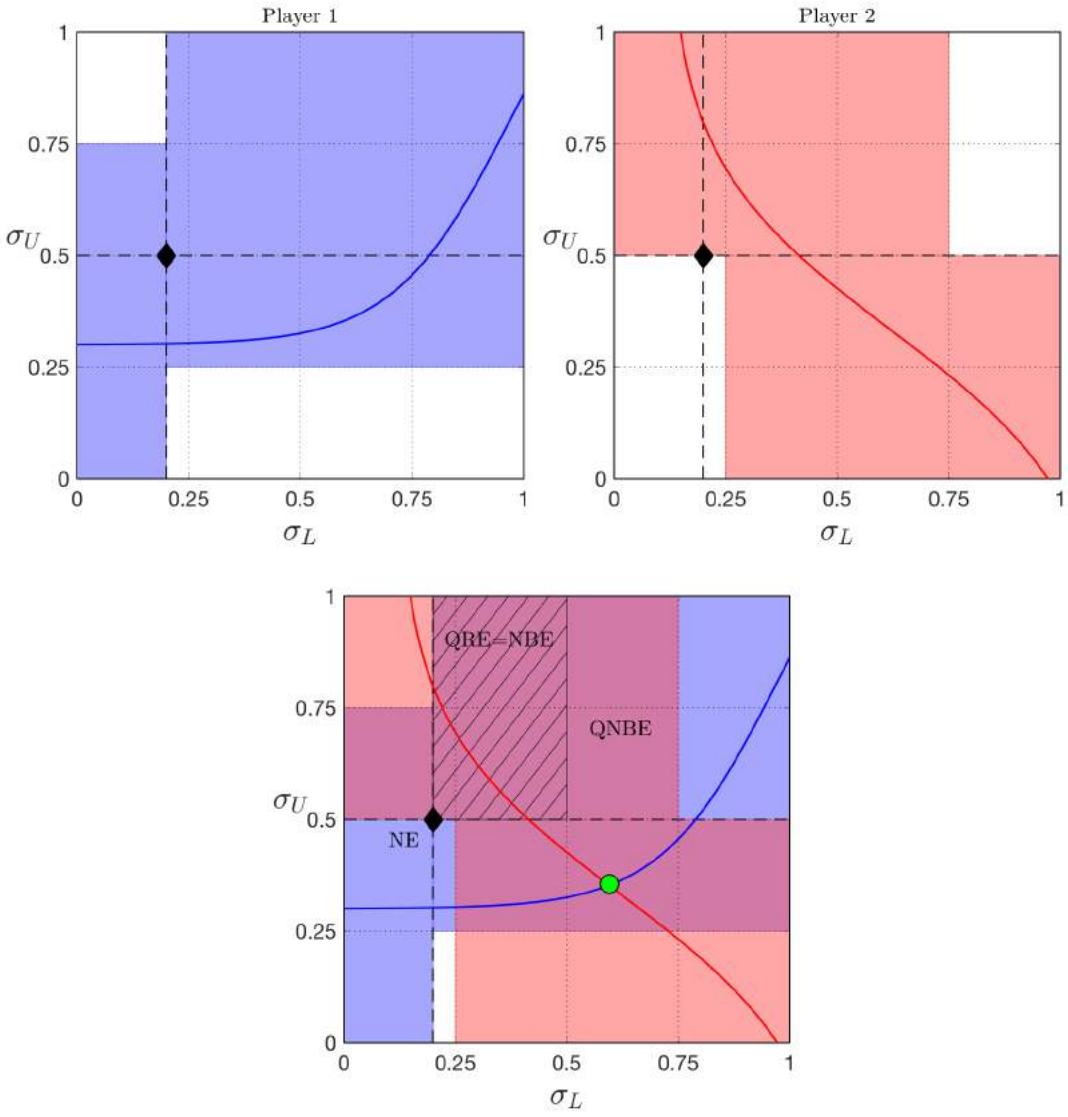

Figure 1: *QNBE in game* $X = 80$. The first panel gives the region in which player 1's QNBE reaction must lie, with an example drawn in blue. The second panel gives the region in which player 2's QNBE reaction must lie, with an example drawn in red. The third panel plots the intersection of the two regions which gives the set of QNBE mixed action profiles that can be attained for some primitives. The black diamond is the Nash equilibrium, the cross-hatched rectangle gives the sets of attainable QRE and NBE, which coincide, and the green dot is an example QNBE mixed action profile.

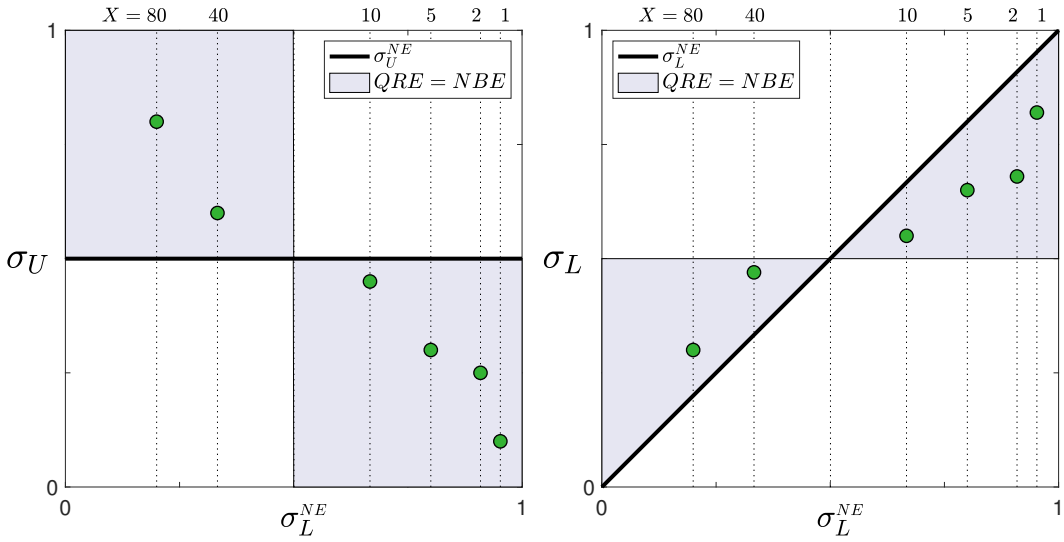

Figure 2: *QRE and NBE in the X-games as a function of σ_L^{NE}.* This figure plots a hypothetical dataset $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ of mixed actions for the set of X-games with $X \in \{80, 40, 10, 5, 2, 1\}$. The left panel plots $\hat{\sigma}_U^X$ and the right panel plots $\hat{\sigma}_L^X$ (green dots), both as functions of σ_L^{NE}. The data can be supported as QRE or NBE outcomes for some primitives (held fixed across games) if and only if the data is in the gray region, decreasing in the left panel, and increasing in the right panel.

(iii) $\hat{\sigma}_U^X$ *is strictly decreasing in* $\sigma_L^{NE,X}$*, and*
(iv) $\hat{\sigma}_L^X$ *is strictly increasing in* $\sigma_L^{NE,X}$*.*

Proof. See Appendix 11.2. □

Since we are concerned with tracking patterns of behavior across games, we would like a plot to help visualize both the data and model predictions from the entire set of X-games. However, since the set of attainable QNBE mixed action profiles is not rectangular (see Figure 1), it is too cumbersome to plot the QNBE predictions as a function of one variable. For this reason, we provide an analogue of Proposition 5 for QRE and NBE, which *are* rectangular. The characterizations for both models coincide.

Proposition 6. *Let $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ be a dataset of mixed actions for any finite number of X-games. The data can be supported as QRE or NBE outcomes for some primitives (held fixed across games) if and only if*

(i) $\hat{\sigma}_U^X \in (\frac{1}{2}, 1)$ *for* $\sigma_L^{NE,X} < \frac{1}{2}$*;* $\hat{\sigma}_U^X \in (0, \frac{1}{2})$ *for* $\sigma_L^{NE,X} > \frac{1}{2}$*,*
(ii) $\hat{\sigma}_L^X \in (\sigma_L^{NE,X}, \frac{1}{2})$ *for* $\sigma_L^{NE,X} < \frac{1}{2}$*;* $\hat{\sigma}_L^X \in (\frac{1}{2}, \sigma_L^{NE,X})$ *for* $\sigma_L^{NE,X} > \frac{1}{2}$*,*
(iii) $\hat{\sigma}_U^X$ *is strictly decreasing in* $\sigma_L^{NE,X}$*, and*
(iv) $\hat{\sigma}_L^X$ *is strictly increasing in* $\sigma_L^{NE,X}$*.*

Proof. See Appendix 11.2. □

Figure 2 plots the sets of attainable QRE and NBE as functions of σ_L^{NE}, which are given in the proposition. The vertical dotted lines correspond to specific values of X (marked at the top). We also plot a hypothetical dataset $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ as green dots: the left panel plots $\hat{\sigma}_U^X$ and the right panel plots $\hat{\sigma}_L^X$–both as functions of σ_L^{NE}. The proposition says that a dataset can be supported as QRE or NBE outcomes if and only if it looks qualitatively like the green dots in the figure: in the gray regions, decreasing in the left panel, and increasing in the right.

3 Experimental Design

Recall that the goal of our experiment is to make observable the empirical action- and belief-maps, which we pursue through collecting actions and beliefs data for a family of games. An important consideration is to be able to interpret within-subject variations in actions and beliefs as the result of idiosyncratic "noise" as opposed to other predictable variations.

3.1 Overall structure

The experiment consisted of two treatments, which we label [A,BA] and [A,A]. Our sessions were run in the Columbia Experimental Laboratory in the Social Sciences (CELSS). Subjects were mainly undergraduate students at Columbia and Barnard Colleges, all of whom were recruited via the Online Recruitment System for Economics Experiments (ORSEE) (Greiner [2015]).

The main treatment is [A,BA], which we describe here. The treatment [A,A] is similar, but does not involve belief elicitation. It was included to test whether belief elicitation itself has an effect on behavior, and we defer its discussion to Section 8.

The experiment involved 2×2 matrix games, and at the beginning of the experiment, subjects were divided into two equal-sized subpopulations of row and column players, which we refer to as players 1 and 2, respectively. The [A,BA] treatment consisted of *two stages*. Each round of the first stage involved taking actions, and each round of the second stage involved stating a belief *and* taking an action. The name of the treatment reflects this ("A" for "action", "BA" for "belief-action").

In each of the 20 rounds of the first stage, subjects were anonymously and randomly paired and took actions simultaneously. In each of the 40 rounds of the second stage, subjects were

presented with a payoff matrix that appeared in the first stage. Subjects then stated their beliefs over opponents' expected action choices before taking actions.[13] These beliefs were elicited over actions taken by subjects in the first stage, and these actions were similarly paired against randomly selected actions (from the relevant games) taken in the first stage. In this way, subjects in the second stage were both forming beliefs about and playing against subjects from the first stage whose actions had already been recorded. Subjects in the second stage were not paired since they were playing against subjects from the first stage. For this reason, subjects in the second stage were not required to wait for all subjects to finish a round before moving on to the next, though in both stages subjects were required to wait for 10 seconds before submitting their answers. Screenshots of the experimental interface are given in Appendix 11.5.

Before the start of the first stage, instructions (see Appendix 11.1) were read aloud accompanied by slides. These instructions described the strategic interaction and taught subjects how to understand 2×2 payoff matrices. Subjects then answered 4 questions to demonstrate understanding of how to map players' actions in a game to payoff outcomes. All subjects were required to answer these correctly. Subjects then played 4 unpaid practice rounds before proceeding to the paid rounds. After the first stage, additional instructions for the second stage were given. Only at that point were subjects introduced to the notion of a belief and the elicitation mechanism described. Subjects then played 3 unpaid practice rounds before proceeding to the paid rounds of the second stage.

We are interested in observing the stochasticity inherent in beliefs, not changes in beliefs that are due to new information. For this reason, at no point during the experiment (including the unpaid practice rounds) were subjects provided any feedback. In particular, no feedback was provided about other subjects' actions, the outcomes of games, or the accuracy of belief statements. Only at the end of the experiment did subjects learn about the outcomes of the games and belief elicitations that were selected for payment. This also simplifies the analysis because subjects could not condition on the history of play.

We also wish to avoid other non-inherent sources of stochasticity in beliefs. Since we elicited beliefs about the first-stage actions which had already been recorded, multiple elicitations for a given game all refer to the same event. Hence, variation in an individual subject's beliefs for a given game cannot be due to a higher-ordered belief that other subjects were learning. To avoid stochasticity in stated beliefs due to mechanical error, belief

[13] After entering a belief for the first time in a round, subjects could freely modify both their actions and beliefs in any order before submitting. In any case, we see very few revisions of stated beliefs.

statements had to be entered as whole numbers into a box rather than via a slider.

Each game was played multiple times. This was necessary because we wish to analyze stochasticity and patterns in individual subjects' belief data. However, we took several measures to approximate a situation in which each game was seen as if for the first time. First, there was no feedback, as described. Second, there was a large "cross section", i.e. more distinct games than the number of times each game was played. Third, the games appeared in a random order which is described in Section 3.2.

In addition to a \$10 show-up fee, subjects were paid according to one randomly selected round (based on actions) from the first stage and four randomly selected rounds from the second stage–two rounds based on actions and two rounds based on beliefs (see Section 3.3 for details on belief payments). Since there were twice as many rounds in the second stage as in the first stage, this equated the incentives for taking actions across the stages. Each unit of payoff in the matrix corresponded to a probability point of earning \$10 (e.g. 20 is a lottery that pays \$10 with probability 20% and \$0 otherwise). This was to mitigate the effects of risk aversion as expected utility is linear in probability points.[14] This is important for our purposes since several of our tests require that utilities are identified.

To allay any hedging concerns, all five payments were based on different matrices and this was emphasized to subjects.

Treatment	Player 1-subjects	Player 2-subjects	Total
[A,BA]	54	56	110
[A,A]	27	27	54
Total	81	83	164

Table 2: *Overview of experiment.*

Table 2 summarizes the number of subjects who participated in the experiment by treatment and player role.[15] On average, the experiment took about 1 hour and 15 minutes, and the average subject payment was \$19.5.

[14]Evidence suggests that this significantly, but not completely, linearizes payoffs in the sense that people still behave as if they have a utility function over probability points with some curvature. See for example, Harrison et al. [2012].

[15]There are two fewer player 1-subjects than player 2-subjects in [A,BA]. This is because two subjects (in separate sessions) had to leave early. They left after the first stage, and since the whole experiment was anonymous and without feedback and the second stage was played asynchronously, this had no effect on the rest of the subjects. These two subjects' data was dropped prior to analysis.

3.2 The games

As discussed in Section 2.5, the X-games take center-stage since they are predicted to give rise to systematic variation in actions and beliefs. Henceforth, we say "$X80$" to refer to the game $X = 80$ and similarly for the other games.

The X-games have other important features for the experiment. Since they are very simple and fully mixed, we would not expect there to be much no-feedback learning (Weber [2003]). This is important since we are studying stochasticity in beliefs, and so want to minimize variation in beliefs due to learning. The payoffs are also "sparse" in the sense of having many payoffs set to 0. This makes the games' structure more transparent and easier to calculate best responses. The fact that one player's payoffs are symmetric and fixed across games also makes it easier to perceive differences across games.

For the experiment, we chose the six values of X given in Table 3. These correspond to the vertical lines in Figure 2. They were chosen so that the corresponding values of σ_L^{NE} are relatively evenly spaced on the unit interval and come close to the boundary at one end. The values of X also go well above and well below 20 so that across the set of games, one player does not always expect to receive higher payoffs. Games $X80$ and $X5$ as well as $X40$ and $X10$ are symmetric-pairs in that $\sigma_L^{NE,X80} = 1 - \sigma_L^{NE,X5}$ and $\sigma_L^{NE,X40} = 1 - \sigma_L^{NE,X10}$. This does not, however, imply the same relation for QNBE without additional conditions.[16]

X	80	40	10	5	2	1
σ_L^{NE}	0.2	0.333	0.667	0.8	0.909	0.952
σ_U^{NE}	0.5	0.5	0.5	0.5	0.5	0.5

Table 3: *Selection of X-games.*

In addition to the X-games, we also included the games whose payoffs are in Appendix Table 16. $D1$ and $D2$ are dominance solvable games, which are identical up to which player faces which set of payoffs. These are included in order to derive, for each subject, a measure of strategic sophistication (using a level k-type framework), which we conjectured would help to rationalize observed deviations from theoretical predictions. We discuss the dominance

[16] If Q is *scale invariant* ($Q_i(\beta v_i) = Q_i(v_i)$ for $\beta > 0$) and *label invariant* ($Q_{i1}((v,w)) = Q_{i2}((w,v))$ for any payoffs $v, w \in \mathbb{R}$) and σ^* is *label invariant* ($F_i(\bar{\sigma}_j|\sigma_j) = 1 - F_i(\bar{\sigma}_j|1-\sigma_j)$ for all $\sigma_j, \bar{\sigma}_j \in (0,1)$), then $\sigma_L^{QNBE,X} = 1 - \sigma_L^{QNBE,X'}$ and $\sigma_U^{QNBE,X} = 1 - \sigma_U^{QNBE,X'}$ if $\sigma_L^{NE,X} = 1 - \sigma_L^{NE,X'}$.

Stage	Games	Rounds of each	Rounds
A	$X80$, $X40$, $X10$, $X5$, $X2$, $X1$	2	20
	$D1$, $D2$	1	
	$X80s$	2	
	$R1$, $R2$	2	
BA	$X80$, $X40$, $X10$, $X5$, $X2$, $X1$	5	40
	Di	3	
	Dj	2	
	$X80s$	5	

Table 4: *Games by section.*

solvable games at length in Section 6. For brevity, we will not discuss the data from the three remaining games, $X80s$, $R1$, and $R2$, in this paper.[17]

Table 4 summarizes the games played in both stages of the experiment and the number of rounds for each. Note that, for each of the X-games, there are two rounds in the first stage and five rounds in the second stage. The dominance solvable games appeared at fixed, evenly spaced rounds.[18] The other games appeared in random order subject to the same game not appearing more than once within 3 consecutive rounds. Subjects were told nothing about what games to expect, the number of times each was to appear, or their order.

3.3 Eliciting beliefs using random binary choice

We used the random binary choice (RBC) mechanism (Karni [2009]) to incentivize subjects to state their beliefs accurately.[19] In an RBC, subjects are asked which option they prefer from a list of 101 binary choices, as in Table 5 with option A on the left and option B on the right. If a subject holds belief $b\%$ over the probability that event E occurs and her preferences respect stochastic dominance (in particular, she does not have to be risk-

[17] These were included to test specific hypothesis related to the models under scrutiny. $X80s$ ("s" for "scale") is the same as $X80$, except with all payoffs divided by 10. This was included because QRE and NBE make very different predictions with respect to scaling payoffs (Friedman [2019]), and this gives direct insight into the effects of scale on beliefs and the effects of scale on actions given beliefs. $R1$ and $R2$ are similar to $X5$, except the symmetry of player 2's payoffs have been broken. These were included to determine if payoff symmetry is driving results.

[18] For a subject in role i in the first stage, Di and Dj appeared in rounds 7 and 14 or 14 and 7 with equal probability. In the second stage, Di appeared in rounds 7, 21, and 35, and Dj appeared in rounds 14 and 28.

[19] Another popular method for incentivizing beliefs is the quadratic scoring rule (see, for example, Nyarko and Schotter [2002]), which has advantages but requires risk neutrality for incentive compatibility. Schotter and Trevino [2014] and Schlag et al. [2015] review these and other elicitation mechanisms.

neutral), it is optimal to choose option A for questions numbered *less than b* and option B for questions numbered *greater than b*. Otherwise, the subject is failing to choose the option that she believes gives the highest probability of receiving the prize.

	Would you rather have:		
	Option A:		Option B:
Q.0	\$5 if the event E occurs	or	\$5 with probability 0%
Q.1	\$5 if the event E occurs	or	\$5 with probability 1%
Q.2	\$5 if the event E occurs	or	\$5 with probability 2%
	⋮		⋮
Q.99	\$5 if the event E occurs	or	\$5 with probability 99%
Q.100	\$5 if the event E occurs	or	\$5 with probability 100%

Table 5: Random binary choice.

Beliefs were elicited in the second stage of the experiment in which the event E was that a randomly selected subject from the first stage chose a particular action. Specifically, subjects were shown a matrix that appeared in the first stage and told that "The computer has randomly selected a round of Section 1 in which the matrix below was played." Player 1 (blue) subjects were then asked "What do you believe is the probability that a randomly selected red player chose L in that round?", and similarly for player 2 (red) subjects (see Appendix 11.5 for screenshots). By entering a belief into a box, a whole number between 0 and 100 inclusive, the rows of the table were filled out optimally given the stated belief (indifference broken in favor of option B). The table did not appear on subjects' screens by default, but they could see it by "scrolling down".

For each round selected for a subject's belief payment, one of the 101 rows was randomly selected and she received her chosen option. If she chose option A in the selected row, a subject of the relevant type was randomly drawn and she received \$5 if the randomly drawn subject chose the relevant option. If she chose option B in the selected row, she received \$5 with the probability given. Since each row was selected for payment with positive probability, subjects were incentivized to state their beliefs accurately. In addition, subjects were told explicitly that it was in their best interest to state their beliefs accurately.

4 Overview of the data

Prior to testing the axioms, we examine the data at a high level. Since our procedures are novel, we benchmark our findings against those from other experiments in which data is

collected in a more conventional way. Since this paper is concerned primarily with noisy behavior, we also explore measures of variability in stated beliefs.

4.1 Actions

Throughout the paper, we refer to actions data from various parts of the experiment and in some cases pool across treatments. For clarity, we use the following notation to indicate the data source:

- [A, ∘]: first-stage actions, pooled across [A,BA] and [A,A]
- [A,BA]: second-stage actions from [A,BA]
- [A,BA]: first-stage actions from [A,BA]
- [A,A]: first-stage actions from [A,A]
- [A,A]: second-stage actions from [A,A]

We focus primarily on [A, ∘] and [A,BA]. We consider [A, ∘] because, in testing axioms on the belief-map, we must compare beliefs to the actions they refer to, and beliefs refer to the first stage. Since there is no feedback provided to subjects and the first stages are identical in [A,BA] and [A,A], we pool across treatments to arrive at [A, ∘]. We consider [A,BA] because, in testing axioms on the action-map, we must associate to each belief statement a corresponding action.

Table 6 gives the empirical frequencies from [A, ∘] and [A,BA]. We observe some differences between the two sets of frequencies. In Section 8, we show that this difference is caused by the process of belief elicitation itself and discuss the implications for our results. This does not affect our main conclusions, but requires that we be careful about what data sources we are using for different tests. In particular, we cannot pool actions data across the two stages.

Even in the first-stage, before we elicit beliefs, our procedures for collecting actions data are somewhat unusual in that we play a large number of games, without feedback, and without the same game appearing consecutively. How does our actions data compare to actions data that is collected under more standard experimental conditions? Figure 3 plots our action frequencies from [A, ∘], superimposed with those from three studies, Ochs [1995], McKelvey et al. [2000], and Rutstrom and Wilcox [2009]. For inclusion, we sought studies

		$X80$	$X40$	$X10$	$X5$	$X2$	$X1$	$D1$	$D2$
[A, ∘]	$\hat{\sigma}_U$	0.50	0.42	0.51	0.40	0.40	0.31	0.04	0.22
	$\hat{\sigma}_L$	0.27	0.25	0.66	0.74	0.74	0.74	0.78	0.02
[A,BA]	$\hat{\sigma}_U$	0.38	0.39	0.65	0.61	0.51	0.49	0.10	0.30
	$\hat{\sigma}_L$	0.21	0.22	0.74	0.78	0.83	0.82	0.80	0.04

Table 6: *Empirical action frequencies.*

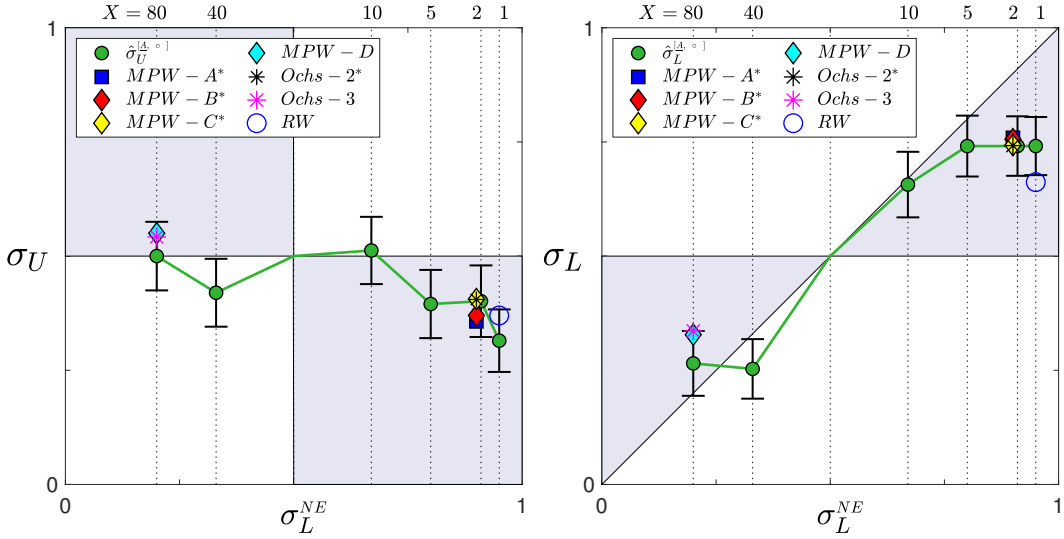

Figure 3: *Actions data.* This figure plots the first stage empirical frequencies [A, ∘] with 90% confidence bands (clustered by subject), superimposed with the empirical frequencies from other studies.

that played games with "sparse" payoffs[20] and $\sigma_L^{NE} = \frac{1}{2}$ (after relabelling). This latter feature allows us to plot their data in our figure as a function of σ_L^{NE}. In these studies, a single game was played 36-50 times consecutively with feedback against either randomly re-matched opponents or a fixed opponent. We find that our data is remarkably close to theirs despite the differences in procedures.

That being said, we cannot find precedents in the literature for games closely matching our more symmetric games–those with σ_L^{NE} relatively close to $\frac{1}{2}$. The only surprising behavior is for $X40$ in which the data falls significantly outside of the QRE-NBE region. In all cases, however, the empirical frequencies from individual games can be supported as QNBE outcomes, as we show in Appendix Figure 26.

[20] Goeree and Holt [2001] played similar games one-shot without sparse payoffs and found data that deviated much farther from NE than in any of these other papers.

4.2 Rates of best response

As we did with the actions data, we compare our findings on beliefs to a benchmark from the literature on belief elicitation. To this end, we look at rates of best response to stated beliefs, which this literature has suggested as a method for validating elicited beliefs (see Schotter and Trevino [2014] for a discussion of this view). Appendix Figure 27 plots histograms of subjects' rates of best response from the X-games. Compared to the study of Nyarko and Schotter [2002] who report an average rate of 75% for an asymmetric matching pennies game played many times with feedback, we find lower rates for player 1 (64%) and higher rates for player 2 (85%).

Appendix Table 15 shows the average rates of best response for each game. Our relatively low rates for player 1 are driven by the very asymmetric games with low values of X. For games with higher values of X that resemble the games from Nyarko and Schotter [2002] more closely, we have very similar rates. That our rates are higher for player 2 is unsurprising since player 2 faces symmetric payoffs and thus has an easier choice to make for any given belief.

4.3 Are beliefs noisy?

To our knowledge, this is the first study to have multiple belief elicitations per subject-game without feedback. A natural question is: are beliefs noisy?

For each subject and X-game, we calculate the spread of her beliefs–the highest belief minus the lowest–across the five belief statements. We average this across the six X-games for each subject to to get an average spread measure. Figure 4 plots histograms of subjects' spreads by player role. There is considerable heterogeneity in spreads, and there is a right tail of very noisy subjects. The average spreads are 24 and 21 belief-points for player 1- and player 2-subjects, respectively. This seems large to us, though these are lower than that expected of the benchmark of uniformly randomizing over a range of 50 belief points (expected value of 33).[21] In unreported results, ANOVA reveals much larger between- than within-subject variance in beliefs for all games and roles. This suggests that, while subjects do have noisy beliefs, patterns in individual subjects' beliefs are relatively stable but heterogeneous.

[21] We consider this a natural benchmark because it results from believing one action is more likely than another but otherwise reporting beliefs randomly subject to that constraint.

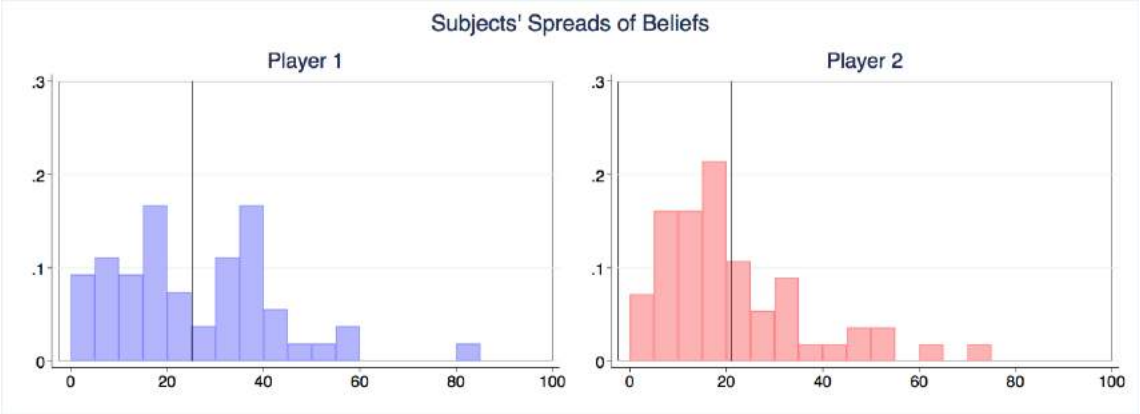

Figure 4: *Subjects' spreads of beliefs.* This figure gives histograms of subjects' spreads of beliefs, averaged across all X-games.

4.4 Actions given expected payoffs (given beliefs)

The premise of quantal response is that beliefs determine actions only insofar as they pin down expected payoff vectors. Hence, we visualize the variation in expected payoffs observed in the data and the extent to which it is predictive of the actions subjects take.

The left panel of Figure 5 plots the convex hull of all expected payoff vectors that we may observe in the data. (v_{i1}, v_{i2}) is a vector of expected payoffs. In the case of player 1, v_{i1} and v_{i2} are the payoffs to U and D, respectively. In the case of player 2, v_{i1} and v_{i2} are the payoffs to L and R, respectively. Each of the straight black lines refer to expected payoffs given beliefs that can be observed in different player-game combinations. The line labelled "$X1$" refers to player 1 in $X1$, the line labelled "$X2$" refers to player 1 in $X2$, and similarly for lines labelled "$X5$", "$X10$", "$X40$," and "$X80$". Recalling that player 2's payoff matrix is fixed across games, the line labelled "$P2$" refers to player 2 in any of these games. The right panel plots, as black dots, the empirical expected payoff vectors (i.e. given stated beliefs) and associated actions, where U and L are coded as 1 and D and R are coded as 0. We also plot a surface that gives the expected action as a function of payoff vectors based on a local linear (lowess) regression. The left panel gives the associated level sets.

From this exercise, we conclude that there is a wide range of belief statements–and thus of expected payoff vectors–both within and across games. Furthermore, this variation is predictive of the actions subjects take.

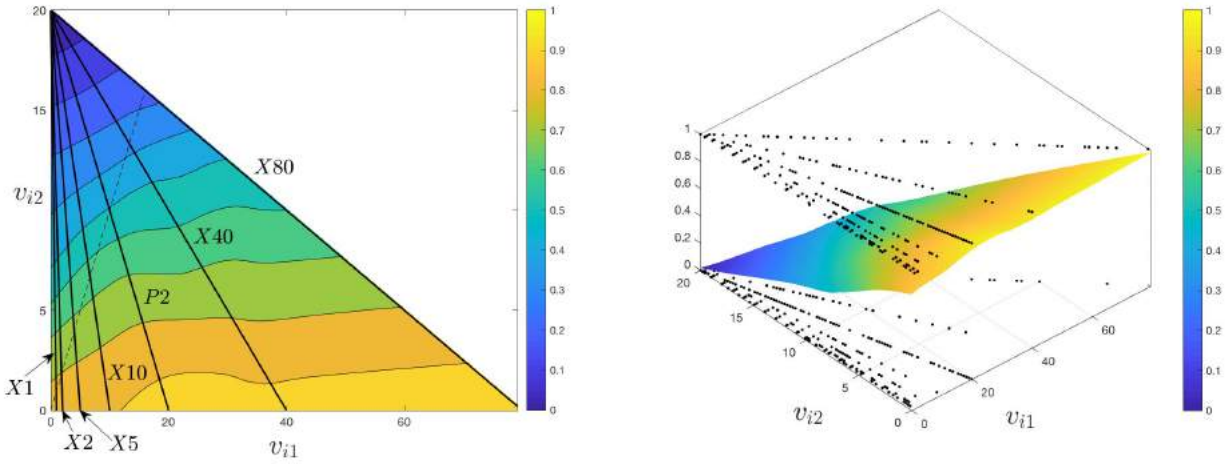

Figure 5: *Actions given expected payoffs (given beliefs).* The left panel gives the convex hull of all expected payoff vectors that may be observed in the data in any of the X-games. (v_{i1}, v_{i2}) refers to either the payoffs to (U, D) for player 1 or the payoffs to (L, R) for player 2. "X1", "X2", ... , and "X80" refer to player 1's vectors in the corresponding games, and "P2" refers to player 2's vectors in any of the games. The right panel plots the action taken as a function of the expected payoff vectors observed in the data, with U and L coded as 1 (D and R coded as 0). The surface is the predicted action from a local linear (lowess) regression (smoothing parameter 0.85). The left panel gives the corresponding level sets.

5 Testing the Axioms

We test each axiom by formulating a statistical hypothesis test with the axiom as the null hypothesis.

5.1 Responsiveness

Responsiveness states that an all-else-equal increase in the expected payoff to some action increases the probability that action is played. To test this, we must associate actions with their expected payoffs given beliefs, and so we use the data from [A,BA].

Since player 1's payoff parameter X is different in each game and there is variation in beliefs across games, not all of player 1's expected payoff vectors across games can be ordered by an all-else-equal increase in the payoff to some action. In such cases, *responsiveness* imposes no restrictions on stochastic choice. With additional conditions, one can complete the order, but we do not pursue that here.[22] Instead, we first consider tests game-by-game.

[22]Consider two unordered vectors, $v_i = (v_{i1}, v_{i2}) = (5, 2)$ and $w_i = (w_{i1}, w_{i2}) = (3, 1)$. One can complete the order with additional restrictions. For instance, if the quantal response function is translation

Then, we consider player 2 only, whose payoff parameters are fixed across games, allowing us to pool data across the entire set of games. In all cases, the variation in expected payoffs is through variation in beliefs.

Since expected payoffs are one-to-one with beliefs within a game, *responsiveness* is easily translated into a condition on beliefs. For player 1 and game x, we state the hypothesis[23] as

$$H_o : Q_U(\bar{u}_1^x(\sigma_L^{'})) \text{ is everywhere weakly } \textit{increasing} \text{ in } \sigma_L^{'}.$$

Similarly, for player 2:

$$H_o : Q_L(\bar{u}_2^x(\sigma_U^{'})) \text{ is everywhere weakly } \textit{decreasing} \text{ in } \sigma_U^{'}.$$

We visualize the relevant data in Figure 6, which plots estimates of $\hat{Q}$ for games $X80$ and $X5$ for both players. Appendix Figure 28 gives the plots for all six games. These are simply the predicted action frequencies from regressing actions on beliefs using a flexible specification (see figure caption for details). Recall that, for each game and player role, there are five observations per subject and so these plots represent a mix of between- and within-subject variation. The vertical dashed line gives the *indifferent belief* $\sigma_j^{'} = \sigma_j^{NE}$ and the horizontal dashed line is set to one-half.

To get a better sense of the raw data, the plots also include belief histograms and the average action within each of ten equally spaced bins (black dots). In some of these bins, there are very few datapoints and so the average action is not very meaningful. The predicted $\hat{Q}$ uses data much more efficiently.

Responsiveness is equivalent to an increasing slope for player 1 and a decreasing slope for player 2. Applying the non-parametric monotonicity test of Bowman et al. [1998] (see Appendix 11.3 for details of implementation),[24] we reject this for both players in all six games with p-values close to 0. However, we must be careful in interpreting this result. Different subjects form different beliefs, and hence the $\hat{Q}$-curves plotted in Figure 6 are

invariant, then $Q_{i1}(v_i) > Q_{i1}(w_i)$ since $Q_{i1}((5,2)) = Q_{i1}((4,1)) > Q_{i1}((3,1))$ where the inequality is due to *responsiveness*. If the quantal response function is scale invariant, then $Q_{i1}(w_i) > Q_{i1}(v_i)$ since $Q_{i1}(3,1) = Q_{i1}(6,2) > Q_{i1}(5,2)$ where the inequality is due to *responsiveness*.

[23] The null hypothesis, by allowing for weak monotonicity, is slightly weaker than the axiom, but it allows for the use of more standard tests.

[24] It is a bootstrap-based test where the data generating process is, heuristically, the best-fit (non-parametrically estimated) monotonic (upward sloping for player 1, downward sloping for player 2) function plus noise, and the p-value is constructed as the fraction of simulations for which a non-parametric regression estimator is non-monotonic.

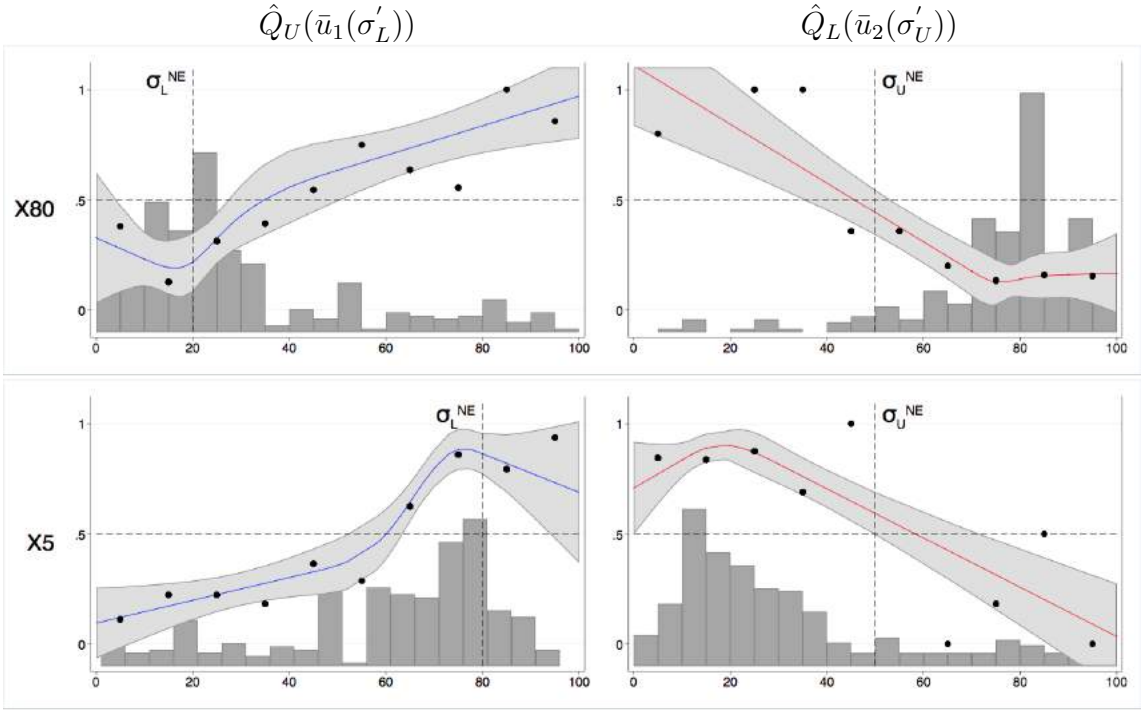

Figure 6: *Action frequencies predicted by beliefs.* For each player and games $X80$ and $X5$, we plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms appear in gray and the average action within each of ten equally spaced bins appear as black dots. The vertical dashed line is the indifferent belief $\sigma_j' = \sigma_j^{NE}$, and the horizontal line is set to one-half.

patched together from different subjects representing different parts of the domain. Hence, the violation could result from individual subjects who violate *responsiveness* to variations in their own beliefs or it could be, in the case of player 1 (and similarly for player 2), there are subjects who tend to hold lower beliefs and favor taking U (whose payoff increases in beliefs). This latter possibility could lead to violations of *responsiveness* even if all individual subjects are responsive to variations over the range of their own stated beliefs.

To determine if individual subjects are responsive to variations in their own stated beliefs, we run fixed effects regressions for different regions of stated beliefs (*responsiveness* is a local property, so we wish to maintain some flexibility in the specification). Let $\{a_{sl}^{ix}, b_{sl}^{ix}\}$ be the lth action-belief pair of subject s in role i in game x. As has been our convention, the actions of U and L are coded as 1, and D and R are coded as 0 (e.g. $a_{sl}^{ix} = 1$ if player 1 takes U). Let $\bar{a}_s^{ix} \equiv \frac{1}{5}\sum_l a_{sl}^{ix}$ and $\bar{b}_s^{ix} \equiv \frac{1}{5}\sum_l b_{sl}^{ix}$ be the subject-level averages. For each role i and game x, we run the following regression for each tercile of belief statements $\{b_{sl}^{ix}\}_{sl}$, which we label as "low", "medium", and "high" beliefs:[25]

$$a_{sl}^{ix} - \bar{a}_s^{ix} = \beta(b_{sl}^{ix} - \bar{b}_s^{ix}) + \varepsilon_{sl}^{ix}. \tag{3}$$

Since there is no difference across subjects in the averages of their demeaned variables (by construction), the coefficient estimate $\hat{\beta}$ reflects within-subject variation.

The results are displayed in Table 7. Consistent with *responsiveness*, we find that every slope is positive for player 1 and all but one (which is extremely close to 0 and insignificant) are negative for player 2, with many of these being highly statistically significant. Furthermore, the magnitudes are large: a majority of slopes have an absolute value greater than 0.005,[26] indicating that a 1 percentage point change in belief is associated with a 0.5 percentage point change in the probability of taking an action. Since the slopes all have the sign predicted by *responsiveness*, this suggests that individual subjects are overwhelmingly responsive.

Following our discussion at the beginning this section, we now turn to player 2 data pooled across all games. Using this data, the top panel of Figure 7 reproduces Figure 6, and Appendix Table 17 presents results of the fixed effects regressions. Since we have much more data that is distributed more uniformly within the space of possible beliefs, we run regressions for each belief-quintile instead of tercile (first column) and also present a version

[25] Results are largely unchanged, but a bit underpowered, if instead use 4 or 5 bins.

[26] For player 1, the absolute slopes average 0.065 and range from 0.000-0.015. For player 2, the average is 0.065 and range from 0.000-0.018.

Player 1

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1
low beliefs	0.000	0.006*	0.008**	0.010***	0.005**	0.004**
	(0.958)	(0.077)	(0.017)	(0.002)	(0.035)	(0.043)
medium beliefs	0.007**	0.010**	0.015***	0.005	0.006	0.006*
	(0.033)	(0.020)	(0.000)	(0.153)	(0.141)	(0.051)
high beliefs	0.005*	0.010***	0.004	0.005	0.004	0.007***
	(0.052)	(0.002)	(0.448)	(0.164)	(0.283)	(0.004)
Observations	270	270	270	270	270	270

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Player 2

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1
low beliefs	-0.010***	-0.018***	-0.005**	-0.007*	-0.004*	-0.004
	(0.001)	(0.000)	(0.038)	(0.074)	(0.071)	(0.133)
medium beliefs	-0.013***	-0.006	-0.002	-0.000	-0.010*	0.000
	(0.000)	(0.174)	(0.490)	(0.930)	(0.058)	(0.969)
high beliefs	-0.008**	-0.004	-0.010***	-0.007**	-0.004	-0.005***
	(0.015)	(0.199)	(0.003)	(0.046)	(0.127)	(0.009)
Observations	280	280	280	280	280	280

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 7: *Fixed effect regressions of actions on beliefs.* For each game and player, we divide individual belief statements into terciles–low, medium, and high beliefs. For each belief tercile, we run a separate linear regression of actions on beliefs that are both first demeaned by subtracting subject-specific averages. Standard errors are clustered by subject.

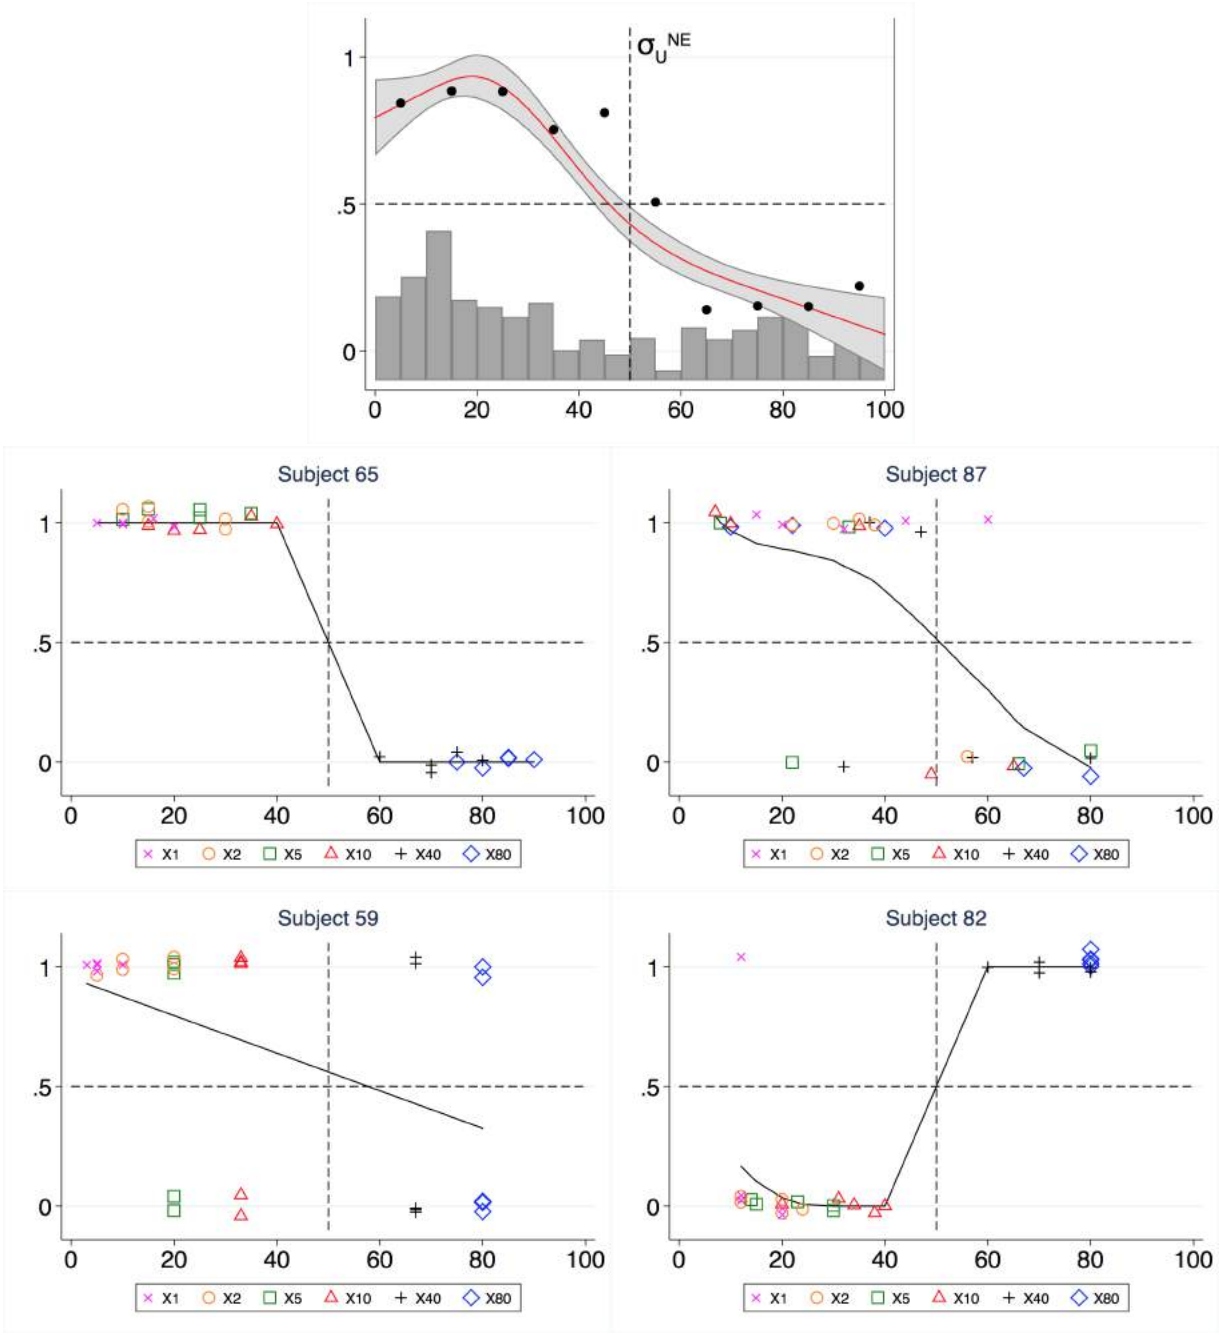

Figure 7: *Player 2 subjects–actions and beliefs, pooled across games.* All plots involve player 2-subjects whose data is pooled across all games. Action L is coded as 1, and action R is coded as 0. The top panel uses all player 2-subjects and gives the predicted action frequencies (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, std. errors clustered by subject) superimposed over the histogram of beliefs. The remaining plots are for specific player 2 subjects. The solid black curve is the parametric regression estimator used in Step 2 of the Bowman et al. [1998] test (Appendix 11.3), and data is separately marked for each game. All datapoints involve a value of 1 or 0 on the vertical axis, but are plotted with a bit of (vertical) noise for visual clarity.

using evenly spaced bins of 20 belief points (second column). Consistent with *responsiveness*, we find that every slope is highly statistically negative and with large magnitudes.

An interesting question is whether within-subject variation in beliefs has predictive power only insofar as beliefs go on one side or the other of the indifferent belief. Inspecting Appendix Table 17, the answer is definitive. Even for player 2, whose indifferent belief is salient, constant across games, and invariant to curvature in the utility function, this variation is highly predictive of actions. Restricting attention to beliefs that are in the bottom or top quintiles–at least 30 points away from the indifferent belief–a 1 percentage point change in belief is associated with a 0.5-0.6 percentage point change in the probability of taking an action.

Being able to pool across games for player 2-subjects results in many more datapoints per subject (30 as opposed to 5) that typically cover much more of the space of possible beliefs. In particular, this allows us to test for *responsiveness* for individual subjects using the Bowman et al. [1998] test. Figure 7 plots some representative individual subjects' data pooled across all six games, superimposed by the non-parametric regression estimator used in Step 2 of the Bowman et al. [1998] test (see Appendix 11.3). The four subjects depicted in the figure are representative of the types of subjects we observe: subject 65 is characterized by step function-like *responsiveness* and always best responds;[27] subject 87 is also responsive, but has action-probabilities that are more continuous in beliefs; subject 59 is similar to subject 87 but noisier, and the non-parametric test rejects *responsiveness*; subject 82 is an "opposite type" who fails *responsiveness* trivially. In all, *responsiveness* is rejected in only 19% of player 2-subjects.

5.2 Monotonicity

Monotonicity is a weakening of best response which states that, *given beliefs*, the action with a higher expected payoff is played more often than not and, if players are indifferent, they uniformly randomize. Since we must associate expected payoffs given beliefs to actions, we again use the data from [A,BA].

For the games studied in this paper, since players are indifferent when their beliefs equal the opponent's Nash equilibrium strategy, *monotonicity* takes a particularly simple form. For player 1 and game x, we state the hypothesis as

[27] Many subjects look similar to subject 65 except with up to 2 "mistakes", which typically does not lead to a rejection of *responsiveness*.

$$H_o : Q_U(\bar{u}_1^x(\sigma_L^{'})) \gtreqless \frac{1}{2} \text{ if and only if } \sigma_L^{'} \gtreqless \sigma_L^{NE,x}.$$

Similarly, for player 2 and game x:

$$H_o : Q_L(\bar{u}_2^x(\sigma_U^{'})) \gtreqless \frac{1}{2} \text{ if and only if } \sigma_U^{'} \lesseqgtr \sigma_U^{NE,x}.$$

In order to visualize potential *monotonicity* violations, we appeal once again to Figure 6, which plots estimates of $\hat{Q}$ for games $X80$ and $X5$ for both players (see figure caption for details; see Appendix Figure 28 for all six games). The vertical dashed line gives the indifferent belief $\sigma_j^{'} = \sigma_j^{NE}$ and the horizontal dashed line is set to one-half. As opposed to *responsiveness* that concerns the slope, *monotonicity* concerns the levels of the graph. Specifically, for player 1 (left panels), $\hat{Q}_U$ should be less than $\frac{1}{2}$ to the left of the vertical line and greater than $\frac{1}{2}$ to the right of the vertical line; for player 2 (right panels), $\hat{Q}_L$ should be greater than $\frac{1}{2}$ to the left of the vertical line and less than $\frac{1}{2}$ to the right of the vertical line.

In testing *monotonicity*, we conduct the analysis at the aggregate level since we have only 5 belief statements for each subject-game. Unlike for *responsiveness*, there is no issue in aggregation. Since *monotonicity* is a condition that holds pointwise, if all subjects have monotonic quantal response over the range of their stated beliefs (even if different subjects form very different beliefs), the aggregate will also be monotonic.

Our test for *monotonicity* is the natural one suggested by eyeballing Figure 6. After running flexible regressions of actions on beliefs, we calculate the standard error of the prediction (clustering by subject), which we use to calculate error bands for the estimated $\hat{Q}$. From the figure, one can observe rejections of the null at the given level of significance. For instance, in the top left panel (game $X80$, player 1), we see that for beliefs just above 20, whereas *monotonicity* requires that Q should be above $\frac{1}{2}$, we observe that the estimated $\hat{Q}$ is significantly below $\frac{1}{2}$. Since it is the 90% error band that is plotted, inspection reveals that *monotonicity* is rejected with a p-value less than 0.1. Similarly, if the 95% error band still leads to a violation, then the p-value is less than 0.05. By considering error bands of increasing size, all violations will eventually disappear. Hence, we calculate the p-value as c, where the $100(1-c)\%$ error band is the smallest which results in no violations.

One weakness of the test is that it is sensitive to the regression specification, so we report the results (p-values) of the statistical tests in Table 8 for 5 different specifications (see table caption for details). The second panel of the table gives a reduced-form measure of the degree of *monotonicity* violations–the total area enclosed between $\hat{Q}$ and the one-half line

Tests of Monotonicity (p-values)

	Player 1 (Q_U)						Player 2 (Q_L)					
	$C4^*$	$L4^*$	$C5$	$C6$	$C7$	Avg	$C4^*$	$L4^*$	$C5$	$C6$	$C7$	Avg
$X80$	0.00***	0.02**	0.01***	0.03**	0.04**	0.02**	0.35	0.36	0.47	0.42	0.50	0.42
$X40$	0.01***	0.00***	0.00***	0.00***	0.01***	0.00***	0.05**	0.05**	0.23	0.47	0.87	0.34
$X10$	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.09*	0.02**	0.13	0.66	0.87	0.35
$X5$	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.11	0.10*	0.22	0.41	0.58	0.28
$X2$	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.07*	0.06*	0.13	0.17	0.22	0.13
$X1$	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.30	0.34	0.86	0.77	0.97	0.65
Avg	0.00***	0.00***	0.00***	0.00***	0.01***	–	0.16	0.16	0.34	0.48	0.67	–

Size of Monotonicity Violation

	Player 1 (Q_U)						Player 2 (Q_L)					
	$C4^*$	$L4^*$	$C5$	$C6$	$C7$	Avg	$C4^*$	$L4^*$	$C5$	$C6$	$C7$	Avg
$X80$	2.08	1.88	1.69	2.29	2.19	2.03	0.15	0.17	0.20	0.28	0.25	0.21
$X40$	0.82	1.68	1.57	1.27	1.39	1.35	0.32	0.38	0.34	0.19	0.02	0.25
$X10$	2.33	2.86	3.33	3.67	3.45	3.13	0.19	0.29	0.26	0.01	0.00	0.15
$X5$	5.27	5.82	5.50	5.68	5.85	5.62	0.34	0.42	0.81	0.46	0.30	0.46
$X2$	6.54	6.63	6.83	6.49	6.47	6.59	1.35	1.33	1.90	1.81	1.81	1.64
$X1$	5.85	5.84	5.65	5.29	5.62	5.65	0.51	0.48	0.06	0.02	0.00	0.21
Avg	3.81	4.12	4.10	4.12	4.16	–	0.48	0.51	0.60	0.46	0.40	–

Table 8: *Testing monotonicity.* For each player and game, we test for *monotonicity* in the manner described in Section 5.2 using 5 different regression models to estimate $\hat{Q}$. The 5 models are based on restricted splines: cubic with 4 knots based on belief quintiles ($C4^*$); linear with 4 knots based on belief quintiles ($L4^*$); and cubic with 5, 6, or 7 equally spaced knots ($C5$, $C6$, and $C7$, respectively). The top panel reports p-values, as well as the p-values averaged across games for a given model and averaged across models for a given game. The bottom panel reports a reduced-form measure of *monotonicity* violations–the total area enclosed between $\hat{Q}$ and the one-half line over beliefs that lead to (not necessarily significant) violations.

over beliefs that lead to (not necessarily significant) violations.

We find that *monotonicity* cannot be rejected for player 2 in any game, and this is consistent across regression specifications. In particular, it is not rejected with very high p-values for the most flexible specifications (see table caption for details). For player 1, on the other hand, we observe consistent and highly significant violations of *monotonicity* in all games that occur over a region of 5-30 belief points, depending on the game. Moreover, based on the belief histograms in Figure 6, it is clear that a large mass of beliefs (including the mode) fall in the regions with *monotonicity* violations.

From Figure 6, it is clear that the nature of player 1's *monotonicity* violations is systematic. For $X > 20$, the violations occur over an interval of beliefs just "right of" the indifferent belief, and for $X < 20$, the violations are over an interval of beliefs just "left of" the indifferent belief. We consider explanations for this pattern in Section 6.1.

A weak implication of *monotonicity* is that best responses will be taken with probability greater than one-half. As shown previously in Appendix Table 15, best responses are taken with probability greater than one-half in all games. Thus, even though subjects tend to best respond to the beliefs that they form, they systematically fail to best respond to beliefs that realize in particular regions of the belief-space. Hence, our analysis expands upon previous studies using elicited beliefs (e.g. Costa-Gomes and Weizsacker [2008] and Rey-Biel [2009]) that have focused only on rates of best response.

5.3 Belief-responsiveness

Belief-responsiveness states that, if the frequency of player j's action increases, so too does the distribution of player i's beliefs in the sense of first-order stochastic dominance. Recalling that the beliefs are elicited about behavior in the first stage and that the first stages are identical across the treatments, we use the beliefs data from [A,$\underline{\text{B}}$A] and the actions data from [$\underline{\text{A}}$, $\circ$].

Across games x and y, we seek tests of the form[28]

$$\begin{aligned} H_o : &\sigma_j^x > \sigma_j^y \text{ and } F_i(\cdot|\sigma_j^x) \succ_{FOSD} F_i(\cdot|\sigma_j^y), \text{ or} \\ &\sigma_j^y > \sigma_j^x \text{ and } F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x). \end{aligned} \tag{4}$$

Prior to testing, we visualize the data in Figure 9, which plots histograms of stated beliefs, superimposed with median beliefs (solid vertical lines) and the corresponding empirical

[28]The null hypothesis here is slightly weaker than used in the axiom, but it allows for the use of more standard tests.

frequencies of actions (dashed vertical lines). It appears that the distributions of beliefs shift monotonically in X in the direction predicted by QNBE: as X increases, player 2 believes that player 1 will play U more often and player 1 believes player 2 will play L less often. Furthermore, plotting the CDFs of beliefs in Figure 8 suggests that the belief distributions are ordered by stochastic dominance. The empirical action frequencies also typically, but not always, move in the same direction, consistent with *belief-responsiveness*.

Our test of hypothesis (4) is simple and conservative in the sense of not over-rejecting. To this end, we perform one-sided tests of the null hypotheses $H_0 : \sigma_j^x > \sigma_j^y$ and $H_0 : F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x)$. If σ_j^y is greater than σ_j^x, *belief-responsiveness* dictates that $F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x)$, so we say that there is a failure of the axiom if we reject both that $\sigma_j^x > \sigma_j^y$ and that $F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x)$.[29] Table 9 reports the p-values of the one-sided tests of $H_0 : \sigma_j^x > \sigma_j^y$ and $H_0 : F_i(\cdot|\sigma_j^x) \succ_{FOSD} F_i(\cdot|\sigma_j^y)$ for every combination of games x and $y \neq x$ and for each player i and her opponent j (see table caption for details of these tests). These are reported in matrix form as entries in row x and column y. We find only one significant violation across the many comparisons. This can be seen from the p-values in bold, indicating rejections of both $\sigma_1^{X40} > \sigma_1^{X10}$ and $F_1(\cdot|\sigma_1^{X10}) \succ_{FOSD} F_1(\cdot|\sigma_1^{X40})$. We conclude that *belief-responsiveness* cannot be rejected in our data.

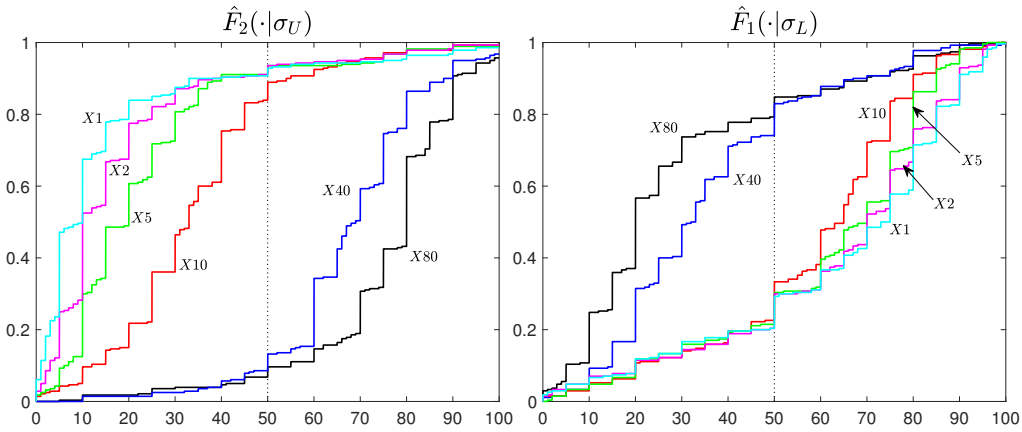

Figure 8: *CDFs of belief distributions.* We plot the empirical CDFs of belief distributions. The left panel is for player 2's beliefs about U, and the right panel is for player 1's beliefs about L.

[29] This would be a conservative test if a situation in which belief distributions were *un*ordered by stochastic dominance did not lead to rejection of the axiom, and in most cases the ordering is clear (Figure 8).

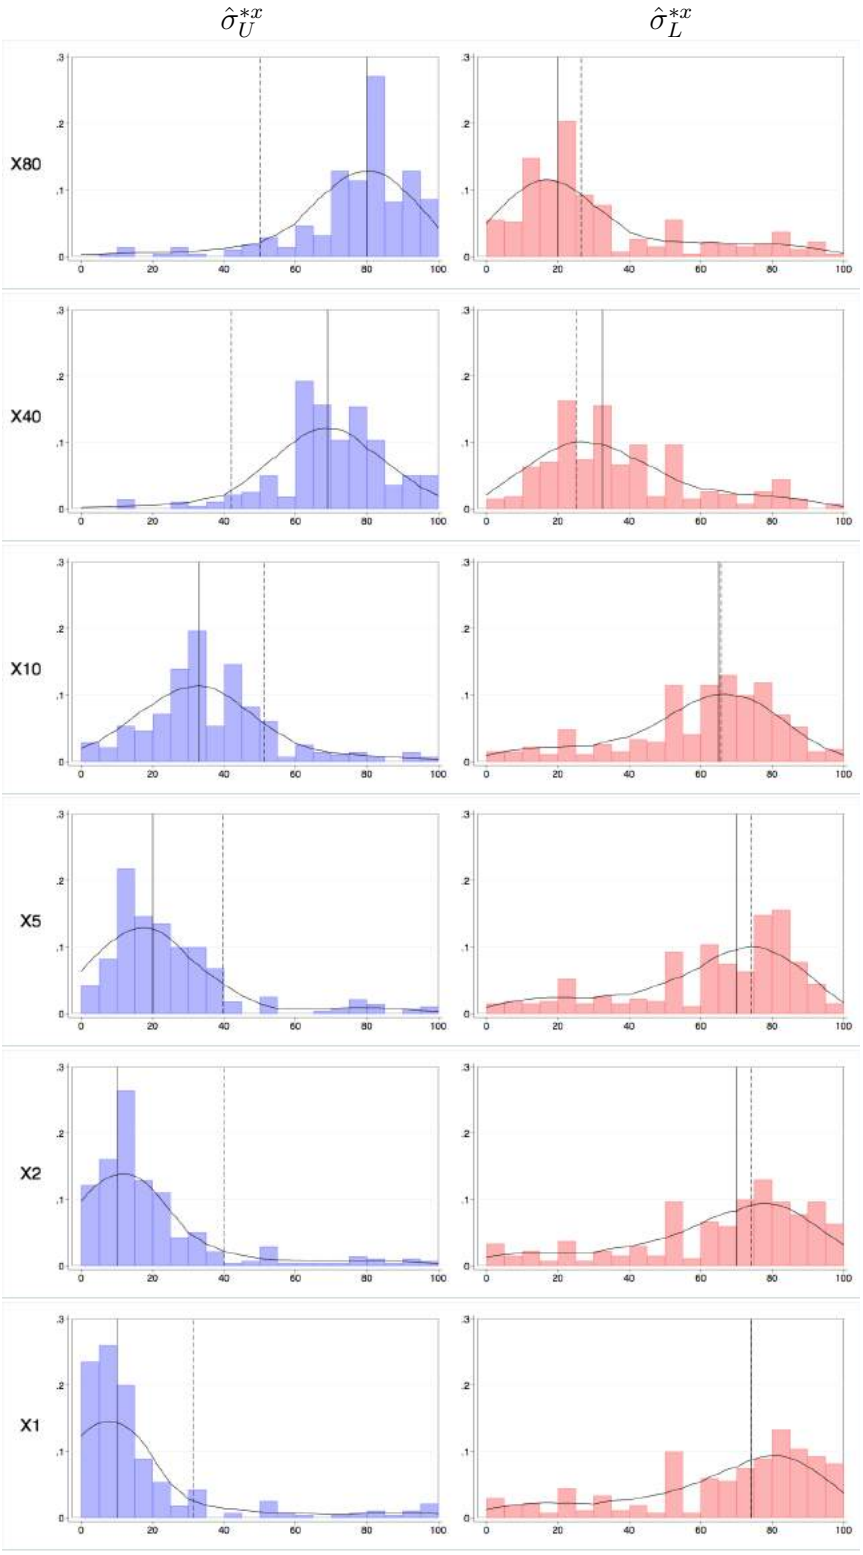

Figure 9: *Belief distributions.* The left panel is for player 2's beliefs about U, and the right panel is for player 1's beliefs about L. The solid lines mark the median of i's beliefs and the dashed line marks the empirical frequency of j's actions.

Actions

$H_0: \sigma_j^x > \sigma_j^y$	Players $i=1, j=2$ (p-values)						Players $i=2, j=1$ (p-values)					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
$X80$	–	0.58	0.00***	0.00***	0.00***	0.00***	–	0.89	0.42	0.95	0.93	1.00
$X40$	0.42	–	0.00***	0.00***	0.00***	0.00***	0.11	–	**0.07***	0.65	0.61	0.96
$X10$	1.00	1.00	–	0.08*	0.08*	0.07*	0.58	0.93	–	0.97	0.95	1.00
$X5$	1.00	1.00	0.92	–	0.50	0.50	0.05*	0.35	0.03**	–	0.46	0.90
$X2$	1.00	1.00	0.92	0.50	–	0.50	0.07*	0.39	0.05**	0.54	–	0.91
$X1$	1.00	1.00	0.93	0.50	0.50	–	0.00***	0.04**	0.00***	0.10*	0.09*	–

Beliefs

$H_0: F_i(\cdot\|\sigma_j^x) \succ_{FOSD} F_i(\cdot\|\sigma_j^y)$	Players $i=1, j=2$ (p-values)						Players $i=2, j=1$ (p-values)					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
$X80$	–	0.00***	0.00***	0.00***	0.00***	0.00***	–	0.87	1.00	1.00	1.00	1.00
$X40$	0.80	–	0.00***	0.00***	0.00***	0.00***	0.00***	–	1.00	1.00	1.00	1.00
$X10$	0.96	0.93	–	0.00***	0.00***	0.00***	0.00***	**0.00*****	–	0.84	0.98	0.73
$X5$	0.97	0.92	0.76	–	0.00***	0.00***	0.00***	0.00***	0.00***	–	0.95	0.77
$X2$	1.00	0.72	0.79	0.68	–	0.01**	0.00***	0.00***	0.00***	0.00***	–	0.78
$X1$	1.00	0.83	0.66	0.72	0.73	–	0.00***	0.00***	0.00***	0.00***	0.00***	–

Table 9: *Testing belief-responsiveness.* The top panel reports p-values from tests of $H_0 : \sigma_j^x > \sigma_j^y$ across games x (row) and y (column). This is from standard t-tests, clustering by subject. The bottom panel reports p-values from tests of $H_0 : F_i(\cdot|\sigma_j^x) \succ_{FOSD} F_i(\cdot|\sigma_j^y)$ across games x (row) and y (column). This is from non-parametric Kolmogorov-Smirnov-type tests in which the test statistic is bootstrapped following Abadie [2002] in a way that preserves the within-subject correlation of beliefs observed in the data (see Appendix 11.3 for details). The entries in bold correspond to the only rejection: $\sigma_1^{X40} \not\succ \sigma_1^{X10}$ and $F_1(\cdot|\sigma_1^{X10}) \not\succ_{FOSD} F_1(\cdot|\sigma_1^{X40})$.

5.4 Unbiasedness

Unbiasedness states that beliefs are correct on median. Hence, for each player j and game x, we test the hypothesis

$$H_o : med(\sigma_j^{*,x}) = \sigma_j^x,$$

where $med(\cdot)$ denotes the median of a random variable. Once again, we use the beliefs data from [A,BA] and the actions data from [A, $\circ$].

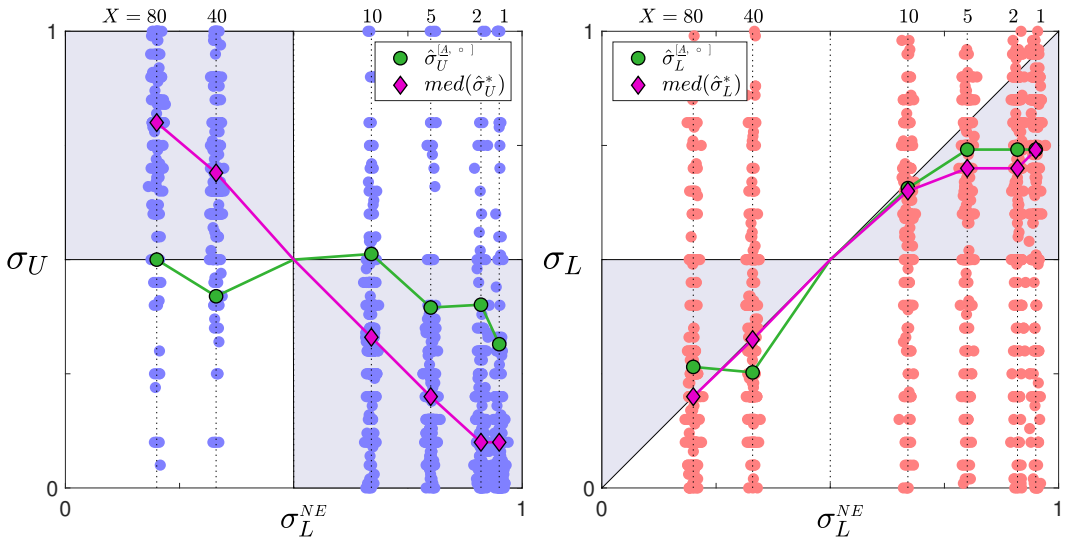

Figure 10: *Bias in beliefs.* The left panel gives player 1's action frequency from [A, $\circ$] and the median of player 2's beliefs about player 1. Blue circles are individual belief statements. The right panel gives player 2's action frequency from [A, $\circ$] and the median of player 1's beliefs about player 2. Red circles are individual belief statements.

Unbiasedness requires that beliefs are correct on median, so we plot in Figure 10 the aggregate action frequencies and median beliefs as well as the individual belief statements. Appendix Table 18 reports the bias in both median- and mean-beliefs with p-values of the hypothesis that beliefs are unbiased (see caption for details).

We find that player 1's beliefs about player 2 ($\hat{\sigma}_L^*$) are remarkably unbiased in that we fail to reject *unbiasedness* for most games individually. When using the mean belief instead of median, we find that there is a small "conservative" bias in the sense that mean beliefs are closer to the uniform distribution than the actual distribution of actions. Such bias has been documented by Huck and Weizsacker [2002] and Costa-Gomes and Weizsacker [2008] in other settings, and is relatively common in experiments in which beliefs are elicited.

More interestingly, we find that player 2's beliefs about player 1 ($\hat{\sigma}_U^*$) are very "extreme" and we reject *unbiasedness* for all games (and similarly for mean-unbiasedness). Whereas player 1's actions are relatively close to uniform for all values of X, player 2 overwhelmingly believes player 1 takes U when X is large and similarly takes D when X is small. Hence, the nature of bias depends qualitatively on player role, which we take as one of the key facts to be explained in the next section.

6 Explaining the Failure of the Axioms

6.1 Monotonicity

We observe, for player 1 only, a systematic failure of *monotonicity*, as shown in Figure 6. For $X > 20$, the violations are for beliefs just "right of" the indifferent belief, and for $X < 20$, the violations are for beliefs just "left of" the indifferent belief. The proposition below states that, with concavity in the utility function over payoffs in the matrix (again, not risk aversion since the payoffs are in probability points), player 1's indifferent belief deviates systematically from that under linear utility. Since monotonic quantal response predicts that $\hat{Q}$ should cross the one-half line at the indifferent belief, this can explain precisely the types of violations we observe.

Proposition 7. *Let w and v be any strictly increasing Bernoulli utility functions. For player 1 in game X, w induces expected utility vectors $\bar{w}^X = (\bar{w}_U^X, \bar{w}_D^X)$. Let $\tilde{\sigma}_L^{w,X}$ be the unique indifferent belief such that $\bar{w}_U^X(\tilde{\sigma}_L^{w,X}) = \bar{w}_D^X(\tilde{\sigma}_L^{w,X})$. (i) If w is more concave than v ($w = f(v)$ for f concave), then $\tilde{\sigma}_L^{w,X} > \tilde{\sigma}_L^{v,X}$ for $X > 20$ and $\tilde{\sigma}_L^{w,X} < \tilde{\sigma}_L^{v,X}$ for $X < 20$. (ii) if w is concave, then $\tilde{\sigma}_L^{w,X} \in (\sigma_L^{NE,X}, \frac{1}{2})$ for $X > 20$ and $\tilde{\sigma}_L^{w,X} \in (\frac{1}{2}, \sigma_L^{NE,X})$ for $X < 20$.*

Proof. See Appendix 11.2. □

Since player 2 faces symmetric payoffs, curvature does not affect her indifferent belief. Hence, both players having concave utility is qualitatively consistent with the whole of the data.

To test for concavity, we fit a random utility model (e.g. Luce [1959]) with curvature to each player 1-subject's actions data given belief statements. Since we will fit random utility models to both players' data later on, we keep the notation general by using i for player role.

The data of subject s in role i is a set of 30 action-belief pairs $\{\hat{a}_{sl}^{iX}, \hat{b}_{sl}^{iX}\}_{lX}$ where $l \in \{1, ..., 5\}$ indexes each elicitation and X indexes the game. We assume that the utility function is the constant relative risk aversion (CRRA) utility function with curvature

parameter ρ, which has been modified to allow for 0 monetary payoffs and normalized so that utility is between 0 and 1:[30]

$$w(z;\rho) = \frac{(z+\epsilon)^{1-\rho} - \epsilon^{1-\rho}}{(80+\epsilon)^{1-\rho} - \epsilon^{1-\rho}}.$$

This utility function induces, for each game X and stated belief $\hat{b}_{sl}^{iX}$, a vector of expected utilities $\bar{w}_i^X(\hat{b}_{sl}^{iX};\rho) = (\bar{w}_{i1}^x(\hat{b}_{sl}^{iX};\rho), \bar{w}_{i2}^X(\hat{b}_{sl}^{iX};\rho))$. We assume that the probability of taking the first action (U in the case of player 1, L in the case of player 2) depends only on this vector, based on the Luce quantal response function with sensitivity parameter $\mu_a > 0$:[31]

$$p_X(\hat{a}_{sl}^{iX}|\hat{b}_{sl}^{iX};\rho,\mu_a) = \frac{\bar{w}_{i1}^X(\hat{b}_{sl}^{iX};\rho)^{\frac{1}{\mu_a}}}{\bar{w}_{i1}^X(\hat{b}_{sl}^{iX};\rho)^{\frac{1}{\mu_a}} + \bar{w}_{i2}^X(\hat{b}_{sl}^{iX};\rho)^{\frac{1}{\mu_a}}}. \tag{5}$$

For subject s in role i, we choose ρ and μ_a to maximize the log-likelihood of observed actions given stated beliefs:

$$L^s(\hat{a}|\hat{b};\rho,\mu_a) = \sum_X \sum_{l=1}^{5} ln(p_X(\hat{a}_{sl}^{iX}|\hat{b}_{sl}^{iX};\rho,\mu_a)).$$

We find that for 37 of 54 player 1-subjects (69%), a likelihood ratio test rejects the restriction of linear utility, that $\rho = 0$, at the 5% level. For 31 of those 37 subjects (84%), the estimated $\hat{\rho}$ is positive, indicating concavity.

We also fit ρ and μ_a to the player 1 data pooled across all subjects and games, i.e. to maximize $L(\hat{a}|\hat{b};\rho,\mu_a) = \sum_s L^s(\hat{a}|\hat{b};\rho,\mu_a)$. We find the estimate $\hat{\rho} = 0.87$, indicating concavity. In Figure 11, we reproduce Figure 6 for player 1, but we now plot the indifferent beliefs implied by the best-fit utility function as solid vertical lines (see Appendix Figure 29 for all six games). Each such line intersects $\hat{Q}_U$ near to where it crosses the horizontal one-half line. Hence, if the subjects emitted a representative agent with this concave utility, nearly all of the *monotonicity* violations would disappear. This also captures the fact that the regions of violations are larger for the more asymmetric games (compare, for example, $X10$ and $X1$ in Appendix Figure 29).

There are several potential explanations as to why subjects' behavior can be rationalized

[30] $\epsilon > 0$ was pre-set arbitrarily to $\epsilon = 0.001$. By construction, $w(0;\rho) = 0$ and $w(80;\rho) = 1$.

[31] The Luce rule (5) fits the data much better than the logit quantal response function, but is undefined when one of the expected utilities is 0. This happens if and only if the stated belief is 0 or 100, which occurs very few times in the data. When this occurs, we instead use 1 or 99, respectively, to calculate the expectations.

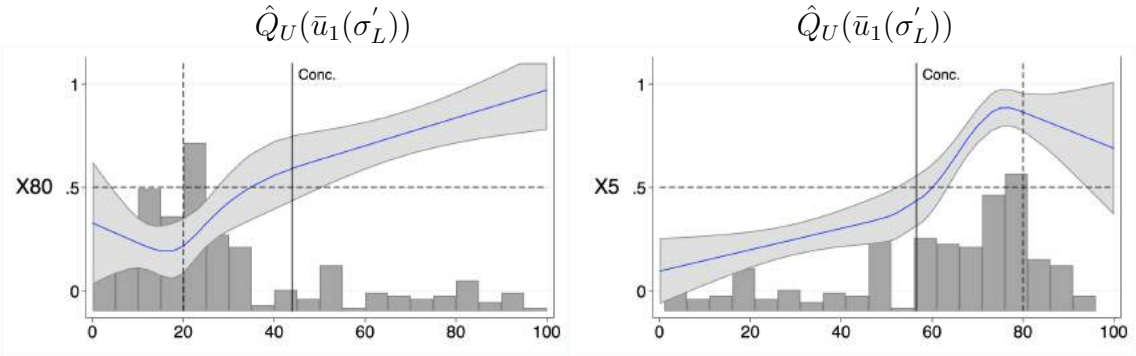

Figure 11: *Concave utility explains monotonicity failures.* For player 1 and games $X80$ (left panel) and $X5$ (right panel), we plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms appear in gray, the vertical dashed line is the risk neutral indifferent belief $\sigma_j' = \sigma_j^{NE}$, and the horizontal line is set to one-half. The solid vertical line is the indifferent belief with concave utility that is estimated from fitting a single curvature parameter to player 1-subjects' data from all X-games.

by concave utility over matrix payoffs. First, it could be that subjects thought of probability points as money and were risk averse. We do not believe, however, that subjects were confused about the nature of payoffs: they had to answer four questions demonstrating understanding of how to map players' actions in a game to payoff outcomes (see Section 3), and these emphasized that payoffs were in probability points. Second, it could be that subjects simply wanted to "win" the game by taking the action that maximized the probability of earning positive probability points unless the other action was sufficiently attractive.

6.2 Unbiasedness

Unbiasedness is rejected in the X-games: player 1 forms unbiased/conservative beliefs whereas player 2 forms extreme beliefs. That the players systematically form qualitatively different biases is mysterious, though the number of models that can rationalize this observation is potentially large. What is the true mechanism?

We argue that the roles subjects find themselves in causally induce different degrees of strategic sophistication in the level k sense (e.g. Nagel [1995] and Stahl and Wilson [1995]). In particular, player 1-subjects are made more sophisticated than player 2-subjects, and this generates precisely the biases we observe: whereas player 2-subjects overwhelmingly attribute the low-level action to player 1 (U when X is large, D when X is small), a sizable fraction of player 1-subjects correctly anticipate this.

We provide two types of corroborating evidence for this sophistication hypothesis. First, all subjects stated beliefs in both roles of a dominance solvable game, from which we derive a subject-level measure of strategic sophistication that is collected identically for all subjects. By this measure, player 1-subjects are much more sophisticated than player 2-subjects. Second, player 1-subjects have much longer response times, suggestive of deeper thinking.

We first present this evidence. Then, we argue that player 1-subjects' stated beliefs in the X-games also suggest much higher levels of strategic sophistication. To this end, we fit a structural model to each subject's beliefs data from the X-games and show, in the context of the model, that this implies much higher levels of sophistication for player 1-subjects. We then validate this finding by showing that the implied measures of sophistication correlate strongly with those measured in the dominance solvable games.[32]

The data tell us that player role itself has a causal effect on sophistication but it cannot tell us the precise mechanism. We conclude this section by discussing potential mechanisms and suggestions for future work.

6.2.1 Dominance solvable games and a measure of sophistication

If player 1-subjects are truly made more sophisticated because of their role in the X-games, we conjecture that this should "spill over" to other games.[33] To this end, we consider the dominance solvable games reproduced in Figure 12.

D1

	$\mathbf{L}_{k\geq 2}$	$\mathbf{R}_{k=1}$
U	0 6	20 0
$\mathbf{D}_{k\geq 1}$	20 8	4 20

D2

	L	$\mathbf{R}_{k\geq 1}$
$\mathbf{U}_{k=1}$	0 20	20 4
$\mathbf{D}_{k\geq 2}$	6 0	8 20

Figure 12: *Dominance solvable games.* In game Di, player i has a strictly dominant action (taken by levels $k \geq 1$). Player j can either best respond to a uniform distribution ($k = 1$) or to player i's dominant action ($k \geq 2$).

Games $D1$ and $D2$ ("Dominant 1" and "Dominant 2") are identical up to which player faces which set of payoffs. In the former, player 1 has a strictly dominant action and in the latter, player 2 has a strictly dominant action. Furthermore, in game Di, one of player j's

[32] This is internally consistent in that the structural model predicts such a correlation between the measures.

[33] One concern is that, since experience in the fully mixed X-games affects behavior in the dominance solvable $D1$ and $D2$, these latter games may also have an affect on behavior in the former. However, we find this implausible since the X-games take up a large majority of the experiment, so we think of behavior in $D1$ and $D2$ as reflections of the cognitive processes used in the X-games.

actions is the best response to a uniform distribution and the other is the best response to i's dominant action.

Assuming level 0 types uniformly randomize, the following characterizes level-types in Di. Player i: levels $k \geq 1$ take the dominant action. Player j: level 1 best responds to a uniform distribution and levels $k \geq 2$ best respond to i's dominant action.

This suggests two benchmark beliefs: (1) i's belief that j takes her dominant action in Dj and (2) i's belief that j best responds to i's dominant action in Di. Assuming i believes j is drawn from a distribution of level types, for any fixed probability that i believes j is level 0, these correspond to increasing functions of i's belief that j is any level $k \geq 1$ and i's belief that j is any level $k \geq 2$, respectively. We call these benchmark beliefs $\beta(k \geq 1)$ and $\beta(k \geq 2)$, and they are readily seen as coarse measures of sophistication as they measure the belief that the opponent is of a sufficiently high level.

Throughout the paper, we aggregate $\hat{\beta}(k \geq 1)$ and $\hat{\beta}(k \geq 2)$ to the subject level by averaging beliefs across instances of Di and Dj, respectively (player i sees Di three times and Dj two times in the second stage). Figure 13 gives histograms of these measures for each player role.

From the top panel of Figure 13, we see that both players have very similar distributions of $\hat{\beta}(k \geq 1)$ that are highly concentrated toward the right of the space with modes close to 100 and very similar means of approximately 85 (solid lines). The corresponding action frequencies (from $[\underline{A}, \circ\,]$) are even higher: greater than 95 for both players (dashed lines). From the bottom panel of Figure 13, we see that player 1's distribution of $\hat{\beta}(k \geq 2)$ is relatively uniform whereas that of player 2 is concentrated below 50, and the respective means are 56 and 33 (solid lines). The corresponding action frequencies are nearly the same for both players at approximately 78 (dashed lines).

Our main takeaways from Figure 13 are twofold. First, there is much more variation in $\hat{\beta}(k \geq 2)$ than in $\hat{\beta}(k \geq 1)$. In other words, subjects overwhelmingly believe other subjects respond to incentives, but vary greatly in how many additional steps of reasoning they perform. Therefore, we will use $\hat{\beta}(k \geq 2)$ as our measure of sophistication.

Second, player 1 is *much more sophisticated* than player 2 by this measure, with an average difference in $\hat{\beta}(k \geq 2)$ of 22. In Table 19 of Appendix 11.7, we show that this sophistication gap is highly significant, robust to various controls, and not driven by erratic subjects.

Importantly, since $D1$ and $D2$ are exactly the same up to which player faces which payoffs, the sophistication measure is derived in exactly the same way for both players. Furthermore,

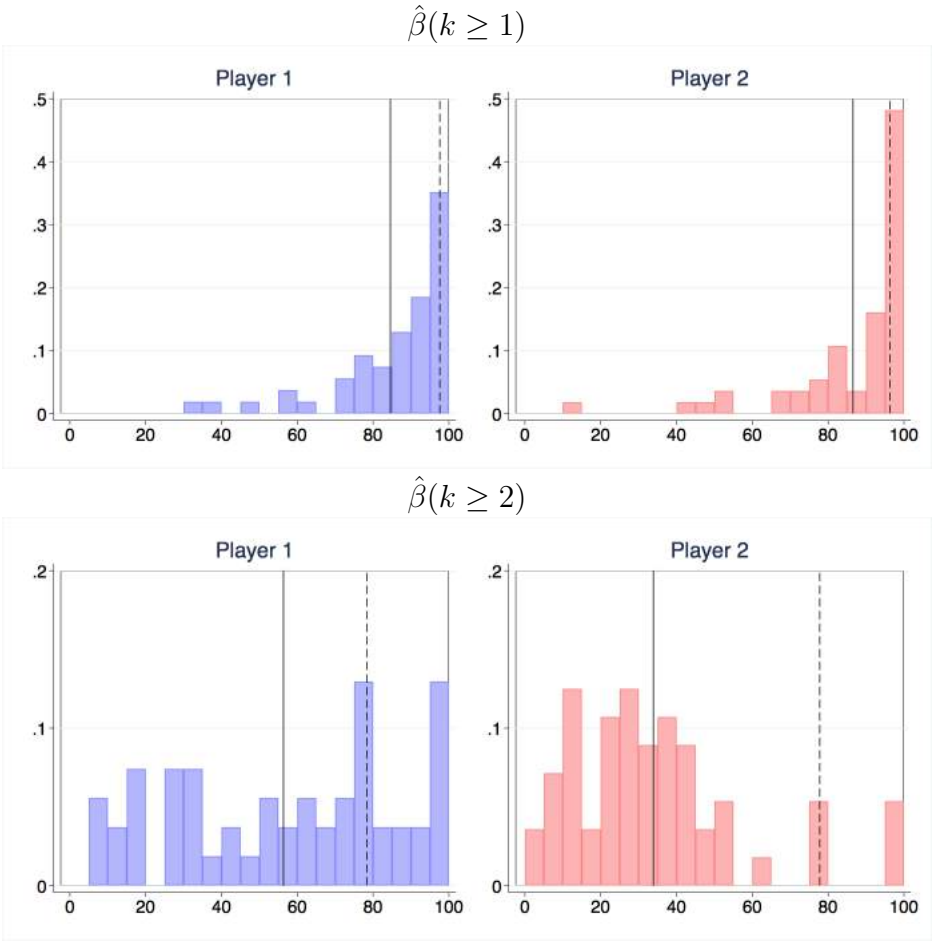

Figure 13: *Sophistication by player.* The top panel gives histograms of $\hat{\beta}(k \geq 1)$, i's belief that j best responds to his dominant action in Dj, across subjects. The bottom panel gives histograms of $\hat{\beta}(k \geq 2)$, i's belief that j best responds to i's dominant action in Di (as opposed to the a uniform distribution), across subjects. The solid lines mark i's average beliefs, and the dashed lines mark j's corresponding action frequencies from [A, ∘].

all subjects observed exactly the same games throughout the experiment and were randomly assigned to their roles. Thus, the difference in measured sophistication must be caused by their experience in different roles of the X-games.

The frequency of actions taken in the dominance solvable games are nearly identical across player 1- and player 2-subjects in the first stage. That stated beliefs (and to some extent actions) differ across player roles in the second stage suggests that role-dependent no-feedback learning took place. Interestingly, however, there is no evidence of learning throughout the experiment in the sense that, within each stage of the experiment, there is no within-player trend in actions or beliefs across multiple rounds of the same game. Hence, we believe that there was some no-feedback "belief learning" in the first stage that did not manifest in actions. In the second stage, player 1-subjects' stated beliefs already indicated higher levels of sophistication in the very first instance of Di, so we believe all of the learning had already taken place by that point.

6.2.2 Response times

For additional support, we consider response times. If the player 1-role induces greater strategic sophistication, we would expect for player 1 to also take longer to form beliefs since they go farther in terms of strategic reasoning.

In Appendix Figure 30, we plot the average time to finalize belief statements for each game and player role. Player 1 takes longer on average for all games. That player 1-subjects take longer on games Di and Dj is further suggestion that their experience in the role of player 1 in the X-games spills over to these new environments even though they are the same for both players.

6.2.3 The relationship between sophistication and behavior

If differential sophistication across player roles is to explain behavior in the X-games, sophistication measured in dominance solvable games should be predictive of behavior in the X-games *within*-role. We establish this before formally modeling the relationship between sophistication and beliefs in the next subsection.

To this end, we divide player 1-subjects into low and high sophistication groups based on having values of $\hat{\beta}(k \geq 2)$ below and above the player 1-median, and similarly for player 2-subjects. We then compare behaviors across the sophistication groups, focusing on beliefs data and first-stage actions data [A,BA] (results using the second-stage actions are similar). Appendix Table 20 summarizes our results, which we discuss in detail below. In each column,

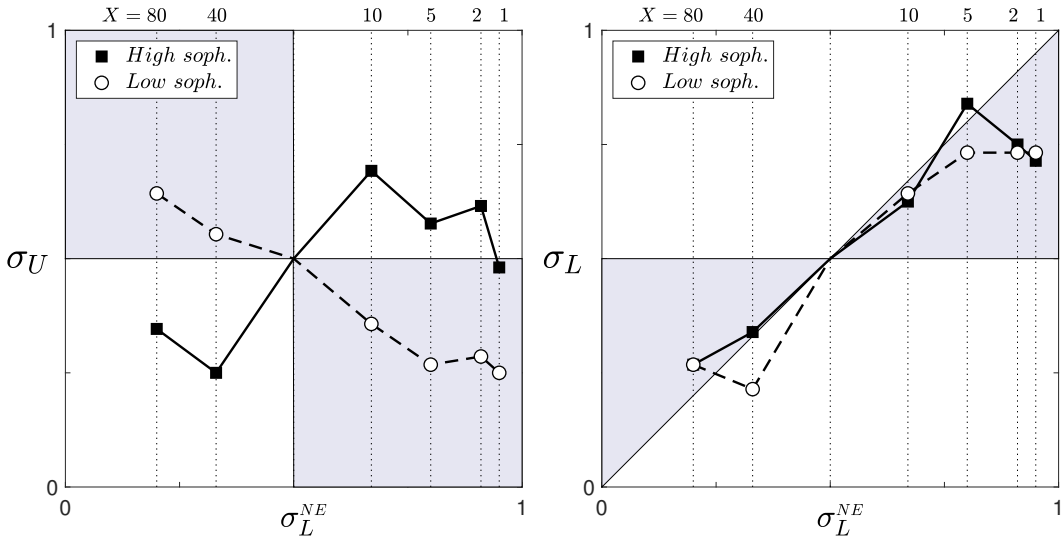

Figure 14: *Actions data by sophistication group.* For each player, we plot the empirical action frequencies [A,BA] for high- and low-sophistication groups. Player 1 is on the left, and player 2 is on the right.

we regress beliefs or actions on indicators for each of the six X-games (omitted from the table) and indicators for each of the six games interacted with an indicator for the high sophistication group. The results are robust to alternate groupings[34] and using $\hat{\beta}(k \geq 2)$ as a continuous variable.

Compared to less sophisticated player 1s, more sophisticated player 1s tend to believe that player 2 plays L less often for $X > 20$ (more often for $X < 20$). Compared to less sophisticated player 2s, more sophisticated player 2s tend to believe that player 1 plays U less often for $X > 20$ (more often for $X < 20$). In Appendix Figure 31, we plot histograms of beliefs by sophistication group for both players and all games. This shows that sophisticated player 1s have more extreme beliefs while for player 2, it is the opposite. Hence, low sophistication does not simply proxy for more conservative beliefs.

In Figure 14, we plot the empirical action frequencies from [A,BA]. For player 1, the difference between high and low sophistication groups is quantitatively very large, highly significant (column 2 of Appendix Table 20), and qualitatively surprising. Consistent with the differences in their beliefs, the low sophistication group tends to take U for $X > 20$ and D for $X < 20$, while for the high sophistication group, it is the opposite. Interestingly, the low sophistication group is very consistent with the joint QRE-NBE predictions and the high

[34] We tried terciles and quartiles as well as using the median of $\hat{\beta}(k \geq 2)$ across all subjects for both players instead of player-specific medians.

sophistication group looks bizarre from the perspective of this theory. On average, however, the behavior is not far from Nash (left panel of Figure 3). For player 2, the two sophistication groups have very similar actions data which cannot be distinguished statistically (column 4 of Appendix Table 20).

6.2.4 Beliefs and sophistication in the X-games

Our analysis shows that player 1-subjects are more sophisticated than player 2-subjects in the dominance solvable games and that beliefs in the dominance solvable games predict behavior in the X-games. However, this does not imply in of itself that player 1-subjects form more sophisticated beliefs in the X-games. To determine if this is the case, we introduce a simple model of belief formation that provides a formal link between beliefs in the X-games and sophistication. The goal is not to propose a general theory of belief formation, but to infer sophistication from the X-game data to test the hypothesis that player 1 is more sophisticated than player 2 and determine if this can generate the biases we observe. As such, the model is highly specialized to the X-games.

		k							
		0	1	2	3	4	5	6	...
$X > 20$	Player 1 (σ_U^k)	$\frac{1}{2}$	1	1	0	0	1	1	...
	Player 2 (σ_L^k)	$\frac{1}{2}$	$\frac{1}{2}$	0	0	1	1	0	...
$X < 20$	Player 1 (σ_U^k)	$\frac{1}{2}$	0	0	1	1	0	0	...
	Player 2 (σ_L^k)	$\frac{1}{2}$	$\frac{1}{2}$	1	1	0	0	1	...

Table 10: *Levels in game X.*

To this end, we will use a modified cognitive hierarchy framework (Camerer et al. [2004b]) in which each subject believes she faces opponents drawn from a distribution of level types (e.g. Stahl and Wilson [1995] and Nagel [1995]). Level 0 is assumed to uniformly randomize, level 1 best responds to level 0, and so on, with level k best responding to level $k-1$. In forming beliefs, subjects believe they face a distribution of level types. We say that a subject is *sophisticated* if her beliefs imply that she believes she faces types with high levels.

In Table 10, we write out the actions taken by different level k-types in the X-games, written as σ_U^k and σ_L^k. In the case of indifference we assume uniform tie-breaking. We make two observations. First, the levels "cycle" in the sense that $\sigma_U^k = \sigma_U^{k-1}$ for even k and $\sigma_L^k = \sigma_L^{k-1}$ for odd k. Second, X matters only insofar as it is greater than or less than 20.

The fact that levels cycle causes problems for identification: any action can be interpreted as coming from an arbitrarily high level and therefore any belief can be rationalized as being

arbitrarily sophisticated. We restore identification by truncation.

Camerer et al. [2004b] find that an average of 1.5 steps of reasoning fits the data from many games, so we assume player i forms beliefs according to the following three step procedure. First, player i imagines what j would do naively, which we assume is a best response to uniform play. Second, player i imagines her own best response to that action. Third, player i imagines j's best response to that action. During this process, player i imagines player j taking the level 1 and level 3 actions, so we assume player i believes she faces a fraction $(1-\alpha)$ of level 1s and a fraction α of level 3s. Player i's belief in game X is thus given by

$$\bar{\sigma}_j^X(\alpha) = (1-\alpha)\cdot\sigma_j^{1,X} + \alpha\cdot\sigma_j^{3,X}, \tag{6}$$

where α is a free parameter that we interpret as sophistication as it is the belief that the opponent is of a high level. Our decision to begin the process of introspection at level 1, as opposed to level 0, is motivated by the fact that subjects overwhelmingly expect their opponents to take the dominant action in Dj (top panel of Figure 13), which is unsurprising in simple games. That the introspection process skips level 2 is a consequence of iterated best response in asymmetric games, not an arbitrary restriction.

The fact that $\sigma_j^{k,X}$ only depends on X insofar as X is greater or less than 20 means that beliefs formed as in (6) will also have this property. However, this is counterfactual: the analysis of Section 5.3 shows that beliefs change systematically across all values of X. For this reason, we generalize level k to allow for each level type to make payoff sensitive errors.[35] Player i believes player j of level k faces a vector of payoffs $v_j^k = (v_{j1}^k, v_{j2}^k)$ and takes an action according to

$$Q_j^\mu(v_j^k) = \frac{(v_{j1}^k)^{1/\mu}}{(v_{j1}^k)^{1/\mu} + (v_{j2}^k)^{1/\mu}}, \tag{7}$$

where parameter $\mu > 0$ controls the sensitivity to payoff differences. The action of player j of level k is thus defined recursively according to $\sigma_j^{k,X}(\mu) \equiv Q_j^\mu(\bar{u}_j(\sigma_i^{k-1}(\mu;X)))$ and $\sigma_i^0 = \sigma_j^0 = \frac{1}{2}$. We replace $\sigma_j^{k,X}$ in (6) with $\sigma_j^{k,X}(\mu)$ which yields

$$\bar{\sigma}_j^X(\alpha,\mu) = (1-\alpha)\cdot\sigma_j^{1,X}(\mu) + \alpha\cdot\sigma_j^{3,X}(\mu). \tag{8}$$

We still interpret α as sophistication and μ is the payoff sensitivity player i attributes to j. We favor the Luce form of quantal response (7) because it is scale invariant, which implies

[35] The idea of injecting noise into the description of levels is not new. See, for example, Capra [1999], Weizsacker [2003], and Rogers et al. [2009].

beliefs described by (8) are symmetric in σ_L^{NE}–a feature we will show matches the data.

To gain some intuition for the types of beliefs implied by (8), we first consider the cases of $\mu = 0$ and $\mu = 1$. When $\mu = 0$, player i believes each level-type of player j best responds to their beliefs: (8) collapses to (6) and hence X only matters insofar as X is greater than or less than 20 or equivalently if σ_L^{NE} is less than or greater than $\frac{1}{2}$. Beliefs thus follow a step pattern:

$$\bar{\sigma}_U^X(\alpha, 0) = \begin{cases} (1-\alpha) & \sigma_L^{NE,X} < \frac{1}{2} \\ \alpha & \sigma_L^{NE,X} > \frac{1}{2} \end{cases}$$

$$\bar{\sigma}_L^X(\alpha, 0) = \begin{cases} \frac{1}{2}(1-\alpha) & \sigma_L^{NE,X} < \frac{1}{2} \\ \frac{1}{2}(1+\alpha) & \sigma_L^{NE,X} > \frac{1}{2} \end{cases}.$$

For $\mu > 0$, player i believes each level-type of player j makes payoff sensitive errors in best responding to her beliefs. Hence, i's beliefs are sensitive to all changes in X and therefore also to changes in σ_L^{NE}. When $\mu = 1$, properties of (7) imply that i's beliefs are linear in σ_L^{NE}:

$$\bar{\sigma}_U^X(\alpha, 1) = (1-\alpha) \cdot (1 - \sigma_L^{NE,X}) + \alpha \cdot \frac{1}{2}$$

$$\bar{\sigma}_L^X(\alpha, 1) = (1-\alpha) \cdot \frac{1}{2} + \alpha \cdot \sigma_L^{NE,X}.$$

Thus, when beliefs are viewed as functions of σ_L^{NE}, identification of sophistication is based on levels for $\mu = 0$ and based on the slope for $\mu = 1$. Note also that, in the $\mu = 1$ case, beliefs coincide with Nash equilibrium when players are fully sophisticated ($\alpha = 1$). In general, it is easy to generate parameter values for which the beliefs fall in the interior of the QRE-NBE region.

6.2.5 Structural model

We adapt the model of the previous section to be fit to individual subjects' belief data from the X-games. By fitting the model to each subject's data, we infer a measure of strategic sophistication for each subject.

We recast the belief $\bar{\sigma}_j^X$ defined in (8) as the central tendency of beliefs and assume that beliefs are noisy with a parametric error structure. For player i with parameters α and μ,

we assume belief $b \in \{0, 1, ..., 100\}$ is drawn in game X according to

$$p_X(b; \alpha, \mu, \lambda)) = \frac{e^{-\lambda(b-100\cdot\bar{\sigma}_j^X(\alpha,\mu))^2}}{\sum_{b'\in\{0,1,...,100\}} e^{-\lambda(b'-100\cdot\bar{\sigma}_j^X(\alpha,\mu))^2}},$$

so that the belief closest to $100 \cdot \bar{\sigma}_j$ is the mode and $\lambda > 0$ is a precision parameter.

The data of subject s in role i is a set of 30 belief statements $\{\hat{b}_{sl}^{iX}\}_{lX}$ where $l \in \{1, ..., 5\}$ indexes each elicitation and X indexes the game. For each subject, we choose α, μ, and λ to maximize the log-likelihood of stated beliefs:

$$L^s(\hat{b}; \alpha, \mu, \lambda) = \sum_X \sum_{l=1}^{5} ln(p_X(\hat{b}_{sl}^{iX}; \alpha, \mu, \lambda)).$$

We find that for 50 out of 110 subjects (45%), a likelihood ratio test rejects the restriction $\mu = 0$ at the 5% level, meaning that a substantial fraction of subjects are significantly sensitive to variations in X in ways allowed by the model.

Figure 15 plots the beliefs of some representative subjects, superimposed with the best-fit $\bar{\sigma}_j^X(\hat{\alpha}, \hat{\mu})$ for $X \in (0, \infty)$. The top panels are for player 2-subjects forming beliefs about player 1, and the bottom panels are for player 1-subjects forming beliefs about player 2. Clearly, there is considerable heterogeneity across subjects, but the model is flexible enough to accommodate their diverse belief patterns.

The top panel of Figure 15 features player 2-subjects who, from left to right, are increasing in inferred sophistication $\hat{\alpha}$. Subject 105 believes, overwhelmingly, that player 1 will take U when X is large and take D when X is small, or in other words to engage in level 1 behavior. For this reason, the model infers the low level of sophistication $\hat{\alpha} = 0.04$. At the other extreme, subject 104 believes that player 1 will mostly take D when X is large and take U when X is small, or in other words to engage in level 3 behavior, and so the model infers a high level of sophistication $\hat{\alpha} = 0.76$. Interestingly, it is the very sophisticated subjects whose beliefs systematically fall outside of the QRE-NBE region.

The bottom panel of Figure 15 features player 1-subjects who, from left to right, are increasing in inferred sophistication $\hat{\alpha}$. Subject 26 believes that player 2 will tend to take R when X is large and take L when X is small, consistent with a mix of level 1 and level 3 behavior. The same can be said of Subject 5, however, the model infers that Subject 5 is much more sophisticated than Subject 26, with values of $\hat{\alpha} = 1$ and $\hat{\alpha} = 0.47$, respectively. Inspecting their beliefs more closely, we observe that Subjects 26 and 5 have estimated

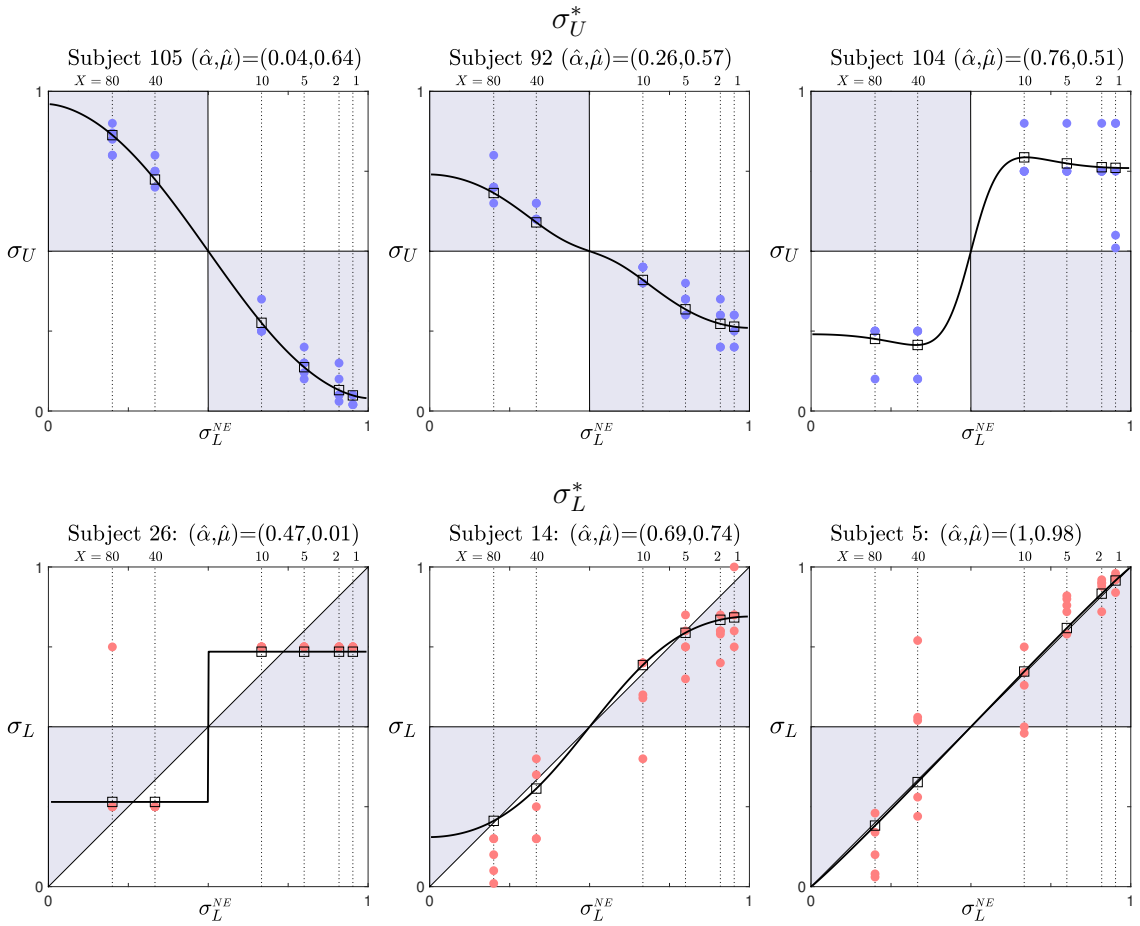

Figure 15: *Individual subjects' beliefs.* We plot representative individual subjects' stated beliefs, superimposed with the best fit $\bar{\sigma}_j^X(\hat{\alpha},\hat{\mu})$. The top row is for player 2-subjects forming beliefs about player 1, and the bottom row is for player 1-subjects forming beliefs about player 2.

sensitivities of $\hat{\mu} \approx 0$ and $\hat{\mu} \approx 1$, respectively, corresponding to step-like and linear belief patterns. Hence, while the two subjects may have similar beliefs when averaged across games $X > 20$ (and similar beliefs when averaged across games $X < 20$), Subject 5 believes in a much higher fraction of level 3 opponents, albeit much noisier ones.

6.2.6 Player 1 is more sophisticated than player 2 in the X-games

We show that the structural model applied to the X-games implies that player 1-subjects tend to be much more sophisticated by $\hat{\alpha}$ than player 2-subjects. This can be seen from Figure 16, which replicates the bottom panel of Figure 13 by plotting histograms of inferred sophistication $\hat{\alpha}$ for player 1- and player 2-subjects.

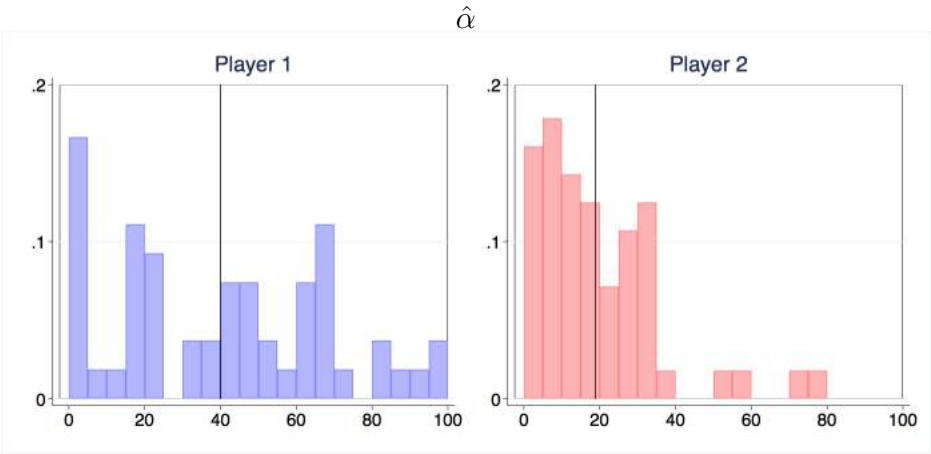

Figure 16: *Inferred sophistication by player.*

To further validate this finding, we show that $\hat{\alpha}$ estimated from the fully mixed X-games and $\hat{\beta}(k \geq 2)$ directly measured in dominance solvable Di are strongly correlated. Applying the structural model to Di, α and μ cannot be separately identified in the sense that $\beta(k \geq 2)$ may be consistent with different (α, μ)-pairs.[36] However, for any fixed μ, $\beta(k \geq 2)$ is an increasing function of α. Hence, if μ is sufficiently uncorrelated with α, $\hat{\beta}(k \geq 2)$ is predicted to correlate with $\hat{\alpha}$. We are therefore justified in comparing $\hat{\alpha}$ and $\hat{\beta}(k \geq 2)$ to validate the structural model.

Figure 17 gives a scatter plot of inferred versus directly measured sophistication. We find there is a strong positive correlation for player 1-subjects, player 2-subjects, and all subjects; and this is confirmed in Table 11 which presents the correlations.

[36] This is not an issue in the X-games because of variations in X.

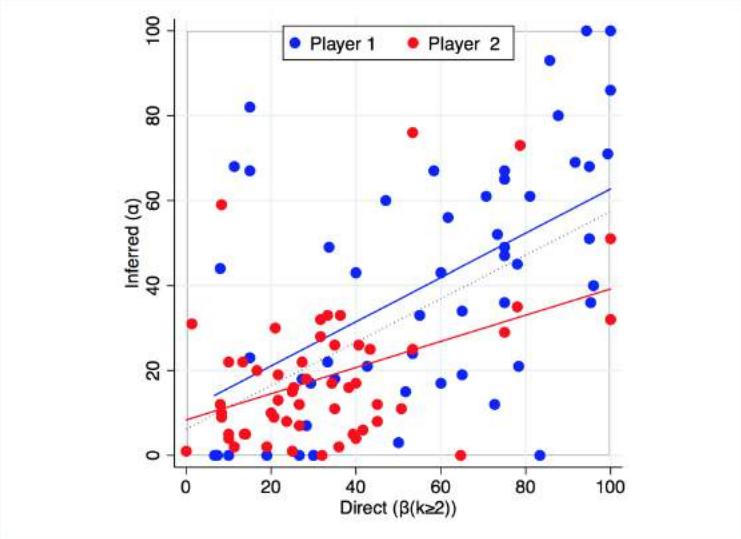

Figure 17: *Inferred versus directly measured sophistication.* We give a scatterplot of subjects' $\hat{\alpha}$ versus $\hat{\beta}(k \geq 2)$, with best-fit lines for player 1-subjects, player 2-subjects, and all subjects.

		Player 1	Player 2	Both
α, μ	Pearson	0.53^{***}	0.44^{***}	0.57^{***}
	spearman	0.54^{***}	0.34^{***}	0.54^{***}
$\alpha, \mu = 0$	pearson	0.48^{***}	0.41^{***}	0.50^{***}
	spearman	0.50^{***}	0.35^{***}	0.46^{***}
$\alpha, \mu = 1$	pearson	0.52^{***}	0.41^{***}	0.57^{***}
	spearman	0.54^{***}	0.34^{***}	0.55^{***}

Table 11: *Correlation between inferred and directly measured sophistication*

For robustness, in addition to the 2-parameter (α, μ) model, we also consider the 1-parameter restrictions $(\alpha, \mu = 0)$ and $(\alpha, \mu = 1)$ discussed in the previous section. We also report both Pearson (linear) coefficients and spearman (rank-based) coefficients. We find that the correlations are in all cases highly significant, with large magnitudes ranging from 0.34 to 0.63. The correlations are a bit higher for player 1 than for player 2, and highest for both players pooled together. The 2-parameter model and restricted $\mu = 1$ model lead to very similar correlations, which are slightly higher than the correlations implied by the $\mu = 0$ restricted model.

Taken together, the results of this section provide support for the sophistication hypothesis. The structural model is qualitatively consistent with the patterns observed in subjects' beliefs data, it captures the stylized fact of the sophistication gap, and it implies degrees of sophistication that correlate strongly with the direct measures.

6.2.7 Discussion

Our results indicate that player 1 forms more sophisticated beliefs than player 2 in the X-games. Since the subjects in the two roles were ex-ante identical, this suggests a model of *endogenous* role-dependent sophistication. It is beyond the scope of this paper, but we consider eveloping such models collecting datasets to differentiate between them is a promising direction for future research.

It seems difficult to reconcile our data with a rational, optimizing model of endogenous sophistication (e.g. the model of Alaoui and Penta [2015]) for the reason that one player or the other faces much higher average payoffs depending on X, and yet it is always player 1 who forms more sophisticated beliefs. We believe psychological explanations related to the salience of player 1's payoffs are more promising.

6.3 Modeling actions and beliefs jointly

In Sections 6.1 and 6.2, we offered explanations for the failures of *monotonicity* and *unbiasedness*, respectively. These explanations came with structural models that were fit to actions given beliefs and then to beliefs, respectively. In this section, we combine the previously introduced elements to maximize the likelihood of actions and beliefs jointly, which we show can rationalize the whole of the data, including the belief biases we observe.

The data of subject s in role i is a set of 30 action-belief pairs $\{\hat{a}_{sl}^{iX}, \hat{b}_{sl}^{iX}\}_{lX}$ where $l \in \{1, ..., 5\}$ indexes each elicitation and X indexes the game. For each player 1-subject s, we choose $\rho, \mu_a, \alpha, \mu, \lambda$ to maximize

$$\begin{aligned} L^s(\hat{a}, \hat{b}; \rho, \mu_a, \alpha, \mu, \lambda) &= \sum_X \sum_{l=1}^{5} ln(p_X(\hat{a}_{sl}^{iX} | \hat{b}_{sl}^{iX}; \rho, \mu_a) p_X(\hat{b}_{sl}^{iX}; \alpha, \mu, \lambda)) \\ &= \sum_X \sum_{l=1}^{5} ln(p_X(\hat{a}_{sl}^{iX} | \hat{b}_{sl}^{iX}; \rho, \mu_a)) + \sum_X \sum_{l=1}^{5} ln(p_X(\hat{b}_{sl}^{iX}; \alpha, \mu, \lambda)) \\ &= L^s(\hat{a}|\hat{b}; \rho, \mu_a) + L^s(\hat{b}; \alpha, \mu, \lambda), \end{aligned}$$

where $L^s(\hat{a}|\hat{b}; \cdot)$ and $L^s(\hat{b}; \cdot)$ are as before. Hence, we find the same parameter estimates as before for each subject. For player 2-subjects, we fit the same model, except under the assumption of linear utility $\rho = 0$ as curvature cannot be identified due to the symmetry of player 2's payoffs (see Section 6.1).

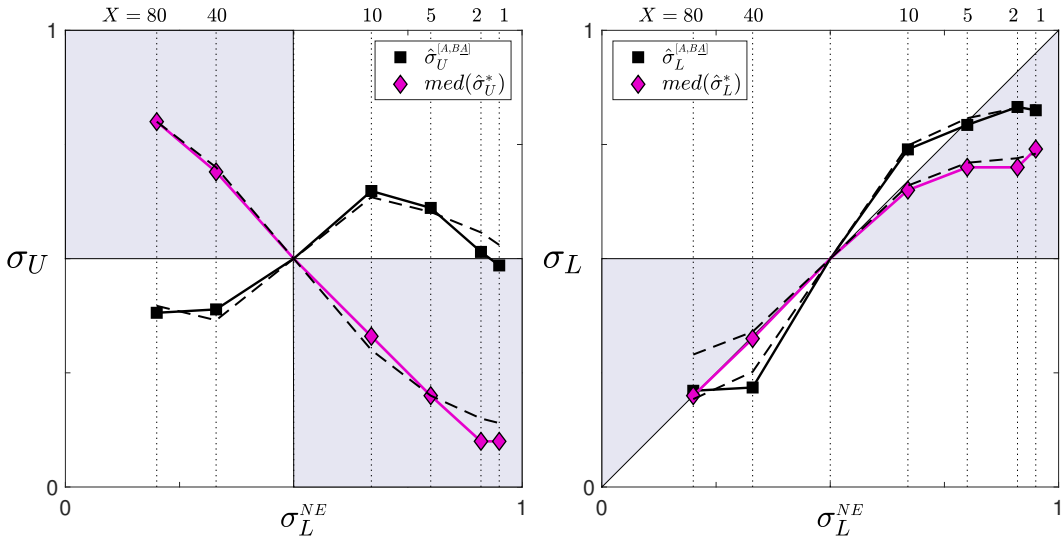

Figure 18: *The fit of the structural model fit to actions and beliefs jointly.*

After fitting the model to each subject, we simulate the aggregate data. In Figure 18, we plot the simulated empirical frequencies and median beliefs, which we compare to the data from [A,BA] to which the model was fit. We find that the model generates the observed belief biases, and we already know from Section 6.2.6 that the fitted model implies much higher levels of sophistication for player 1-subjects.

7 Action-noise or belief-noise?

Clearly, there is considerable noise in both actions and beliefs. What is lost in ignoring either source of noise?

In this section, we explore this question via a counterfactual exercise. Specifically, we construct two counterfactual action frequencies that result from "turning off" just one source of noise in the second-stage data [A,BA] (for which we can associate actions with beliefs). $\hat{\sigma}_i^{best\ response}$ is what we would observe if subjects best responded to every stated belief, and $\hat{\sigma}_i^{correct\ beliefs}$ is what we would observe if subjects had correct beliefs (over first-stage actions [A, $\circ$]). The former can be constructed directly from the data. To construct the latter, we set i's beliefs equal to j's empirical action frequency and assume actions are governed by the random utility model with curvature fitted to each subject's data from Section 6.1.

In Figure 19, we plot $\hat{\sigma}_i^{[A,BA]}$, $\hat{\sigma}_i^{best\ response}$, and $\hat{\sigma}_i^{correct\ beliefs}$ for both players. As a measure

of the performance of each counterfactual, we consider the average absolute differences between actual and counterfactual frequencies across the games for each player, $D_i^{best\ response} \equiv \frac{1}{6}\sum_X |\hat{\sigma}_i^{[\mathrm{A,B\underline{A}}],X} - \hat{\sigma}_i^{best\ response,X}|$ and $D_i^{correct\ beliefs} \equiv \frac{1}{6}\sum_X |\hat{\sigma}_i^{[\mathrm{A,B\underline{A}}],X} - \hat{\sigma}_i^{correct\ beliefs,X}|$. These represent the prediction errors or loss in ignoring action-noise and belief-noise, respectively, and are displayed in Table 12.

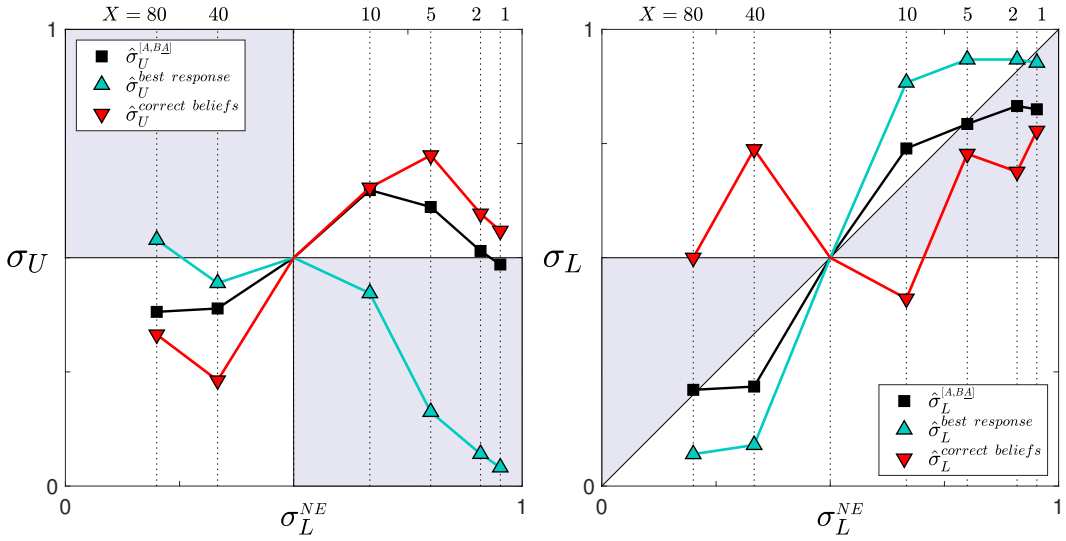

Figure 19: *Ignoring action-noise and belief-noise–a counterfactual.*

We see that, for both players, the two counterfactuals are fairly *in*accurate, indicating that action-noise and belief-noise are both important ingredients for explaining behavior. Models that ignore any one–as indeed the large majority of models applied to experimental data do–may suffer from misspecification. Interestingly, which source of noise is more important depends on player role. For player 1, $\hat{\sigma}_i^{correct\ beliefs}$ is much more accurate than $\hat{\sigma}_i^{best\ response}$, whereas for player 2, it is the opposite. This is intuitive as player 1's stated beliefs are fairly accurate but she faces a difficult decision for any given belief, and it is the opposite for player 2.

	$D_i^{best\ response}$	$D_i^{correct\ beliefs}$
player 1 $(i = 1)$	0.30	0.16
player 2 $(i = 2)$	0.13	0.23

Table 12: *Loss in ignoring action-noise and belief-noise.*

8 Issues of belief elicitation: a discussion

Our analysis depends on making beliefs observable through direct elicitation, so we discuss two well-known and potentially confounding issues of belief elicitation. First, it may be that stated beliefs are only noisy signals of subjects' underlying latent or "true" beliefs. Second, belief elicitation itself may affect the actions subjects take. We argue that these issues do not affect our main conclusions.

8.1 Stated beliefs as noisy signals of true beliefs

Throughout the paper, we have implicitly assumed that stated beliefs equal the latent or "true" beliefs that subjects hold in their minds and guide their actions. More generally, it may be that stated beliefs are noisy signals of the underlying true beliefs due to errors in reporting or noisy introspection about one's beliefs (see Rutstrom and Wilcox [2009] for a discussion). In that case, can we still say that the unobserved true beliefs are noisy? Can we reject the same axioms with respect to the true beliefs? We argue that the answer to both questions is *yes*.

We suppose that, for a given game, b_s^* and b_0^* are stated and true beliefs, respectively. These are (possibly degenerate) random variables whose support is contained in [0,1]. Let b_0 be a realization of true beliefs, and let $b_s^*(b_0)$ be the random stated beliefs conditional on b_0. We assume that actions depend on true belief realizations through the function $Q_i(\bar{u}_i(b_0))$.

Are true beliefs noisy? If within-subject-game, the true belief were fixed and stated beliefs were simply noisy signals of the underlying belief, then within-subject-game variation in stated beliefs would not be predictive of actions. If this were the case, we would see coefficients of 0 in Table 7, but this is strongly rejected. Hence, we conclude that true beliefs are noisy.

As we found using stated beliefs, are *monotonicity* and *unbiasedness* also rejected with respect to true beliefs? To answer this, we require additional structure. To this end, assume that stated beliefs are drawn from a distribution that is centered, in the sense of median, around the true belief realization: $med[b_s^*(b_0)] = b_0$ for all b_0 (and $b_s^*(b_0) = b_0$ w.p. 1 if $b_0 \in \{0, 100\}$). Under this assumption, we argue that it is very unlikely that either axiom holds in true beliefs given our data.

Consider player 1 in game $X5$ (see Figure 6). The indifferent belief is 80, and the *monotonicity* violation occurs in the interval of stated beliefs $[60, 80]$. Suppose that actions given true beliefs are governed by $Q_i(\bar{u}_i(b_0)) = \frac{1}{2}$ for $b_0 \leq 80$ and $Q_i(\bar{u}_i(b_0)) = 1$ for $b_0 >$

80, which is the *monotonic* quantal response function most likely to generate the observed violation. Under the assumption that $med[b_s^*(b_0)] = b_0$ for all b_0, the expected mass of stated beliefs in $[60, 80]$ that is associated with a true belief $b_0 > 80$, and thus the action $Q_i = 1$, is *at most* equal to the mass of stated beliefs greater than 80. It is clear that this is insufficient to rationalize the violation we observe. Hence, the underlying Q_i defined over true beliefs cannot be *monotonic*.

We found that player 2 forms very biased stated beliefs over player 1's actions. For instance, in $X80$, $med(b_s^*) > \hat{\sigma}_U$ (see top left panel of Figure 9). Suppose that, in true beliefs, $med(b_0^*) = \hat{\sigma}_U$. This does not imply that $med(b_s^*) = \hat{\sigma}_U$, but it does imply that $\mathbb{P}(b_s^* > \hat{\sigma}_U) \leq \frac{3}{4}$,[37] and we observe that $\hat{\mathbb{P}}(b_s^* > \hat{\sigma}_U)$ is much greater than three-fourths in the data. Hence, the underlying true beliefs cannot be *unbiased*.

8.2 The effects of belief elicitation

There is little consensus on if, how, and under what conditions belief elicitation has an effect on the actions subjects take. In their recent review articles on belief elicitation, Schlag et al. [2015] describes the evidence as "scanty and contradictory" whereas Schotter and Trevino [2014] state that the "evidence presents a more consistent picture in favor of the idea that belief elicitation is innocuous". We are unaware of studies that elicit beliefs for asymmetric matching pennies without the influence of feedback, and, for the studies that have used feedback, the documented effects have been small.[38]

It has been conjectured that belief elicitation may increase strategic sophistication (see the discussion in Schlag et al. [2015]), but to the best of our knowledge, no previously documented effects can readily be interpreted in this way.[39]

In Appendix 11.4, we show that there are small, but systematic and statistically significant, differences between the first- and second-stage action frequencies in [A,BA]. We find no such differences between the two stages of the [A,A] treatment that did not involve belief

[37] That $med(b_0^*) = \hat{\sigma}_U$ implies that $\mathbb{P}(b_0^* > \hat{\sigma}_U) = \mathbb{P}(b_0^* < \hat{\sigma}_U) = \frac{1}{2}$. Given that $med[b_s^*(b_0)] = b_0$ for all b_0, $\mathbb{P}(b_s^* > \hat{\sigma}_U)$ is maximized when $\mathbb{P}(b_s^*(b_0) > \hat{\sigma}_U | b_0 > \hat{\sigma}_U) = 1$ and $\mathbb{P}(b_s^*(b_0) > \hat{\sigma}_U | b_0 < \hat{\sigma}_U) = \frac{1}{2}$, which implies that $\mathbb{P}(b_s^* > \hat{\sigma}_U) = \frac{3}{4}$.

[38] Nyarko and Schotter [2002] find no effect. Rutstrom and Wilcox [2009] find an effect for only one player and only during early rounds. They claim to be the first to find any such effect in games with unique, mixed strategy Nash equilibria, and we are unaware of any studies to do so since.

[39] Costa-Gomes and Weizsacker [2008] find, in the context of 3×3 dominance solvable games, that beliefs seem more sophisticated than the corresponding actions, but they also find that the belief elicitation has no effect on actions. Schotter and Trevino [2014] suggest that belief elicitation may hasten convergence to equilibrium in games played with feedback, but this is distinct from sophistication.

elicitation, and so we conclude that there *is* a belief elicitation effect. In Appendix 11.4, we provide additional discussion and argue that this does not affect our main conclusions.

9 Relationship to the existing literature

A central goal of behavioral game theory (Camerer [2003]) is to describe how real people play games. This paper contributes to the large sub-literature that focuses on bounded rationality as drivers of behavior (as opposed to, for example, other-regarding preferences (Fehr and Schmidt [1999])).

We fit most squarely in the literature on stochastic equilibrium models that maintain fixed-point consistency between players' actions but allow for random elements. The prominent example is quantal response equilibrium (QRE), a concept that allows for "noise in actions" but maintains that beliefs are correct. Early QRE theory was developed in a series of papers (McKelvey and Palfrey [1995], McKelvey and Palfrey [1998], Chen et al. [1997], and others) and is surveyed in a recent monograph (Goeree et al. [2016]).

Many papers acknowledge that the assumption of correct beliefs is unrealistic, but the large majority of these papers applied to experimental data are non-equilibrium models such as level k (e.g. Nagel [1995] and Stahl and Wilson [1995]; and reviewed in Crawford et al. [2013]) and its many successors (e.g. Camerer et al. [2004b], Alaoui and Penta [2015], and Goeree and Holt [2004]). These models have proven extremely useful in explaining experimental data post-hoc, but their application is sometimes criticized for lacking the discipline that equilibrium consistency brings.

There are many equilibrium models that involve biased or otherwise incorrect beliefs (e.g. Geanakoplos et al. [1989] and Heller and Winter [2018]), but these are typically ill-suited for (nor were they designed for) direct application to experimental data. In terms of models that allow for "noise in beliefs", there are very few. An early example is the parametric sampling equilibrium (Rubinstein and Osborne [2003]) which has been applied to experimental data (Selten and Chmura [2008]). Notably, Friedman and Mezzetti [2005] introduce the notion of a belief-map as part of their random belief equilibrium (RBE). Their focus is on the limiting case in which belief-noise "goes to zero" to develop a theory of equilibrium selection, so their conditions on belief distributions do not impose any testable restrictions beyond ruling out weakly dominated actions. Noisy belief equilibrium (NBE) (Friedman [2019]), on which our paper builds, was developed to study the case in which belief noise is bounded away from zero, so it maintains the structure of the belief-map but imposes behavioral axioms to impose

testable restrictions.

We are novel in studying an equilibrium model that allows for both noise in actions and noise in beliefs. Because it is a non-parametric, axiomatic model that makes set-predictions and we study the restrictions imposed by the axioms, there is a clear relationship to the literature on the empirical content of QRE (e.g. Haile et al. [2008], Goeree et al. [2005], Melo et al. [2017], Goeree and Louis [2018], and Goeree et al. [2018]).[40]

We also make several contributions to the literatures on belief elicitation and strategic sophistication. By having multiple belief elicitations per subject-game without feedback, we are able to study noise in beliefs. By eliciting beliefs for a family of closely related games, we are able to track how beliefs vary within-subject across games and compare these belief patterns across individual subjects. This distinguishes us from experiments that elicit beliefs once for each game in a set of seemingly unrelated games (e.g. Costa-Gomes and Weizsacker [2008] and Rey-Biel [2009]) as well as studies that elicit beliefs for the same game repeatedly with feedback (e.g. Nyarko and Schotter [2002] and Rutstrom and Wilcox [2009]). In terms of analysis, we focus not just on rates of best response, but also on how these rates of best response vary across every neighborhood of stated beliefs.[41] In addition to establishing that subjects' beliefs are noisy, we show that within-subject variations in beliefs predict actions.

Studies on strategic sophistication, primarily using the level k framework, typically use dominance solvable games to get around the non-identifiability issues discussed in Section 6.2.4 (e.g. Costa-Gomes and Crawford [2006] and Kneeland [2015]). Hence, little is known about how sophistication manifests itself in the important class of fully mixed games. By combining our rich subject-level beliefs data with a direct measure of sophistication, we provide some of the first evidence. In finding a correlation between sophistication measures from fully mixed and dominance solvable games, our analysis also suggests a "persistence of strategic sophistication" across these two domains, adding positive evidence to a literature that has found largely negative results (e.g. Georganas et al. [2015]).

[40] This literature was jumpstarted with the negative results of Haile et al. [2008] who showed that *structural* QRE can rationalize the data from any one game without strong restrictions on the error distributions. Work since has focused on studying the restrictions imposed by other variants of QRE.

[41] Hyndman et al. [2013] provides some evidence in 3×3 games that beliefs toward the corners of the simplex are best responded to more often.

10 Conclusion

Motivated to contribute to a more realistic game theory, we study the beliefs people form over opponents' behavior and the actions they take conditional on these beliefs.

We begin by characterizing an equilibrium model with "noise in actions" and "noise in beliefs"–a benchmark model that avoids unrealistic deterministic assumptions that would be trivially rejected. By injecting noise into both actions and beliefs, the model runs the risk of becoming vacuous, so we restrict both types of noise to satisfy axioms that are stochastic generalizations of "best response" and "correct beliefs".

Using a laboratory experiment, we collect actions data and elicit beliefs for a canonical family of games with systematically varied payoffs. By having multiple elicitations per subject-game without feedback, our design allows us to *(i)* observe noise in both actions and beliefs and *(ii)* test the axioms of the benchmark model.

We find that both sources of noise are important for explaining features of the data, which suggests that deterministic assumptions may be an important source of misspecification. In particular, this calls into question the common practice of applying models with deterministic beliefs to experimental data.

Interestingly, despite the axioms being relatively weak, we find rejections. The most striking violation comes in the form of belief biases that depend on player role. Using a structural model applied to our subject-level beliefs data, we argue that the player role itself induces a higher degree of strategic sophistication in the player who faces more asymmetric payoffs and that this can explain the pattern of bias. This structural feature is not captured by any existing models and, in our view, merits further study.

References

Alberto Abadie. Bootstrap tests for distributional treatment effects in instrumental variable models. *Journal of the American Statistical Association*, 2002. 9

Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *Review of Economic Studies*, 2015. 6.2.7, 9

A.W. Bowman, M.C. Jones, and I. Gijbels. Testing monotonicity of regression. *Journal of Computational and Graphical Statistics*, 1998. 5.1, 7, 5.1, 42

Colin Camerer. Behavioral game theory: Experiments in strategic interaction. *Princeton University Press*, 2003. 1, 9

Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong. Behavioral game theory: thinking, learning, and teaching. *Advances in Understanding Strategic Behavior, Essays in Honor of Werner Guth*, 2004a. 1

Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 2004b. 6.2.4, 9

Monica Capra. Noisy expectation formation in one-shot games. *PhD thesis*, 1999. 35

Hsiao-Chi Chen, James W. Friedman, and Jacques-Francois Thisse. Boundedly rational nash equilibrium: A probabilistic choice approach. *Games and Economic Behavior*, 1997. 1, 9

Miguel Costa-Gomes and Vincent Crawford. Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review*, 2006. 9

Miguel Costa-Gomes and Georg Weizsacker. Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 2008. 6, 1, 5.2, 5.4, 39, 9

Vincent Crawford, Miguel Costa-Gomes, and Nagore Iriberri. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 2013. 7, 9

Ernst Fehr and Klaus Schmidt. A theory of fairness, competition and cooperation. *The Quarterly Journal of Economics*, 1999. 9

Evan Friedman. Stochastic equilibria: Noise in actions or beliefs? *Working Paper*, 2019. 1, 3, 1, 2, 2.2, 10, 2.5, 2.5, 17, 9, 11.2

James Friedman and Claudio Mezzetti. Random belief equilibrium in normal form games. *Games and Economic Behavior*, 2005. 2, 1, 9

John Geanakoplos, David Pearce, and Ennio Stacchetti. Psychological games and sequential rationality. *Games and Economic Behavior*, 1989. 9

Sotiris Georganas, Paul Healy, and Roberto Weber. On the persistence of strategic sophistication. *Journal of Economic Theory*, 2015. 9

Jacob Goeree and Charles Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 2001. 20

Jacob Goeree and Charles Holt. A model of noisy introspection. *Games and Economic Behavior*, 2004. 9

Jacob Goeree and Philippos Louis. M equilibrium. *Working paper*, 2018. 9

Jacob Goeree, Charles Holt, and Thomas Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 2005. 1, 2, 2.1, 2.5, 2.5, 9, 11.2

Jacob Goeree, Charles Holt, and Thomas Palfrey. Quantal response equilibrium. *Princeton University Press*, 2016. 9

Jacob Goeree, Charles Holt, Philippos Louis, Thomas Palfrey, and Brian Rogers. Rank-dependent choice equilibrium: A non-parametric generalization of qre. *Handbook of Research Methods and Applications in Experimental Economics*, 2018. 9

Ben Greiner. Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, 2015. 3.1

Philip Haile, Ali Hortacsu, and Grigory Kosenok. On the empirical content of quantal response equilibrium. *American Economic Review*, 2008. 9, 40

Glenn Harrison, Jimmy Martinez-Correa, and J. Todd Swarthout. Inducing risk neutral preferences with binary lotteries: A reconsideration. *Journal of Economic Behavior and Organization*, 2012. 14

Yuval Heller and Eyal Winter. Biased-belief equilibrium. *working paper*, 2018. 9

Steffen Huck and Georg Weizsacker. Do players correctly estimate what others do?: Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization*, 2002. 1, 5.4

Kyle Hyndman, Antoine Terracol, and Jonathan Vaksman. Beliefs and (in)stability in normal-form games. *Working paper*, 2013. 41

Asen Ivanov. Attitudes to ambiguity in one-shot normal-form games: An experimental study. *Games and Economic Behavior*, 2011. 6

Edi Karni. A mechanism for eliciting probabilities. *Econometrica*, 2009. 3.3

Terri Kneeland. Identifying higher-ordered rationality. *Econometrica*, 2015. 9

Robert Duncan Luce. Individual choice behavior. *Wiley*, 1959. 6.1

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 1995. 1, 1, 9

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1998. 9

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. 3, 4.1

Emerson Melo, Kiril Pogorelskiy, and Matthew Shum. Testing the quantal response hypothesis. *Working Paper*, 2017. 9

Rosemarie Nagel. Unraveling in guessing games: an experimental study. *American Economic Review*, 1995. 7, 6.2, 6.2.4, 9

Yaw Nyarko and Andrew Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 2002. 1, 19, 4.2, 38, 9, 27

Jack Ochs. An experimental study of games with unique mixed strategy equilibria. *Games and Economic Behavior*, 1995. 4, 4.1

Pedro Rey-Biel. Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior*, 2009. 6, 1, 5.2, 9

Brian Rogers, Thomas Palfrey, and Colin Camerer. Heterogeneous quantal response and cognitive hierarchies. *Journal of Economic Theory*, 2009. 35

Ariel Rubinstein and Martin Osborne. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 2003. 9

D. Ruppert, S.J. Sheather, and M.P. Wand. An effective bandwidth selector for local least squares regression. *Journal of the American Statistical Association*, 1995. 42

Elisabet Rutstrom and Nathaniel Wilcox. Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test. *Games and Economic Behavior*, 2009. 1, 4.1, 8.1, 38, 9

K.H. Schlag, J. Tremewan, and J.J. Van Der Weele. A penny for your thoughts: A survey of methods for eliciting beliefs. *Experimental Economics*, 2015. 1, 19, 8.2

Andrew Schotter and Isabel Trevino. Belief elicitation in the laboratory. *Annual Review of Economics*, 2014. 1, 19, 4.2, 8.2, 39

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2 games. *American Economic Review*, 2008. 9

Dale Stahl and Paul Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 1995. 7, 6.2, 6.2.4, 9

Roberto Weber. 'learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 2003. 3.2

Georg Weizsacker. Ignoring the rationality of others: Evidence from experimental normal-forms games. *Games and Economic Behavior*, 2003. 35

11 Appendix

11.1 Experimental instructions

Welcome!

This is an experiment in decision making, and you will be *paid for your participation in cash*. Different subjects may earn different amounts of money. What you earn depends partly on your decisions, partly on the decisions of others, and partly on luck. In addition to these earnings, each of you will receive $10 just for participating in and completing the experiment.

It is the policy of this lab that we are strictly forbidden from deceiving you, so you can trust the experiment will proceed exactly as we describe, including the procedures for payment.

The entire experiment will take place through your computers. It is important that you do not talk or in any way try to communicate with other subjects during the experiment.

Please turn off your cellphones now.

On the screen in front of you, you should see text asking you to wait for instructions, followed by a text box with a button that says "ID". Your computer ID is the number at the top of your desk, which is between 1 and 24. In order to begin the experiment, you must enter your computer ID into the box and press 'ID'. Please do that now.

You should all now see a screen that says "please wait for instructions before continuing". Is there anyone that does not see this screen? This screen will appear at various points throughout the experiment. It is important that whenever you see this screen, you do not click 'continue' until told to do so.

The experiment has *two sections*. We will start with a brief instruction period for Section 1, in which you will be familiarized with the types of rounds you will encounter. Additional instructions will be given for Section 2 after Section 1 is complete.

If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

At the beginning of the experiment, each subject will be assigned the color RED or the color BLUE. There will be an equal number of RED and BLUE subjects. If you are assigned RED, you will be RED for the entire experiment. If you are assigned BLUE, you will be BLUE for the entire experiment.

Section 1 consists of *several rounds*. I will now describe what occurs in each round. First, you will be randomly paired with a subject of the opposite color. Thus, if you are a BLUE subject, you will be paired with a RED subject. If you are a RED subject, you will be paired with a BLUE subject. You will not not know who you are paired with, nor will the other subject know who you are. Each pairing lasts only one round. At the start of the next round, you will be randomly re-paired.

[SLIDE 1]

In each round, you will see a *matrix* similar to the one currently shown on the overhead, though the numbers will change every round. In every round, you and the subject you are paired with will both see the *same matrix*, but remember that one of you is BLUE and one of you is RED.

Both subjects in the pair will *simultaneously* be asked to make a choice. BLUE will choose one of the two rows in the matrix, either 'Up' or 'Down', which we write as 'U' or 'D'. RED will choose one of the two columns, either 'Left' or 'Right', which we write as 'L' or 'R'. We refer to these choices as "actions". Notice that each pair of actions corresponds to one of the 4 cells of the matrix. For instance, if BLUE chooses 'U' and RED chooses 'L', this corresponds to the top-left cell, and similarly for the others.

Thus, depending on *both players' actions*, there are 4 possible outcomes:

- If BLUE chooses 'U' and RED chooses 'L', BLUE receives a payoff of 10, since that is the blue number in the UP–LEFT cell, and RED receives 20, since that is the RED number.
- If BLUE chooses 'D' and RED chooses 'R', BLUE receives a payoff of 11 and RED receives 75.
- And the other two cells UP–RIGHT and DOWN–LEFT are similar.

We reiterate: each number in the matrix is a payoff that might be received by one of the players, depending on *both* players' actions. Are there any questions?

In this section, you will play for 20 rounds and *1 of your rounds* will be chosen for your payment. This 1 round will be selected randomly for each subject, and the payment will depend on the actions taken in that round by you and the subject you were paired with. In the selected round, your payoff in the chosen cell denotes the probability with which you will receive \$10. For example, if you receive a payoff of 60, then for that round you would receive \$10 with 60% probability and \$0 otherwise.

Since every round has an equal chance of being selected for payment, and you do not know which will be selected, *it is in your best interest that you think carefully about all of your choices*.

During the experiment, no feedback will be provided about the other player's chosen action. Only *at the end of the experiment* will you get to see the round that was chosen for your payment and the actions taken by you and the player you were paired with in that round.

Before we begin the first section, you will answer 4 training questions to ensure you understand this payoff structure. In each of these 4 questions, you will be shown a matrix and told the actions chosen by both players. You will then be asked with what probability a particular player earns \$10 if this round were to be selected for payment. That is, you are being asked for their payoff in the appropriate cell. To answer, simply type the probability as a whole number into the box provided and click 'continue'. The page will only allow you to 'continue' when your answer is correct, at which point you may proceed to the next question. Please click 'continue' and answer the 4 training questions now.

[SLIDE 2]

Now that you've completed the training questions and understand the payoff matrices, we will proceed to Section 1. In each round of this section you will be randomly paired with another subject. If you are BLUE, you will be paired with a RED subject, and if you are RED, you will be paired with a BLUE subject. Recall that, at the start of each round, you will be randomly re-paired.

In each round, for each pair, the RED player's task will be to select a column of the matrix, and the BLUE player's task will be to select a row of the matrix, and these actions determine both players' payoffs for the round.

[SLIDE 3]

You should now see an example round on the overhead. This shows the screen for a BLUE player, who is asked to choose between 'U' and 'D'. Notice however that the text instructing you to make a choice is faded. This is because you must wait for 10 seconds before you are allowed to make a decision. Once 10 seconds has passed, the text will darken, indicating that you can now make a selection. The number of seconds remaining until you are able to choose is shown in the bottom right corner. Now the overhead shows what the screen will look like after the 10 seconds have passed.

[SLIDE 4]

The 10 seconds is a minimum time limit. There is no maximum time limit on your choices, and you should feel free to take as much time as you need, even after the 10 seconds has passed. In order to make your selection, simply click on the row or column of your choice. Once you have done so, your choice will be highlighted, and a 'submit' button will appear, as we now show on the overhead.

[SLIDE 5]

You may change your answer as many times as you like before submitting. If you would like to undo your choice, simply click again on the highlighted row or column. Once you are satisfied with your choice, click 'submit' to move on to the next round.

Before beginning the paid rounds of Section 1, we will play 4 practice rounds to familiarize you with the interface. These rounds will *not* be selected for payment. Are there any questions about the game, the rules, or the interface before we begin the practice rounds?

Please click 'continue' and begin the practice rounds now. You will notice that you have been assigned either RED or BLUE. This will be your color throughout the experiment. Please continue until you have completed the 4 practice rounds.

You have now completed the practice rounds, and we will proceed to the paid rounds of Section 1. Section 1 consists of 20 rounds, exactly like those you have just played. Recall that, in each round, you will be randomly paired with another subject and that one round will be randomly selected for payment. Are there any questions about the game, the rules, or the interface before we begin?

[SLIDE 6]

Please click 'continue' and play Section 1 now. The rules we discussed for Section 1 will be shown on the overhead as a reminder throughout.

[SLIDE 7]

We will now have a brief instruction period for Section 2, in which you will be familiarized with the types of rounds you will encounter.

If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise once play has begun, raise your hand, and an experimenter will come and assist you.

In this section, each round will be similar to those from Section 1. You will see some of the same matrices and your assignment of RED or BLUE will be the same as before.

Now, however, after being shown a matrix, your task will be to give your *belief* or *best guess* about the probability that a randomly selected subject chose a particular action when playing the same matrix in Section 1. That is, you will be shown a matrix, and the computer will randomly select a round from Section 1 in which the same matrix was played. Then,

- If you are RED, you will be asked for the probability that a randomly selected BLUE player chose 'U' in that round in Section 1.
- If you are BLUE, you will be asked for the probability that a randomly selected RED player chose 'L' in that round in Section 1.

As before, you will be paid for your responses. We will now describe this payment mechanism.

[SLIDE 8]

Consider first the matrix that is shown on the overhead. Please imagine that the computer has randomly selected a round from Section 1 in which this matrix was played. We wish to know your belief about the probability that a randomly selected RED player chose 'L' in that round. Please, take some time now to think carefully about what you believe this probability to be.

[SLIDE 9]

Consider the question that is now shown on the overhead, which asks which of the following you would prefer:

- Under Option A, you receive \$5 if a randomly selected RED player chose 'L' in that round, and \$0 otherwise.

- Under Option B, you receive \$5 with probability 75%, and \$0 otherwise.

Please think carefully about which of these two options you would prefer.

Presumably, if you believe the probability that a randomly selected RED player chose 'L' is greater than 75%, then you would prefer Option A, which you believe gives you the highest probability of a \$5 prize. For example, if you believe this probability is 89%, you would choose Option A since 89 is greater than 75.

If, on the other hand, you believe the probability that a randomly selected RED player chose 'L' is less than 75%, then you would prefer Option B, which you believe gives you the highest probability of a \$5 prize. For example, if you believe this probability is 22%, you would choose Option B since 22 is less than 75.

In this way, your answer to this question will tell us whether you believe this probability is greater than or less than 75%.

[SLIDE 10]

Now imagine we asked you 101 of these questions, with the probability in Option B ranging from 0% to 100%. Presumably you would answer each of these questions as described previously. That is, for questions for which the probability in Option B is *below* your belief, you would choose Option A, and for questions for which the probability in Option B is *above* your belief, you would choose Option B. Imagine, for example, you believe that there is a 64% probability that a randomly selected RED player chose 'L' in the selected round. Then, you would select Option A for all questions before #64, and Option B for all questions after #64. For Question #64, you could make either selection.

[SLIDE 11]

In this case, your selections would be as shown on the overhead, with the chosen options in black and the unchosen options in gray. From these answers, we could determine that you believe the probability that a randomly selected RED player chose 'L' is 64%.

In each round of this section, you will be faced with a table of 101 questions as shown on the overhead. To save time, instead of having you answer each question individually, we will simply ask you to type in your belief, and the answers to these 101 questions will be automatically filled out as above. That is, for rows of the table in which the probability in Option B is below your stated belief you will automatically select Option A, and for rows of the table in which the probability in Option B is at or above your stated belief you will automatically select Option B.

If this round is chosen for payment, one of the 101 rows of the table will be randomly selected and you will be paid according to your chosen option in that row. If you chose Option A in that row, a subject of the relevant color will be randomly chosen, and you will receive \$5 if they played the relevant action in the selected round of Section 1. If you chose Option B in that row, you will receive \$5 with the probability given in that option.

It is thus in your best interest, given your belief, to state your belief accurately. Otherwise, if you type something other than your belief, there will be rows of the table for which you will not be selecting the option that you believe gives you the highest probability of receiving a \$5 prize.

In this section you will play 40 rounds, giving 40 such beliefs. At the end of the section, 2 rounds will be randomly chosen for payment. For each of these rounds, one of the 101 rows of the table will be randomly selected and you will be paid according to your chosen option in that row.

Are there any questions about this?

In addition to stating a belief, in each round you will also be asked to choose an action, as you did in Section 1. Now, however, the other action will not be determined by another subject acting simultaneously. Instead, recall that the computer has randomly selected a round from Section 1 featuring the matrix shown on your screen. The computer will also randomly select a player of the other color and record the action they took in that round. This is the action that you will be paired with. That is:

- If you are RED, the BLUE action will be that which a randomly selected BLUE player chose in the selected round of Section 1.
- If you are BLUE, the RED action will be that which a randomly selected RED player chose in the selected round of Section 1.

Again, the randomly selected round from Section 1 will feature the same matrix shown on your screen, so your payoff is determined as if you were paired with a randomly selected player from Section 1, rather than being paired with a player who chooses an action simultaneously.

As in Section 1, your payoff from taking an action gives the probability of earning $10 if the round is chosen for payment.

At the end of the section, 2 rounds will be randomly chosen for payments based on your actions. This is in addition to the 2 rounds randomly chosen for payments based on your beliefs. Moreover, the randomization algorithm that selects these rounds will ensure that all 4 rounds feature different matrices and that these matrices will be different from that selected for payment in Section 1. In particular, this means that if a round is selected for an action-payment, it cannot also be selected for a belief-payment and vice versa.

As before, since you do not know which round will be selected for payment, nor which type of payment it will be selected for, these payment procedures ensure that, in each round, *it is in your best interest to both state your belief accurately and choose the action that you think is best.*

[SLIDE 12]

You should now see an example round on the overhead. This shows the screen for a BLUE player. As in Section 1, you will see the matrix in the middle of the screen. At the top of the screen, you are told that the computer has randomly selected a round of Section 1 in which this matrix was played.

Below this, the instructions are shown, and are again faded for 10 seconds. Once 10 seconds has passed, the text asking you for your belief will darken as now shown on the overhead.

[SLIDE 13]

You will not be able to select an action until after you have entered your belief.

Once you have entered your belief, the resulting probabilities will appear below or beside the matrix and the text asking you to select your action will darken, as now shown on the overhead.

[SLIDE 14]

Your belief must be a whole number between 0 and 100 inclusive. Once you enter your belief, we will automatically 'fill out' the questions in the 101 rows based on your belief as previously described. If you wish, at any time you may scroll down to observe the 101 rows.

As in Section 1, once you have selected an action, it will be highlighted on the matrix, as now shown on the overhead.

[SLIDE 15]

At this point, you may freely modify both your belief and action as many times as you wish before pressing 'submit'. Remember that there is no upper time limit on your choices, and you should feel free to take as much time as you need, even after the minimum 10 seconds has passed.

Before beginning the paid rounds of Section 2, we will play 3 practice rounds to familiarize you with the interface. These rounds feature the same matrices as the practice rounds from Section 1, and will not be selected for payment. Are there any questions about the game, the rules, or the interface before we begin the practice rounds?

Please click 'continue' to be taken to the first practice round now. Recall that your belief must be a whole number between 0 and 100 inclusive, and at any time you may scroll down to see the table of 101 questions. Please continue until you have completed the 3 practice rounds.

You've now completed the practice rounds, and we will proceed to the paid rounds of Section 2.

[SLIDE 16]

Recall that Section 2 consists of 40 rounds, exactly like those you have just played. 4 rounds will be randomly selected for payment–2 rounds for beliefs and 2 rounds for your actions. Again, these 4 rounds will feature different matrices to each other and to the matrix selected for payment in Section 1. The payment procedures ensure that *it is always in your best interest to both state your belief accurately and choose the action that you think is best*. Unlike Section 1, Section 2 will be played at your own pace without waiting for other subjects between rounds. Once you have completed Section 2, please remain seated quietly until all subjects have finished.

Are there any questions about the game, the rules, or the interface? If you have any questions during the remainder of the experiment, raise your hand, and an experimenter will come and assist you. You may click 'continue' and play Section 2 now. The rules we discussed for Section 2 will be shown on the overhead as a reminder throughout.

You have now completed the experiment. All that remains is to organize payments. To do this, you will be shown a page with all of your randomly selected rounds and your earnings in each. This page will also show you how to fill out the payment receipt at your desks. Before reaching this page, you will see an explanation page describing how the results are determined and how to read them. You may click 'continue' now and read through the explanation page. Then continue to the payments page, where you will see your results and fill out your receipt.

11.2 Proofs

Proof of Proposition 1. Existence follows from Brouwer's fixed point theorem after showing $\Psi_i : [0,1] \to [0,1]$ is continuous. To this end, let $\mu_i(\cdot|\sigma_j)_{\sigma_j \in [0,1]}$ be the family of Borel measures derived from $F_i(\cdot|\sigma_j)_{\sigma_j \in [0,1]}$ that define the belief distributions. From (B1) and (B2), $\mu_i(\cdot|\sigma_j)_{\sigma_j \in [0,1]}$ has the property that $\mu_i(B|\sigma_j)$ is continuous in $\sigma_j \in (0,1)$ for any Borel subset $B \in \mathcal{B}([0,1])$. From this and the fact that $Q_i \circ \bar{u}_i(\sigma_j)$ is continuous in $\sigma_j \in [0,1]$, it is immediate that $\Psi_i(\sigma_j)$ is continuous in $\sigma_j \in (0,1)$. So we need only consider $\sigma_j \to 0^+$ (the case of $\sigma_j \to 1^-$ is similar). From (B4) (and (B1) and (B2)), there are discontinuities at the endpoints: $\mu_i(\{0\}|\sigma_j) = 0$ for $\sigma_j > 0$ but $\mu_i(\{0\}|0) = 1$. However, from (B1) and (B2), $\mu_i(B|\sigma_j)$ is continuous as $\sigma_j \to 0^+$ if $B = [0,\epsilon)$ or $B = (\epsilon, \epsilon_2)$ (i.e. if B or its complement is well-separated from 0), which implies that $\mu_i([0,\epsilon)|\sigma_j) \to 1$ as $\sigma_j \to 0^+$ for any $\epsilon > 0$. Hence, $\Psi_i(\sigma_j)$ is continuous since $\Psi_i(0) = Q_i \circ \bar{u}_i(0)$ and as $\sigma_j \to 0^+$, beliefs become arbitrarily concentrated within a neighborhood of 0 and $Q_i \circ \bar{u}_i(\sigma_j)$ is continuous in $\sigma_j \in [0,1]$.

□

Proof of Proposition 2. From (A3) and (B3), $\Psi_U(\sigma_L)$ is strictly increasing in σ_L: for any belief realizations $\sigma_L^{''} > \sigma_L^{'}$, $Q_U(\bar{u}_1(\sigma_L^{''})) > Q_U(\bar{u}_1(\sigma_L^{'}))$ by (A3), and if σ_L increases, the distribution of σ_L^* increases in the sense of stochastic dominance by (B3). Similarly, $\Psi_L(\sigma_U)$ is strictly decreasing in σ_U, and so the equilibrium is unique.

□

Proof of Proposition 3. Suppose $\sigma_L < \sigma_L^{NE}$. By (B4), it must be that $F_1(\sigma_L|\sigma_L) = \frac{1}{2}$, and hence $\mathbb{P}(\sigma_L^*(\sigma_L) < \sigma_L^{NE}) \in (\frac{1}{2}, 1)$ and $\mathbb{P}(\sigma_L^*(\sigma_L) > \sigma_L^{NE}) = 1 - \mathbb{P}(\sigma_L^*(\sigma_L) < \sigma_L^{NE}) \in (0, \frac{1}{2})$. By (A4), $Q_U \circ \bar{u}(\sigma_L^{'}) \in (0, \frac{1}{2})$ for belief realization $\sigma_L^{'} < \sigma_L^{NE}$ and $Q_U \circ \bar{u}(\sigma_L^{'}) \in (\frac{1}{2}, 1)$ for belief realization $\sigma_L^{'} > \sigma_L^{NE}$. Together, this implies that $\Psi_U(\sigma_L) \in (0, \frac{3}{4})$ for $\sigma_L < \sigma_L^{NE}$. Using similar arguments, it must be that $\Psi_U(\sigma_L)$ must satisfy (1) for all σ_L. Conversely, let $\sigma_L < \sigma_L^{NE}$ and $c \in (0, \frac{3}{4})$ be arbitrary. $\Psi_U(\sigma_L) = c$ can be obtained by setting $\sigma_L^*(\sigma_L) = \begin{cases} 0 & w.p.\ \frac{1}{2} \\ 1 & w.p.\ \frac{1}{2} \end{cases}$ and Q_U so that $\frac{1}{2} Q_U \circ (\bar{u}_1(0)) + \frac{1}{2} Q_U \circ (\bar{u}_1(1)) = c$. The only restrictions are that $Q_U \circ (\bar{u}_1(0)) \in (0, \frac{1}{2})$ and $Q_U \circ (\bar{u}_1(1)) \in (\frac{1}{2}, 1)$, so this is feasible. This construction violates (B1) and (B2), but can be modified to satisfy these axioms by smoothing out the distribution of $\sigma_L^*(\sigma_L)$ arbitrarily little. As σ_L increases to $\sigma_L^{'} < \sigma_L^{NE}$, this construction can be extended so that $\Psi_U(\sigma_L^{'}) = c^{'}$ for any $c^{'} \in (c, \frac{3}{4})$ in such a way that none of the axioms are violated. Proceeding in this fashion, any $\Psi_U : [0,1] \to [0,1]$ that is continuous, strictly increasing, and satisfying (1) can be induced for some $\{Q_U, \sigma_L^*\}$. Part *(ii)* is similar, and part *(iii)* follows since the QNBE is the intersection of the constructed Ψ_U and Ψ_L.

□

Proof of Proposition 4. As X increases, $Q_U(\bar{u}(\sigma_L^{'}))$ strictly increases for any $\sigma_L^{'}$. Hence, $\Psi_U(\sigma_L)$ shifts up strictly. Since $\Psi_U(\sigma_L)$ is strictly increasing and $\Psi_L(\sigma_U)$ is strictly decreasing, it must be that σ_U^{QNBE} strictly increases and σ_L^{QNBE} strictly decreases (see Figure 1).

□

Proof of Proposition 5. The *only if* direction follows immediately from Propositions 3 and 4. We omit the *if* direction because it is very similar to that in the proof of Proposition 6 below as it basically combines the results for QRE and NBE.

□

Proof of Proposition 6. The *only if* direction can be found for very similar games for QRE in Goeree et al. [2005] and for NBE in Friedman [2019] in terms of X. Since σ_L^{NE} is strictly decreasing in X, the result follows. The *if* direction is novel. For QRE, it is very simple. Let $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ satisfy *(i)-(iv)*. For player 1, set Q_U such that $Q_U \circ \bar{u}_1^X(\hat{\sigma}_L^X) = \hat{\sigma}_U^X$ for all X and similarly for Q_L. It is easy to check that, because $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ satisfies *(i)-(iv)*, this does not violate (A3) or (A4). Since this only pins down Q at finite points, the rest of Q can be constructed in a way that satisfies (A1)-(A4). For NBE, it is more involved. Let $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ satisfy *(i)-(iv)*. For player 1, $\Psi_U^X(\sigma_L) = 1 - F_1(\sigma_L^{NE,X}|\sigma_L)$, so we must construct a family of CDFs $F_1(\cdot|\cdot)$ such that $1 - F_1(\sigma_L^{NE,X}|\hat{\sigma}_L^X) = \hat{\sigma}_U^X$ for all X and that satisfies (B1)-(B4). We illustrate this construction in Figure 20. Given that $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ satisfies *(i)-(iv)*, we have that (1) $\hat{\sigma}_L^X < \hat{\sigma}_L^{X'}$ whenever $X > X'$, (2) $\frac{1}{2} > \hat{\sigma}_L^X > \sigma_L^{NE,X}$ if $X > 20$ and $\frac{1}{2} < \hat{\sigma}_L^X < \sigma_L^{NE,X}$ if $X < 20$, and (3) $\hat{\sigma}_U^X > \frac{1}{2}$ if $X > 20$ and $\hat{\sigma}_U^X < \frac{1}{2}$ if $X < 20$. As illustrated in the figure, this allows us to construct, for each $\hat{\sigma}_L^X$, a CDF $F_1(\cdot|\hat{\sigma}_L^X)$ that satisfies $1 - F_1(\sigma_L^{NE,X}|\hat{\sigma}_L^X) = \hat{\sigma}_U^X$, meaning it can generate the data, and also such that each $F_1(z|\hat{\sigma}_L^X)$ is strictly increasing in $z \in [0,1]$, $F_1(z|\hat{\sigma}_L^{X'}) < F_1(z|\hat{\sigma}_L^X)$ for all $z \in (0,1)$ if $\hat{\sigma}_L^X < \hat{\sigma}_L^{X'}$, and $F_1(\hat{\sigma}_L^X|\hat{\sigma}_L^X) = \frac{1}{2}$ for all X. Hence, the constructed $\{F_1(\cdot|\hat{\sigma}_L^X)\}_X$ satisfies (B1)-(B4) and can be extended to $\{F_1(\cdot|\sigma_L^X)\}_{\sigma_L^X \in (0,1)}$. For the extension, order the values of X in the dataset: $X^1 > X^2 > ... > X^n$ and $\hat{\sigma}_L^{X^1} < \hat{\sigma}_L^{X^2} < ... < \hat{\sigma}_L^{X^n}$. For $\sigma_L^X \in (\hat{\sigma}_L^{X^i}, \hat{\sigma}_L^{X^{i+1}})$ set $F_1(z|\sigma_L^X) = \alpha(X) F_1(z|\hat{\sigma}_L^{X^i}) + (1 - \alpha(X)) F_1(z|\hat{\sigma}_L^{X^{i+1}})$ where $\alpha(X)$ is such that $F_1(\sigma_L^X|\sigma_L^X) = \frac{1}{2}$, which is uniquely defined. Finally, extending for $\sigma_L^X < \hat{\sigma}_L^{X^1}$ and $\sigma_L^X > \hat{\sigma}_L^{X^n}$ is straightforward.

□

Proof of Proposition 7. *(i)*: Let w and v with $w = f(v)$ for some concave f. Let $X > 20$. Without loss, normalize so that $w(0) = v(0) = 0$ and $w(X) = v(X) = 1$. For arbitrary utility function u, it is easy to show that $\tilde{\sigma}_L^{u,X} = \frac{u(20)}{u(20)+1}$. Since w is more concave than v, $w(20) > v(20)$ and thus $\tilde{\sigma}_L^{w,X} > \tilde{\sigma}_L^{v,X}$. Similarly, if $X < 20$, normalize without loss so that $w(0) = v(0) = 0$ and $w(20) = v(20) = 1$. This implies that $\tilde{\sigma}_L^{u,X} = \frac{1}{1+u(X)}$. Since w is more concave than v, $w(20) > v(20)$ and thus $\tilde{\sigma}_L^{w,X} < \tilde{\sigma}_L^{v,X}$. Part *(ii)* is the same, except with $v(z) = z$, which implies $\tilde{\sigma}_L^{v,X} = \frac{20}{20+X} = \sigma_L^{NE,X}$.

□

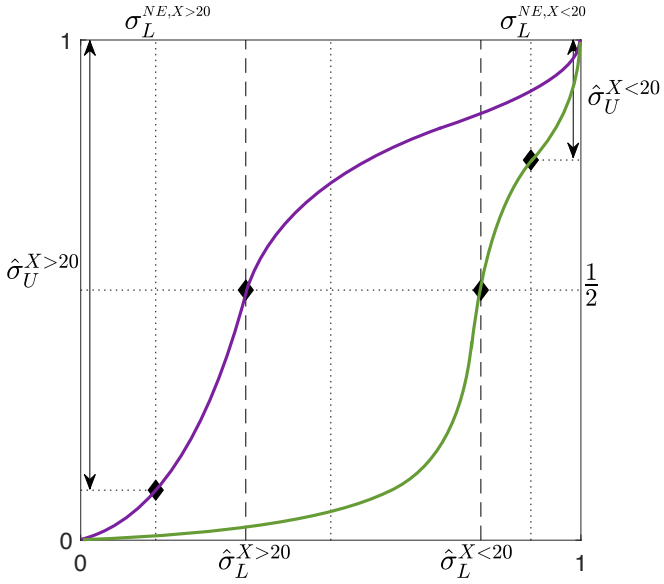

Figure 20: *Construction of belief-map in proof of Proposition 6.* This illustrates the constructed CDFs of player 1's beliefs to rationalize as NBE the actions dataset $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ for $X \in \{X^{'}, X^{''}\}$ with $X^{'} > 20$ and $X^{''} < 20$. The purple CDF is $F_1(\cdot|\hat{\sigma}_L^{X^{'}})$ and the green CDF is $F_1(\cdot|\hat{\sigma}_L^{X^{''}})$.

11.3 Details of statistical tests

For what follows, let N_i^S denote the *set* of subjects in role $i \in \{1,2\}$ in session $S \in \{[\text{A,BA}],[\text{A,A}]\}$.

Bowman et al. (1998) test. The null hypothesis is that a regression function, which in our case is the expected action conditional on beliefs, is weakly monotone.

Let b_{sl}^{ix} be the lth of 5 belief statements for subject s in role i for game $x \in \mathcal{G}$ where $\mathcal{G}$ is a set of games that is either the entire set of six games $\{X80, X40, ..., X1\}$ or any one of these games. Let $a_{sl}^{ix} \in \{0,1\}$ be the action corresponding to belief b_{sl}^{ix}. The method is based on running non-parametric regressions of a_{sl}^{ix} on b_{sl}^{ix}, and we denote such estimators by $\hat{m}(b)$ and $\hat{w}(b)$. We describe the process when aggregating across all subjects of type i indexed by set $\mathcal{I}$, which can be either all subjects $N_i^{[A,BA]}$ or a single subject $\{s\}$.

- Step 1: Find the critical local linear (lowess) regression bandwidth h_c, which is the smallest such that $\hat{m}(b;h_c)$ is weakly monotone (increasing if $i=1$, decreasing if $i=2$).
- Step 2: For each s,x,k, calculate $\hat{\epsilon}_{sl}^{ix} = a_{sl}^{ix} - \hat{w}(b_{sl}^{ix};h_0)$ where $\hat{w}$ is some estimator with bandwidth h_0 chosen to minimize the mean integrated squared error $\sum_{s,x,l}(\hat{m}(b_{sl}^{ix}) - m(b_{sl}^{ix}))$ where $a_{sl}^{ix} = m(b_{sl}^{ix}) + \epsilon_{sl}^{ix}$ is the true model.[42]
- Step 3: Resample the subjects in $\mathcal{I}$ with replacement $|\mathcal{I}|$ times. Conditional on drawing each subject s, resample 5 times with replacement from $\{\hat{\epsilon}_{s1}^{ix}, \hat{\epsilon}_{s2}^{ix}, ..., \hat{\epsilon}_{s5}^{ix}\}$ for each $x \in |\mathcal{G}|$ for a total bootstrap sample of $5 \cdot |\mathcal{G}| \cdot |\mathcal{I}|$ observations $\{\hat{\epsilon}_t^{*i}\}_{t\in\{1,2,...,30\cdot|\mathcal{I}|\}}$ and thus $\{a_t^{*i} = \hat{m}(b_t^i;h_c) + \hat{\epsilon}_t^{*i}\}_{t\in\{1,2,...,30\cdot|\mathcal{I}|\}}$ where b_t^i is the belief associated with $\hat{\epsilon}_t^{*i}$.
- Step 4: Apply $\hat{m}$ using h_c to $\{(b_t^i, a_t^{*i})\}_{t\in\{1,...,5\cdot|\mathcal{G}|\cdot|\mathcal{I}|\}}$ and observe whether or not the result is monotone.
- Step 5: Repeat Steps 3 and 4 $\mathcal{B} = 5,000$ times, constructing the p-value by determining the proportion of estimates at Step 4 which are not monotonic (not everywhere increasing if $i=1$, not everywhere decreasing if $i=2$).

Abadie (2002) test. The null hypothesis is that $F_i(z|\sigma_j^x) \leq F_i(z|\sigma_j^y)$ for $z \in [0,1]$.

- Step 1: Compute Kolmogorov-Smirnov statistic $\hat{T} = \left(\frac{5\left|N_i^{[\text{A,BA}]}\right| \cdot 5\left|N_i^{[\text{A,BA}]}\right|}{5\left|N_i^{[\text{A,BA}]}\right| + 5\left|N_i^{[\text{A,BA}]}\right|}\right)^{\frac{1}{2}} \cdot \sup_{t\in[0,1]}\left(\hat{F}_i(z|\sigma_j^x) - \hat{F}_i(z|\sigma_j^y)\right)$, where $\hat{F}_i(z|\sigma_j^x)$ and $\hat{F}_i(z|\sigma_j^y)$ are the empirical CDFs of beliefs.
- Step 2: Resample the subjects in $N_i^{[\text{A,BA}]}$ with replacement. Conditional on drawing each subject, resample with replacement from her 10 belief statements, i.e. pooled together from x and y, and assign the first 5 to group x and the second 5 to group y. Do this $\left|N_i^{[\text{A,BA}]}\right|$ times to form two bootstrapped CDFs $F_i^*(z|\sigma_j^x)_b$ and $F_i^*(z|\sigma_j^y)_b$ where b is the bootstrap index. Use these to calculate T_b^*.
- Step 3: Repeat previous step $\mathcal{B} = 5,000$ times.
- Step 4: Calculate p-value as $\sum_{b=1}^{\mathcal{B}} \mathbf{1}\{T_b^* > \hat{T}\}/\mathcal{B}$.

[42] Bowman et al. [1998] suggest to use lowess regression with bandwidth selected by the method of Ruppert et al. [1995]. We opt instead for local linear kernel regression with cross-validation based bandwidth selection for its wider availability in statistical packages (Stata 15.0 command *npregress kernel*, which uses optimal bandwidth selection by default).

11.4 The effects of belief elicitation: a closer look

In the top panels of Appendix Figure 21, we plot the action frequencies from [A,BA] and [A,BA]. That is, we are comparing first-stage actions (without belief elicitation) to second-stage actions, each of which was preceded by belief elicitation.[43] For both players, we observe systematic and significant differences between the two stages. This is confirmed in the first 2 columns of Appendix Table 13 in which we regress actions on indicators for each of the six X-games (omitted from the table) and indicators for each of the six games interacted with an indicator for the the second-stage. F-tests reject that the action frequencies are the same across stages.

Our hypothesis is that these differences are *caused by* belief elicitation. However, the two stages differ in which came first, the fact that the games in the second stage are played against previously recorded actions, the number of rounds, and very slightly in their composition of games. To pin down the cause, we ran the additional [A,A] treatment. This is identical to the [A,BA] treatment except beliefs are not elicited (and instructions never mention belief elicitation).

The bottom panels of Appendix Figure 21 plot the action frequencies from [A,A] and [A,A], and Columns 3-4 of Appendix Table 13 replicate columns 1-2 for the [A,A] treatment. We find that the actions data is statistically indistinguishable between the two stages of the [A,A] treatment. In particular, the difference between player 1's first- and second-stage action frequencies completely disappears. We conclude that belief elicitation *does* have an effect on actions.

Our goal in this paper is to study the relationships between beliefs and the associated actions without our own interference as experimenters. How does the fact that belief elicitation affects second-stage actions influence our conclusions? This depends on what is driving the effect.

There are two channels through which belief elicitation may have an effect on the actions subjects take. It could be that *(i)* elicitation affects beliefs or *(ii)* elicitation affects actions conditional on beliefs. If only the former "beliefs channel" is active, only testing axioms on the belief-map would be affected since we condition on second-stage beliefs when testing axioms on the action-map. If only the latter "actions channel" is active, only testing axioms on the action-map would be affected since we are comparing second-stage beliefs to first-stage actions when testing axioms on the belief-map. Since we do not observe the beliefs subjects had in the first stage, there is no way of knowing definitively the degree to which either channel is active, but our previous analysis gives insight.

From the top panels of Appendix Figure 21, the direction of the change in actions due to belief elicitation is systematic. For player 1, subjects are more likely to choose D for $X > 20$ and U for $X < 20$. For player 2, subjects are more likely to choose R for $X > 20$ and L for $X < 20$. While the effect is systematic for both players, the effect for player 2 is rather small.

Suppose the elicitation effect is through the beliefs channel. In the context of our structural model, this is consistent with *increased* sophistication for player 1, with $\hat{\sigma}_U^{[A,B\underline{A}]}$ resembling the first-stage actions of the high sophistication group in Figure 14. For player 2, the effect is consistent with *decreased* sophistication. We find this plausible for player 1 only as it is intuitive that eliciting beliefs may induce subjects to think more deeply about opponents' behavior.

Suppose the elicitation effect is through the actions channel. Player 2's actions become more extreme, so it could be that belief elicitation simply reduces the probability that player 2-subjects make mistakes or "trembles" for any given belief, which we find very plausible. For player 1, this would have an effect in the opposite direction from that which we observe, so it may be at play, but is overwhelmed by an effect in the opposite direction.

For player 1, we believe the beliefs channel dominates. If this is the case, prior to elicitation, player 1-subjects had beliefs closer to the uniform distribution and thus beliefs with a more conservative bias. For player 2, we believe the actions channel dominates, so prior to elicitation, player 2-subjects' actions would have been noisier for any given belief. Under these interpretations, all of our main conclusions would be unchanged.

If belief elicitation affects actions largely through beliefs by making subjects more sophisticated, a natural hypothesis is that more sophisticated subjects will be less affected by belief elicitation. Since we have a measure of sophistication from the dominance solvable games, this is easy to test. In Appendix Figure 22, we plot the first- and second-stage action frequencies from [A,BA] by the sophistication groupings from Section 6.2.3. This seems to indicate that the effect of elicitation is primarily driven by subjects with low sophistication, and this is confirmed in Appendix Table 14.

[43] The results are similar if, instead, we compare the data from [A,∘] and [A,BA], but this would be somewhat confounded by composition effects.

	[A,BA]		[A,A]	
	(1) $\hat{\sigma}_U$	(2) $\hat{\sigma}_L$	(3) $\hat{\sigma}_U$	(4) $\hat{\sigma}_L$
2nd stage × X80	-0.119**	-0.057	0.048	-0.022
	(0.030)	(0.156)	(0.500)	(0.754)
2nd stage × X40	-0.019	-0.059	0.007	0.048
	(0.748)	(0.111)	(0.884)	(0.362)
2nd stage × X10	0.130**	0.105**	-0.056	-0.081
	(0.013)	(0.031)	(0.438)	(0.266)
2nd stage × X5	0.194***	0.007	0.019	0.130*
	(0.000)	(0.850)	(0.746)	(0.070)
2nd stage × X2	0.070	0.091**	0.011	0.000
	(0.202)	(0.040)	(0.867)	(1.000)
2nd stage × X1	0.124**	0.102**	0.015	0.000
	(0.037)	(0.016)	(0.803)	(1.000)
F-test	5.12***	3.24***	0.21	0.92
	(0.000)	(0.004)	(0.972)	(0.485)
[d1,d2]	[6,323]	[6,335]	[6,161]	[6,161]
Observations	2592	2676	1134	1134

p-values in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Table 13: *Effects of belief elicitation.* We regress actions on indicators for all six X-games (omitted) and indicators for each of the six games interacted with an indicator for the second stage. Columns 1-2 are for [A,BA], and columns 3-4 are for [A,A]. We also report the results of F-tests of the hypothesis that all six coefficients are zero. Standard errors are clustered at the subject-game level.

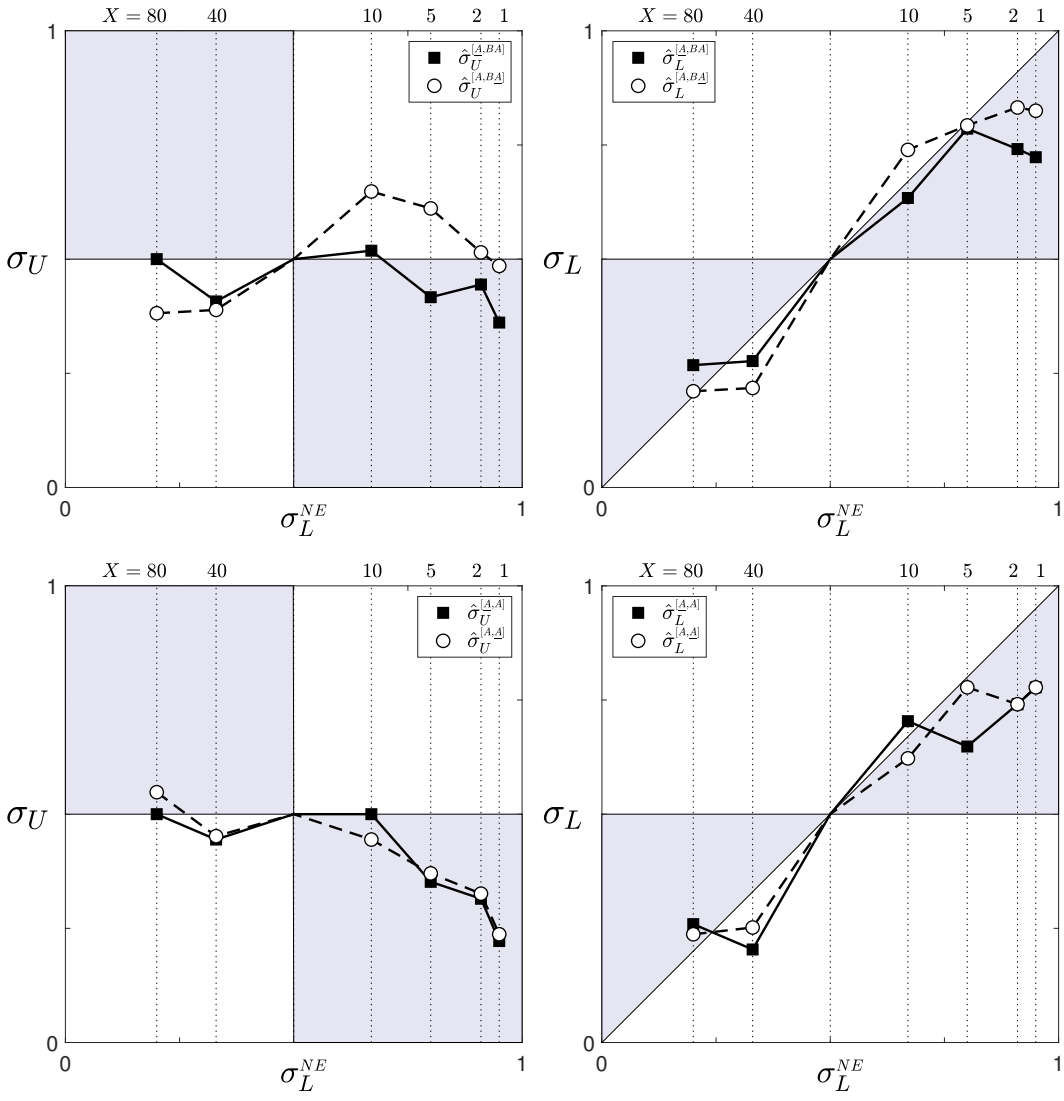

Figure 21: *Effects of belief elicitation.* The top panels plot first-stage and second-stage actions from [A,BA], and shows a systematic difference between the two stages. The bottom panels plot first-stage and second-stage actions from [A,A], and shows no difference between the stages.

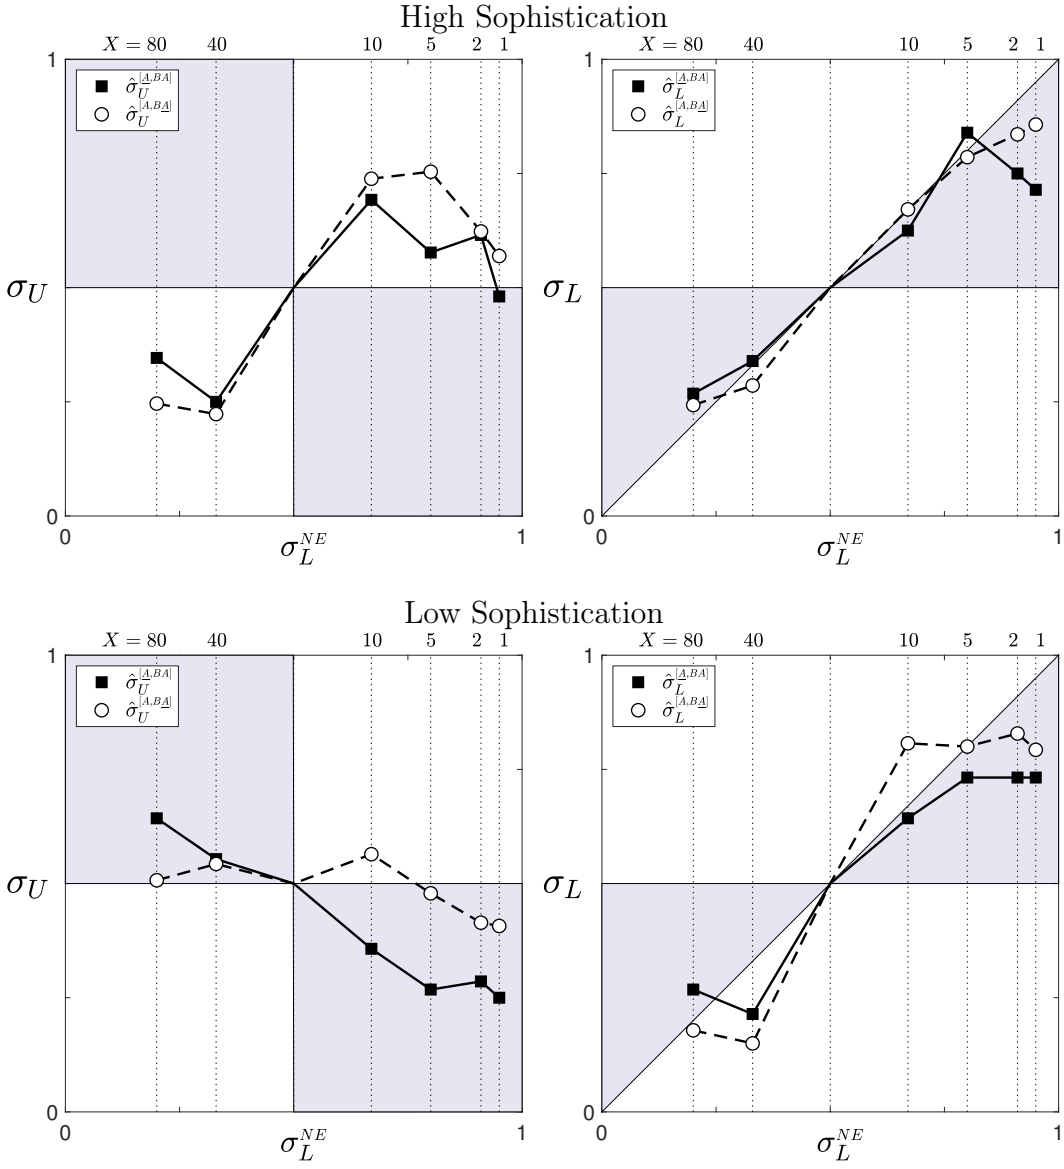

Figure 22: *Effects of belief elicitation by sophistication group.*

	Low Sophistication		High Sophistication	
	(1) $\hat{\sigma}_U$	(2) $\hat{\sigma}_L$	(3) $\hat{\sigma}_U$	(4) $\hat{\sigma}_L$
2nd stage × X80	-0.136	-0.089*	-0.100	-0.025
	(0.121)	(0.092)	(0.120)	(0.679)
2nd stage × X40	-0.011	-0.064	-0.027	-0.054
	(0.907)	(0.301)	(0.696)	(0.188)
2nd stage × X10	0.207***	0.164***	0.046	0.046
	(0.003)	(0.009)	(0.543)	(0.531)
2nd stage × X5	0.211***	0.068	0.177**	-0.054
	(0.005)	(0.289)	(0.020)	(0.152)
2nd stage × X2	0.129*	0.096*	0.008	0.086
	(0.069)	(0.097)	(0.928)	(0.204)
2nd stage × X1	0.157**	0.061	0.088	0.143**
	(0.045)	(0.280)	(0.325)	(0.022)
F-test	4.50***	2.69**	1.58	1.89*
	(0.000)	(0.016)	(0.157)	(0.086)
[d1,d2]	[6,323]	[6,167]	[6,155]	[6,167]
Observations	1176	1176	1092	1176

p-values in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Table 14: *Effects of belief elicitation by sophistication group.* For high and low sophistication groups and for each player, we regress actions from both stages of [A,BA] on indicators for all six X-games (omitted) and indicators for each of the six games interacted with an indicator for the second stage. Columns 1-2 are for the high sophistication group, and columns 3-4 are for the low sophistication group. We also report the results of F-tests of the hypothesis that all six coefficients are zero. Standard errors are clustered at the subject-game level.

11.5 Experimental interface

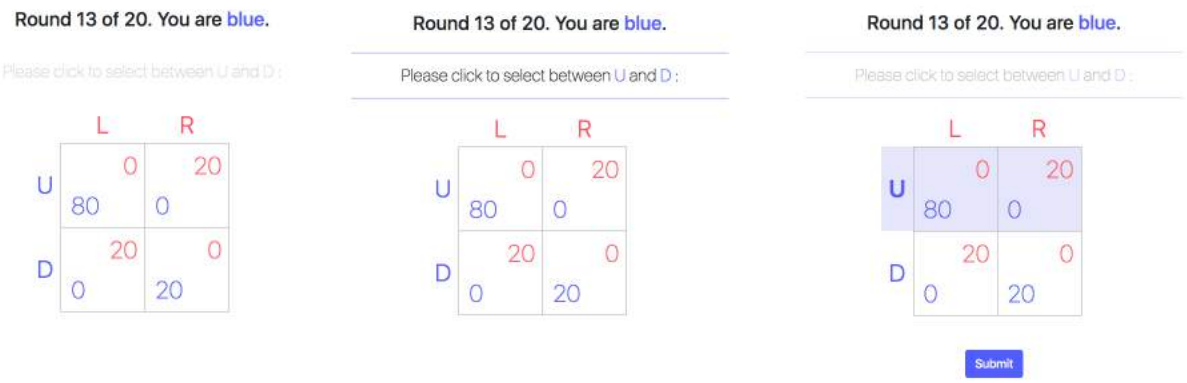

Figure 23: *Screenshots from first stage.* This figure shows an example round from the perspective of a player 1-subject (blue). At the start of the round, the subject sees the payoff matrix (left screen), and a 10 second timer counting down to 0 (not shown here) is seen at the bottom right corner of the screen. After 10 seconds pass, the text "Please click to select between U and D:" darkens (middle screen) indicating that the subject may take an action. To select an action, the subject clicks on a row of the matrix. The row becomes highlighted and a 'Submit' button appears (right screen). At this point, the subject may freely modify his answer before submitting. The subject may undo his action choice by clicking again on the highlighted row.

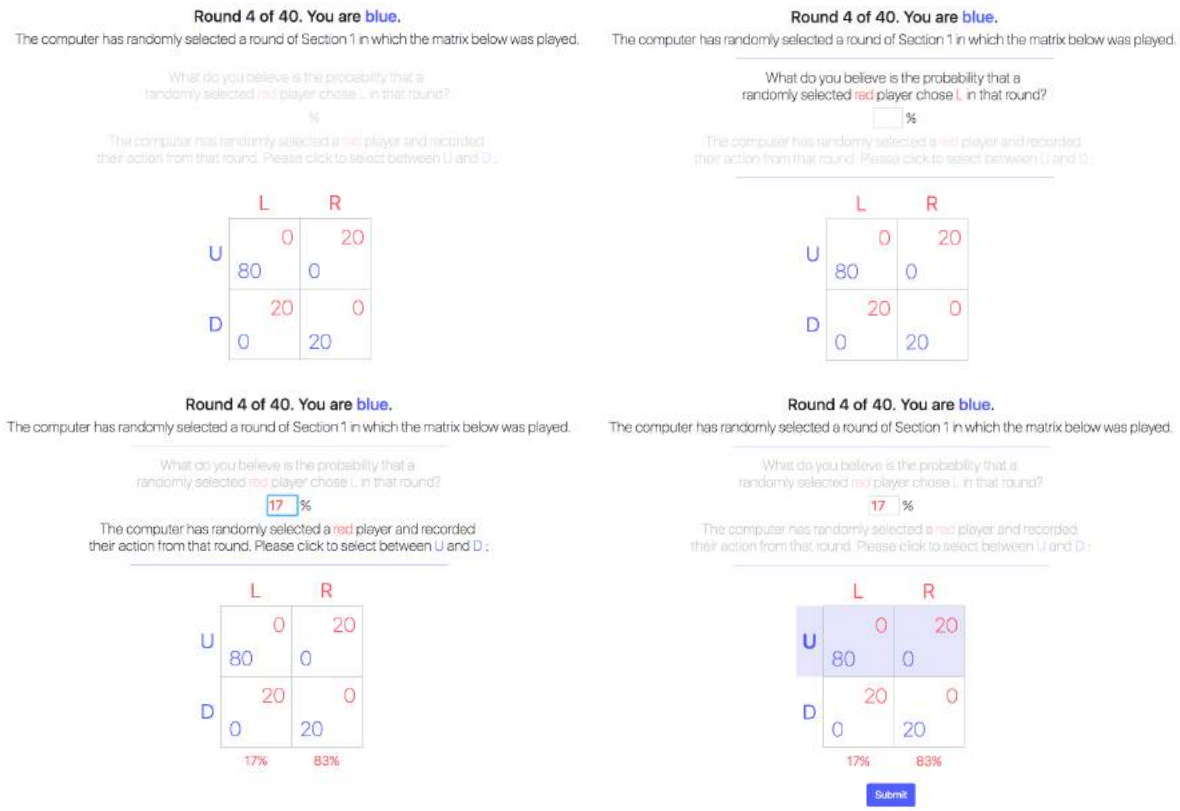

Figure 24: *Screenshots from second stage of [A,BA].* This figure shows an example round from the perspective of a player 1-subject (blue). At the start of the round, the subject sees the payoff matrix (top-left screen) and is told "The computer has randomly selected a round of Section 1 in which the below matrix was played." After 10 seconds pass, the text "What do you believe is the probability that a randomly selected red player chose L in that round?" darkens (top-right screen) indicating that the subject may state a belief. The subject enters a belief as a whole number between 0 and 100. Once the belief is entered, the corresponding probabilities appear below the matrix and the text "The computer has randomly selected a red player and recorded their action from that round. Please click to select between U and D:" darkens (bottom-left screen) indicating that the subject may take an action. Only after stating a belief may the subject select an action, but after the belief is stated, the subject may freely modify both his belief and action before submitting. After a belief is entered and an action is selected, the 'Submit' button appears (bottom-right screen).

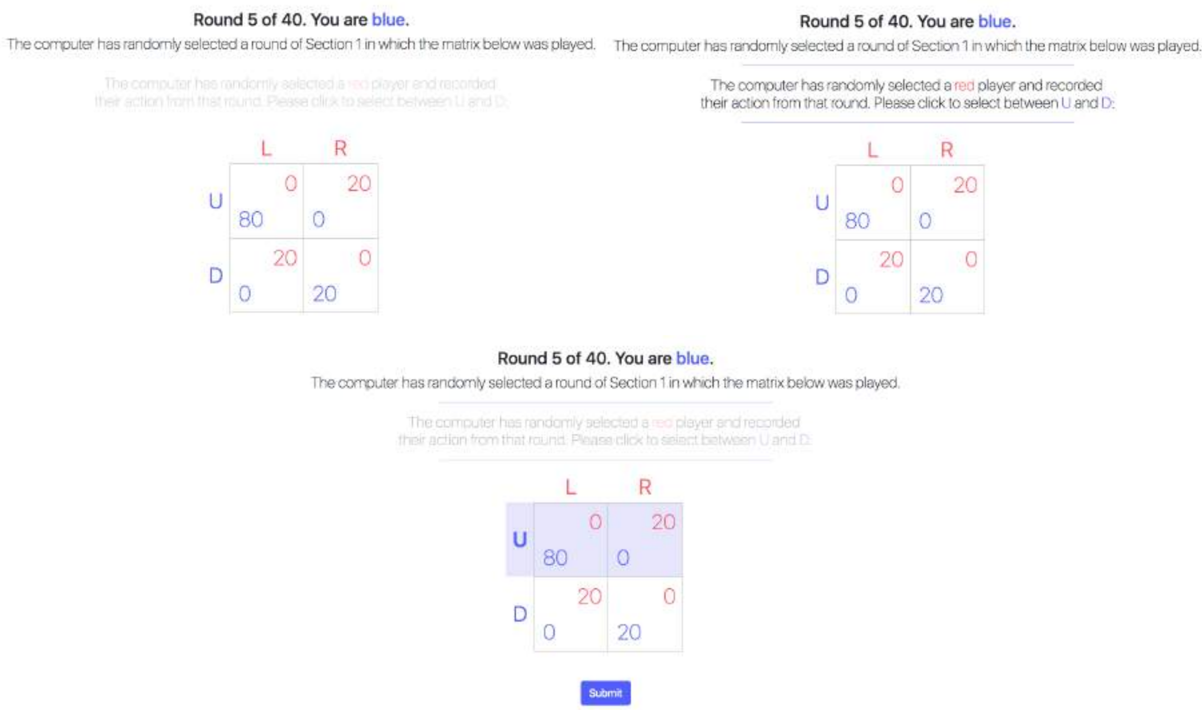

Figure 25: *Screenshots from second stage of [A,A]* The first stage [A,A] is identical to that of the [A,BA]. The second stage of the [A,A] is the same as that of [A,BA], except beliefs are not elicited.

11.6 Additional Figures

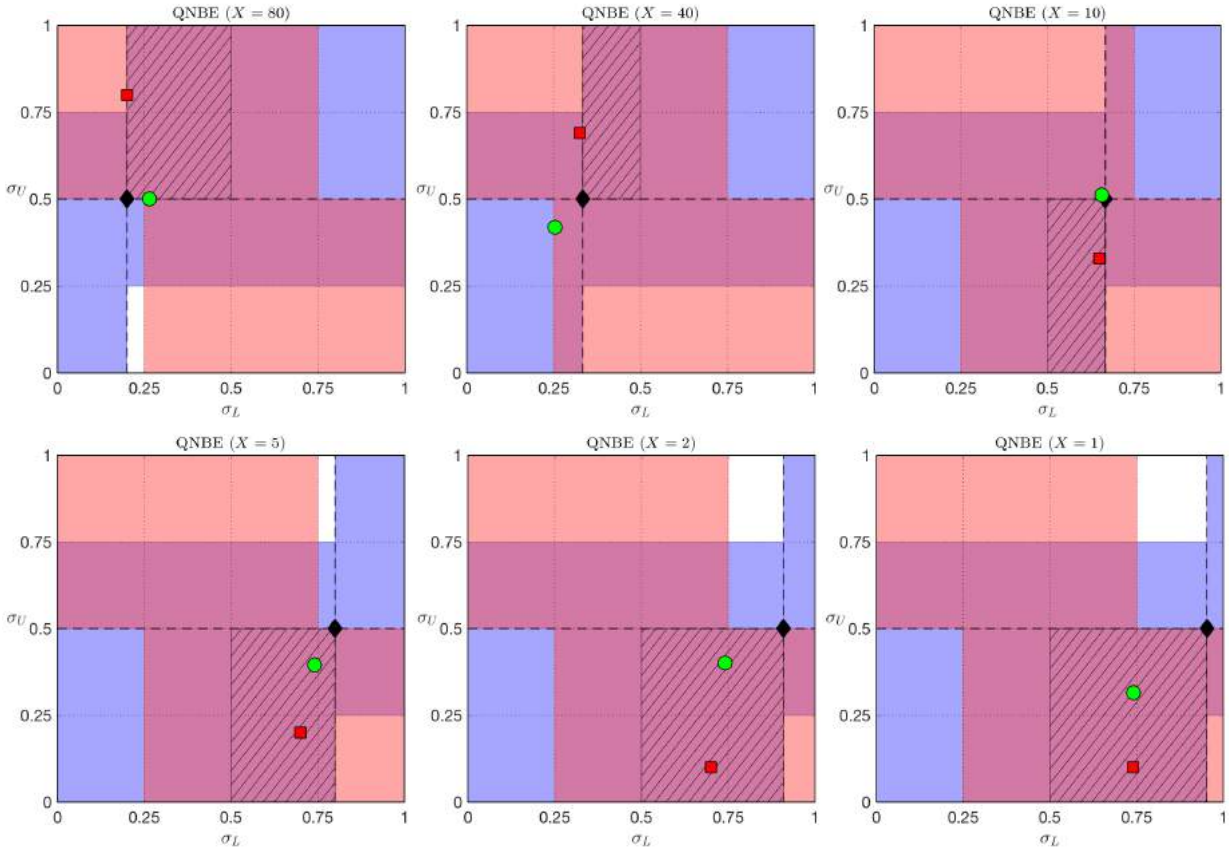

Figure 26: *QNBE and the data.* The green dot gives the empirical action frequencies from [A, ∘], the red square gives the median belief, and the black diamond is the Nash equilibrium.

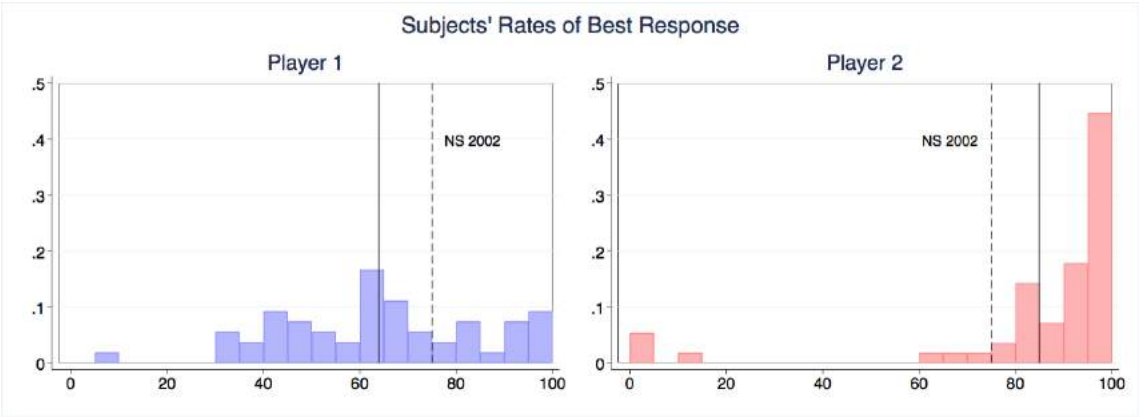

Figure 27: *Subjects' rates of best response.* This figure gives histograms of subjects' rates of best response across all X-games. The solid lines are averages, and the dashed lines in the bottom panel mark the average rate of best response from Nyarko and Schotter [2002].

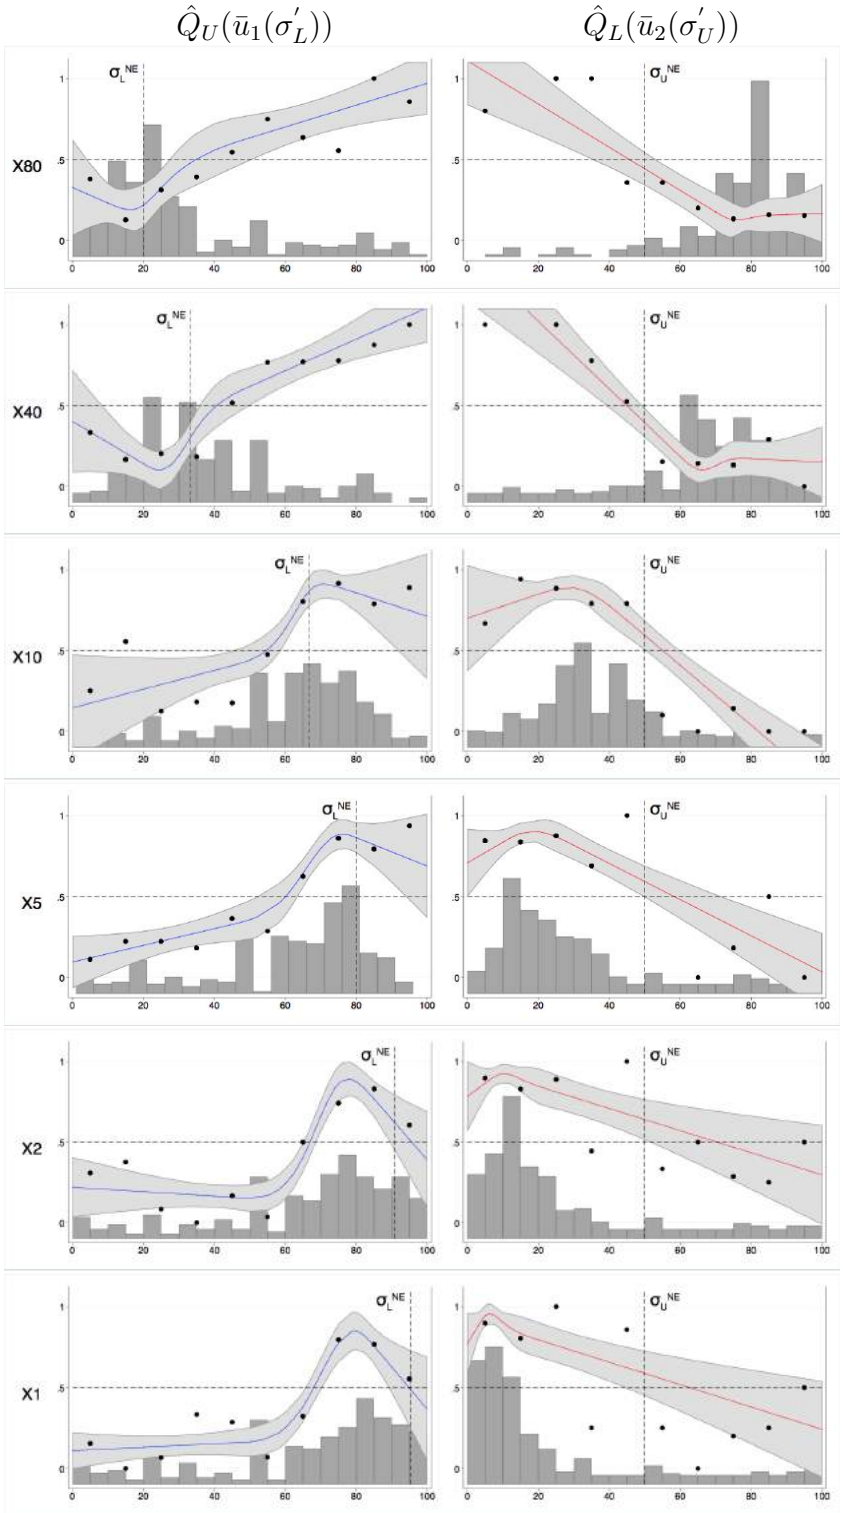

Figure 28: *Action frequencies predicted by beliefs.* We plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, std. errors clustered by subject) superimposed over belief histograms. The vertical dashed line is the indifferent belief, and the horizontal line is set to one-half.

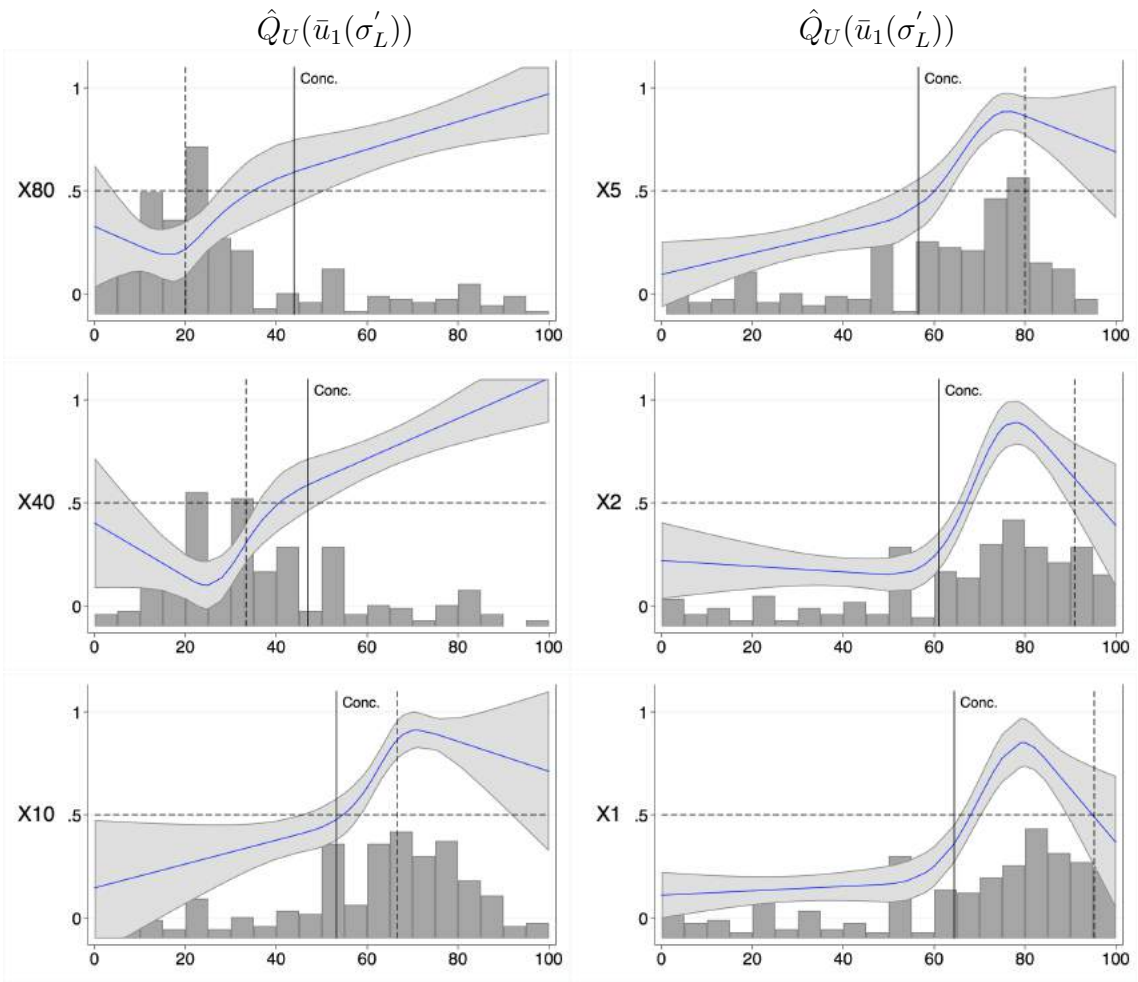

Figure 29: *Concave utility explains monotonicity failures.* For player 1 and each of the X-games, we plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms appear in gray, the vertical dashed line is the risk neutral indifferent belief $\sigma_j' = \sigma_j^{NE}$, and the horizontal line is set to one-half. The solid vertical line is the indifferent belief with concave utility that is estimated from fitting a single curvature parameter to all player 1-subjects' data.

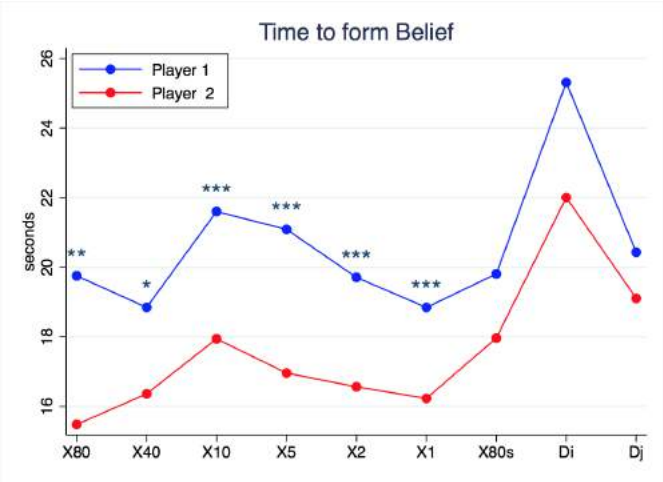

Figure 30: *Average time to form beliefs by game and player.* We plot the average time until stated beliefs are finalized by game and player role.

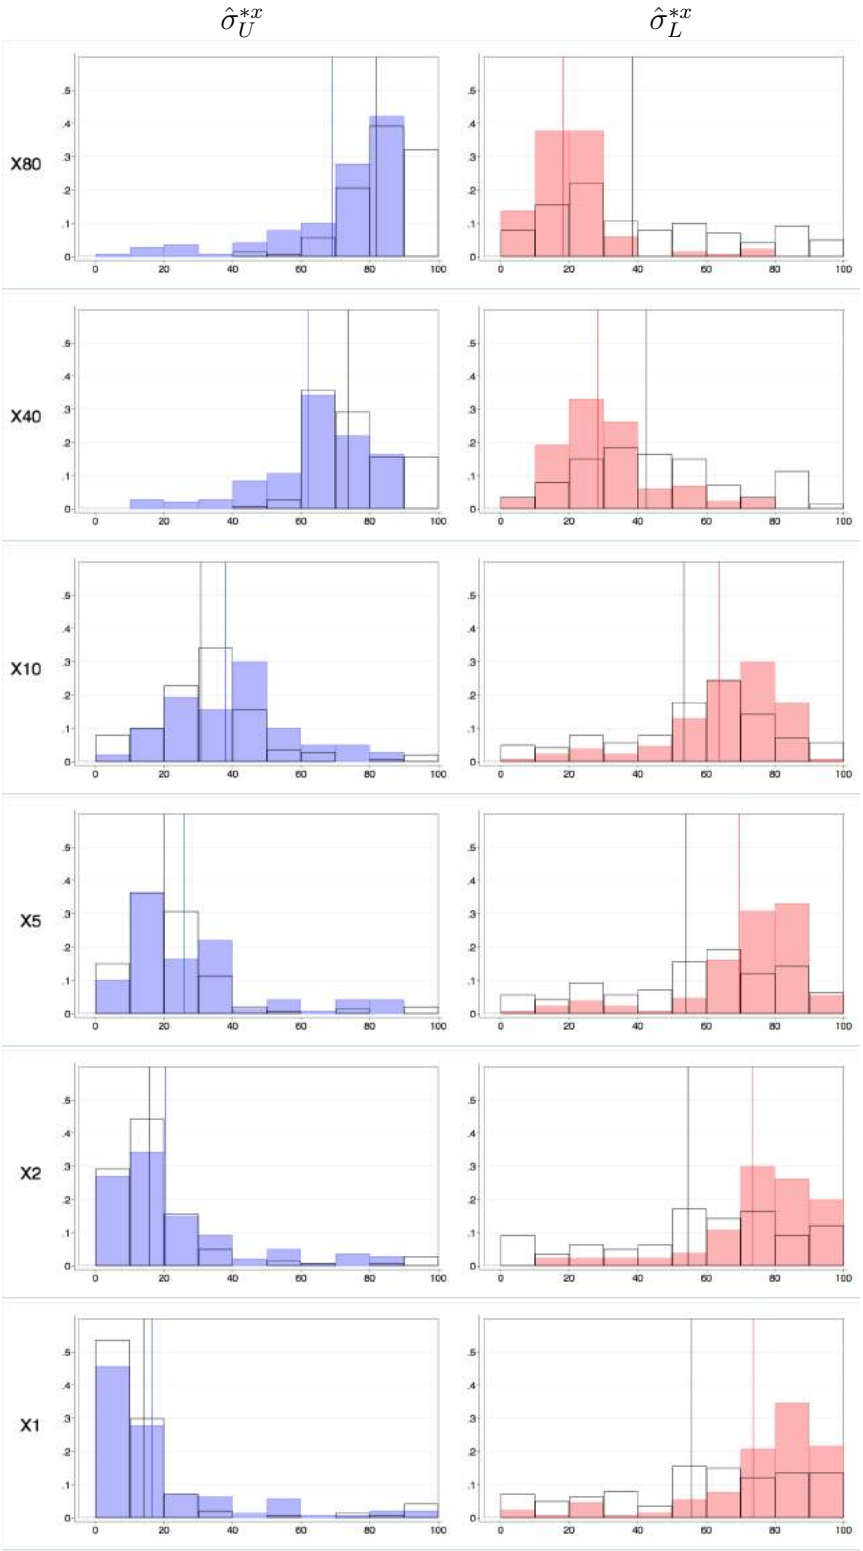

Figure 31: *Belief distributions by sophistication group.* The left panel is for player 2's beliefs about U, and the right panel is for player 1's beliefs about L. The colored histogram is for the high sophistication group, and the white histogram is for the low sophistication group. Mean beliefs are given as colored and black lines, for high and low sophistication, respectively.

11.7 Additional Tables

Player 1

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1	(7) all
best response rate	0.741***	0.737***	0.667***	0.600**	0.544	0.544	0.639***
	(0.000)	(0.000)	(0.000)	(0.026)	(0.356)	(0.414)	(0.000)
Observations	270	270	270	270	270	270	1620

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Player 2

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1	(7) all
best response rate	0.836***	0.857***	0.854***	0.836***	0.854***	0.857***	0.849***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	280	280	280	280	280	280	1680

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 15: *Rates of best response.* This table reports the average rates of best response by player and game. Significance is based on a two-sided t-test of the null hypothesis that the rate of best response is one-half. Standard errors are clustered by subject.

D1

	L	**R**
U	0 6	20 0
D	20 8	4 20

D2

	L	**R**
U	0 20	20 4
D	6 0	8 20

X80s

	L	**R**
U	0 8	2 0
D	2 0	0 2

R1

	L	**R**
U	0 5	20 0
D	10 0	0 20

R2

	L	**R**
U	0 5	20 0
D	40 0	0 20

Table 16: *Additional games.*

	(1) quintile	(2) equally spaced
very low beliefs	-0.006***	-0.006***
	(0.007)	(0.004)
low beliefs	-0.005***	-0.004***
	(0.009)	(0.010)
medium beliefs	-0.006***	-0.010***
	(0.000)	(0.002)
high beliefs	-0.009***	-0.010***
	(0.001)	(0.000)
very high beliefs	-0.005***	-0.005***
	(0.002)	(0.003)
Observations	1680	1680

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 17: *Fixed effect regressions of actions on beliefs–player 2, pooled across games.* For player 2, we pool together the data from all six games. In column 1, we divide the individual belief statements into quintiles–very low, low, medium, high, and very high beliefs. For each belief quintile, we run a separate linear regression of actions on beliefs that are both first demeaned by subtracting subject-specific averages. In column 2, we do the same thing, except the five belief groups are based on evenly spaced bins of 20 belief points. Standard errors are clustered by subject.

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1
med($\hat{\sigma}_U^*$) - $\hat{\sigma}_U$	30.000***	27.025***	-18.235***	-19.506***	-30.124***	-21.482***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
Observations	442	442	442	442	442	442
med($\hat{\sigma}_L^*$) - $\hat{\sigma}_L$	-6.506*	7.199*	-0.663	-4.086	-4.096	-0.096
	(0.079)	(0.050)	(0.418)	(0.345)	(0.172)	(0.433)
Observations	436	436	436	436	436	436
p-values in parentheses						
* $p < .1$, ** $p < .05$, *** $p < .01$						

	(1) X80	(2) X40	(3) X10	(4) X5	(5) X2	(6) X1
mean($\hat{\sigma}_U^*$) - $\hat{\sigma}_U$	25.511***	26.021***	-16.938***	-16.531***	-22.027***	-16.253***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)
Observations	442	442	442	442	442	442
mean($\hat{\sigma}_L^*$) - $\hat{\sigma}_L$	2.146	10.321**	-7.303	-12.615***	-10.441**	-9.670**
	(0.675)	(0.027)	(0.148)	(0.010)	(0.035)	(0.049)
Observations	436	436	436	436	436	436
p-values in parentheses						
* $p < .1$, ** $p < .05$, *** $p < .01$						

Table 18: *Bias in beliefs.* This table reports, for each player and game, the empirical bias in beliefs as measured by the difference between the median or mean belief statement and the empirical action frequency (expressed as percentages). In both cases, we report the p-values from two-sided tests of the null hypothesis that beliefs are unbiased. When using the median, p-values are bootstrapped in a way so as to preserve the within-subject correlation observed in the data. When using the mean, we use standard t-tests, clustering by subject.

	(1) $\hat{\beta}$(k≥2)	(2) $\hat{\beta}$(k≥2)	(3) $\hat{\beta}$(k≥2)	(4) $\hat{\beta}$(k≥2)	(5) $\hat{\beta}$(k≥2)	(6) $\hat{\beta}$(k≥2)
Player 1	22.465***	20.054***	20.626***	24.797***	22.461***	21.743***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Response time		0.729***	0.731***		0.828***	0.857***
		(0.002)	(0.001)		(0.001)	(0.001)
$\hat{\beta}$(k≥1)			0.286**			0.345**
			(0.048)			(0.038)
Constant	33.887***	17.841***	-6.973	33.150***	15.275**	-14.987
	(0.000)	(0.004)	(0.614)	(0.000)	(0.023)	(0.343)
Observations	110	110	110	93	93	93

p-values in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 19: *The sophistication gap.* In column 1, we regress $\hat{\beta}(k \geq 2)$ on an indicator for player 1. Column 2 controls for subject-average response time on the three rounds of Di (since sophistication is measured entirely with beliefs data, we use the time until stated beliefs are finalized). Column 3 additionally controls for $\hat{\beta}(k \geq 1)$. Columns 4-6 are the same, except we first drop subjects who ever took a dominated action in Di throughout the experiment.

	Player 1		Player 2	
	(1) $\hat{\sigma}_L^*$	(2) $\hat{\sigma}_U$	(3) $\hat{\sigma}_U^*$	(4) $\hat{\sigma}_L$
High soph. × X80	-20.172***	-0.297***	-12.607***	-0.000
	(0.000)	(0.006)	(0.000)	(1.000)
High soph. × X40	-14.359***	-0.304***	-11.636***	0.125
	(0.001)	(0.001)	(0.001)	(0.199)
High soph. × X10	10.107**	0.335***	7.321**	-0.018
	(0.032)	(0.001)	(0.041)	(0.871)
High soph. × X5	15.687***	0.309***	5.836	0.107
	(0.001)	(0.005)	(0.124)	(0.254)
High soph. × X2	18.808***	0.330***	4.607	0.018
	(0.000)	(0.003)	(0.239)	(0.857)
High soph. × X1	17.975***	0.231**	2.371	-0.018
	(0.001)	(0.025)	(0.529)	(0.856)
F-test	11.91***	8.53***	5.85***	0.51
	(0.000)	(0.000)	(0.000)	(0.801)
[d1,d2]	[6,323]	[6,323]	[6,335]	[6,335]
Observations	3240	648	3360	672

p-values in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Table 20: *Sophistication and behavior.* We regress beliefs or actions on indicators for all six X-games (omitted) and indicators for each of the six games interacted with an indicator for the high sophistication group. Columns 1-2 use player 1-subjects and columns 3-4 use player 2-subjects. We also report the results of F-tests of the hypothesis that all six coefficients are zero. Standard errors are clustered at the subject-game level.

Stochastic Equilibria: Noise in Actions or Beliefs?*

Evan Friedman†

November 17, 2019

JOB MARKET PAPER #2

Click here for LATEST VERSION

Abstract

Much is known about the empirical content of quantal response equilibrium (QRE) which relaxes the rationality requirement of Nash equilibrium by allowing for "noise in actions" while maintaining that beliefs are correct. By contrast, little is understood of the testable restrictions of equilibrium models which allow for "noise in beliefs" while maintaining best response. We introduce noisy belief equilibrium (NBE) for normal form games in which axioms restrict belief distributions to be *unbiased with respect to* and *responsive to changes in* the opponents' behavior. The axioms impose testable restrictions both within and across games, and we compare these restrictions to those of regular QRE (*Goeree et al. 2005*) in which axioms are placed on the quantal response function as the primitive. We show that NBE generates similar predictions as QRE such as the "own payoff effect", and yet is more consistent with the empirically documented effects of changes in payoff magnitude. Unlike QRE, NBE is a refinement of rationalizability and invariant to affine transformations of payoffs.

Keywords: beliefs; quantal response equilibrium; discrete choice models
JEL Classification: C72, C92, D84

*I thank Alessandra Casella, Yeon-Koo Che, Mark Dean, Ambuj Dewan, Frank Heinemann, Navin Kartik, Judd Kessler, RC Lim, Suanna Oh, Pietro Ortoleva, Tom Palfrey, Andrea Prat, Mike Woodford, and participants at the Columbia Microtheory Colloquium, Columbia Cognition and Decisions lab, Columbia Experimental Lunch, the NYU Student Microtheory Lunch, the Stony Brook Game Theory Conference, and the Summer School of the Econometric Society for helpful comments. I also thank the Columbia Experimental Laboratory in the Social Sciences (CELSS) for its generous support. This paper was previously circulated as "Noisy Belief Equilibrium". All errors are my own, and due to random perturbations in my beliefs.

†Email: ekf2119@columbia.edu

1 Introduction

Game theory rests on Nash equilibrium (NE) as its central concept, but despite its appeal and influence, it fails to capture the richness of experimental data. Systematic deviations from NE predictions have been documented, even in some of the simplest games.

NE rests on two assumptions. First, players form accurate beliefs over their opponents' actions. Second, players best respond to these beliefs. Efforts to reconcile theory with data typically amount to weakenings of these strict assumptions.

One leading example is *quantal response equilibrium* (QRE) (McKelvey and Palfrey [1995]), which is very much like NE, but relaxes the assumption of best response. That is, each player forms correct beliefs over the distribution of opponents' actions, and though he tends to take better actions (by expected utility), he fails to do so with probability one. Simply put, QRE is an equilibrium model with "noise in actions".

In many contexts, however, the assumption of correct beliefs is unrealistic. Therefore, it is natural to consider equilibrium models which relax the other condition of NE by allowing for "noise in beliefs" while maintaining best response. In this paper, we introduce such a model and by comparing it to QRE, we ask: which of action- or belief-noise is more consistent with experimental data?

Since we do not want our conclusions to depend on specific functional forms, we begin by introducing a general class of equilibrium models with noisy beliefs. In a *noisy belief equilibrium* (NBE), players best respond to their beliefs, but their beliefs are drawn from distributions that depend on the opponents' equilibrium behavior. The belief distributions are restricted to satisfy several axioms. The important behavioral axioms are *belief-responsiveness* and *unbiasedness*, which ensure that the belief distributions (1) tend to track changes in opponents' behavior and (2) are appropriately centered around the distribution of opponents' actions in equilibrium.

We study the testable restrictions of NBE, which we compare to those of *regular* QRE (Goeree et al. [2005]) in which axioms embed a *sensitivity to payoff differences* into the primitive quantal response function.[1] This is essentially the most flexible form of QRE which imposes testable restrictions on the data, and so we avoid altogether any concerns that QRE can "explain anything" (see, for example, Haile et al. [2008]).[2] Thus, we compare

[1]For each player, the quantal response function maps his vector of expected utilities (i.e. each element representing the expected payoff to some action) to a distribution over actions. The axioms impose that actions with higher payoffs are played more often (*monotonicity*), and that an increase in the payoff to some action increases the probability it is played (*responsiveness*).

[2]Haile et al. [2008] study *structural* QRE in which quantal response is induced by taking the action that

two families of stochastic equilibrium models, which inject noise into actions and beliefs, respectively.

While the idea of injecting noise into beliefs is not new (see *Related Literature* below), an approach that does not rely on parametric structure brings new insights. For example, some existing parametric models approximately satisfy our axioms and hence give predictions that can be approximated by NBE; and so our results have implications for understanding these models and their relationship to QRE.

In Section 2, we introduce NBE for normal form games and discuss the relationship of NBE to other concepts that relax the assumption of perfect beliefs. In particular, we show that NBE is a refinement of rationalizability (Bernheim [1984] and Pearce [1984]) in the sense that only rationalizable actions are played with positive probability in equilibrium. This distinguishes NBE from QRE, and yet we show that the models make similar predictions in certain types of fully mixed games.

In Section 3, we study the two empirical regularities explained by QRE that lie at the heart of its success. Specifically, in fully mixed games, QRE predicts (1) the commonly observed deviations from NE *within a game*, and (2) the well-known "own payoff effect" *across games*.[3] The best evidence for these regularities comes from generalized matching pennies games, so we specialize results for this context. We begin by showing an equivalence result: NBE imposes the same testable restrictions as QRE within any one of these games. We then show that NBE also predicts the own payoff effect across games. In other words, by adding noise to beliefs, it is *as if* players are sensitive to expected payoff differences–the mechanism behind QRE.

In Section 4, we revisit a sticking point for QRE, that it over-predicts sensitivity to changes in payoff magnitude. The problem is well-known for the parametric logit model: for fixed λ (rationality parameter), equilibrium predictions are sensitive to scaling one or more players' payoffs by positive constants. Such predictions have been tested experimentally by McKelvey et al. [2000], who find that subjects' behavior *within a game* is qualitatively consistent with logit, but the scaling predictions *across games* find little support as subjects' behavior is unaffected by scale. We provide novel results to establish that this "scaling

maximizes the sum of expected utility and a random error. They show that the data from any *one game* can be rationalized as a structural QRE as long as the errors are not i.i.d. across players' actions. On the other hand, the class of regular QRE does impose restrictions and is more general than the class of structural QRE with i.i.d. errors (Goeree et al. [2005]).

[3]Whereas NE predicts that a change in a player's own payoffs does not affect his behavior since the other players have to be kept indifferent, subjects' behavior is systematically affected by *non-affine* transformations of their payoffs. See, for example, Ochs [1995] and Goeree et al. [2003].

issue" is general to all regular QRE in the sense that, if QRE is to explain the empirical regularities discussed in the previous paragraph, it must be non-trivially sensitive to affine transformations of payoffs. In other words, QRE can be made *nearly* invariant to affine transformations of games, but only by being *nearly* insensitive to payoff differences within a game. By contrast, NBE explains the empirical regularities while being invariant to affine transformations, which is more consistent with experimental findings.

In Sections 5 and 6, we consider several datasets to test model predictions. Revisiting the McKelvey et al. [2000] study on scale effects, and using only the structure provided by the models' axioms, we show that NBE is a better qualitative description than QRE of the whole dataset. Both models capture deviations from NE *within* a game, but only NBE can explain the absence of scale effects and other patterns in behavior *across* games. After developing a parametric NBE model based on the logit transform, we compare its performance to that of logit QRE in data from several existing studies on 2×2 and 3×3 games. We find that the models perform similarly when fit to individual games *in-sample*, which is unsurprising due to our equivalence result. However, we show that NBE outperforms QRE in making *out-of-sample* predictions across games of varying scale and in fitting sets of games pooled together.

The paper is organized as follows. In the remainder of this section, we discuss related literature. Section 2 reviews QRE and introduces NBE, Section 3 compares the models' within-game restrictions and studies the own payoff effect, Section 4 establishes payoff magnitude predictions, Section 5 introduces the parametric logit transform NBE, Section 6 tests model predictions in data, and Section 7 concludes.

Related Literature. Early QRE theory was developed in a series of papers (McKelvey and Palfrey [1995], McKelvey and Palfrey [1998], Chen et al. [1997], and others) and is surveyed in a recent textbook (Goeree et al. [2016]). The logit specification was introduced in the original paper and has since found wide application in experimental studies where it is used to reconcile data with theoretical predictions.

Our task in this paper is to study and compare equilibrium models with noise in actions to those with noise in beliefs. For each type of noise, we select a representative family of models.

For noisy actions, we seek a family of QRE models that is both flexible and falsifiable, and so we choose *regular* QRE (Goeree et al. [2005]) in which axioms restrict the quantal response functions directly. The other alternative would have been the family of *structural* QRE (see, for example, Haile et al. [2008]) in which quantal response is induced by players

who choose actions that maximize the sum of expected utility and a random error. However, Haile et al. [2008] show that structural QRE can rationalize the data from any *one game* as long as the errors are not restricted to be i.i.d. across players' actions. By contrast, regular QRE imposes testable restrictions and is strictly more general than the family of structural QRE with i.i.d. errors.

For noisy beliefs, we develop a new model which we call noisy belief equilibrium (NBE). It is analogous to regular QRE in that its primitive, the mapping from opponents' actions to distributions over beliefs, is restricted to satisfy several axioms. Like regular QRE, flexibility in its primitive typically leads to set predictions; and by excluding a measure of possible outcomes, is falsifiable.

For injecting noise into equilibrium beliefs, NBE adapts the basic framework of random belief equilibrium (RBE) of Friedman and Mezzetti [2005] (no relation to the author of this paper). In their model, players best respond to beliefs that depend stochastically on the opponents' behavior, but as they study the case in which belief-noise "goes to zero" to develop a theory of equilibrium selection, their conditions on belief distributions do not impose any testable restrictions beyond ruling out weakly dominated actions. On the other hand, our paper is concerned with characterizing equilibria when belief-noise is bounded away from zero, so we introduce a new model and provide non-overlapping results. After introducing NBE, we give a more detailed discussion in Section 2.4.

Another model similar to NBE is sampling equilibrium of Rubinstein and Osborne [2003], which was applied to experimental data in Selten and Chmura [2008]. Sampling equilibrium is a parametric model of noisy beliefs, which approximately satisfies the NBE axioms and thus (up to technical conditions) is a special case of NBE (see Section 2.4 for details). Our results therefore suggest that it will behave similarly to QRE in certain datasets. Less related to NBE, but similar in spirit, Heller and Winter [2018] and Goncalves and Casanelles [2018] introduce equilibrium models with biased but deterministic beliefs.

We emphasize that NBE, as well as a number of other beliefs-based models, is invariant to affine transformations of payoffs. This is of interest because logit QRE is well-known to over-predict sensitivity to changes in scale (McKelvey et al. [2000]); and to address this "scaling issue", several parametric QRE models have been proposed.[4] However, we show that *all* regular QRE must be non-trivially sensitive to scaling and/or translating payoffs if they are to explain the two empirical regularities for which QRE is renowned. Hence, we argue

[4]Approaches include augmenting logit QRE with risk aversion (Goeree et al. [2003]) or heterogenous λs (McKelvey et al. [2000]); or endogenizing λ as a strategic decision (McKelvey et al. [1997] and Friedman [2018]).

that the scaling issue cannot be adequately addressed within the QRE framework. On the other hand, we show that beliefs-based models can explain the empirical regularities while being invariant to both scale and translation.

Our approach to mistaken beliefs can be contrasted with those that drop the equilibrium assumption altogether. Rationalizability (Bernheim [1984] and Pearce [1984]) is an early concept that allows for any belief not excluded by rationality and common knowledge of rationality.[5] Level k (Nagel [1995] and Stahl and Wilson [1995]) and its successors (Camerer et al. [2004], Alaoui and Penta [2015], and others) assume that subjects' beliefs are determined by their "depths of reasoning" or how many iterations of best response they can calculate. Mauersberger [2018] models beliefs as random draws from a Bayesian posterior.

2 Stochastic Equilibria

We provide the notation for normal form games, review QRE, introduce NBE, and discuss the relationship of NBE to other concepts.

2.1 Normal Form Games

A finite, normal form game $\Gamma = \{N, A, u\}$, is defined by a set of players $N = \{1, ..., n\}$, action space $A = A_1 \times ... \times A_n$ with $A_i = \{a_{i1}, ..., a_{iJ(i)}\}$ such that each player i has $J(i)$ possible pure actions ($J = \sum_i J(i)$ actions total), and a vector of payoff functions $u = (u_1, ..., u_n)$ with $u_i : A \to \mathbb{R}$.

Let Δ_i be the set of probability measures on A_i. Elements of Δ_i are of the form $\sigma_i : A_i \to \mathbb{R}$ where $\sum_{j=1}^{J(i)} \sigma_i(a_{ij}) = 1$ and $\sigma_i(a_{ij}) \geq 0$. For simplicity, set $\sigma_{ij} \equiv \sigma_i(a_{ij})$. Define $\Delta = \Delta_1 \times ... \times \Delta_n$ and $\Delta_{-i} = \times_{k \neq i} \Delta_k$ with typical elements $\sigma \in \Delta$ and $\sigma_{-i} \in \Delta_{-i}$. As is standard, extend payoff functions $u = (u_1, ..., u_n)$ to be defined over Δ via $u_i(\sigma) = \sum_{a \in A} \sigma(a) u_i(a)$. For convenience, we will call any element of Δ an "action", regardless of whether it is pure or mixed, and we use these terms only when the distinction is important.

2.2 Quantal Response Equilibrium

As is standard in the literature on *quantal response equilibrium* (QRE), we use additional notation for expected utilities. Given $\sigma_{-i} \in \Delta_{-i}$, player i's vector of expected utilities is given

[5] Goeree and Holt [2004] introduce *noisy rationalizability* in which "noise is injected into iterated conjectures about others' decisions and beliefs" that is better suited for application to experimental data.

by $\bar{u}_i(\sigma_{-i}) = (\bar{u}_{i1}(\sigma_{-i}), ..., \bar{u}_{iJ(i)}(\sigma_{-i})) \in \mathbb{R}^{J(i)}$ where $\bar{u}_{ij}(\sigma_{-i}) = u_i(a_{ij}, \sigma_{-i})$ is the expected utility to action a_{ij} given behavior of the opponents. We use $v_i = (v_{i1}, ..., v_{iJ(i)}) \in \mathbb{R}^{J(i)}$ as shorthand for an arbitrary vector of expected utilities. That is, v_i is understood to satisfy $v_i = \bar{u}_i(\sigma'_{-i})$ for some σ'_{-i}.

Player i's behavior is modeled via the *quantal response function* $Q_i = (Q_{i1}, ..., Q_{iJ(i)}) : \mathbb{R}^{J(i)} \to \Delta_i$, which maps his vector of expected utilities to a distribution over actions. For any $v_i \in \mathbb{R}^{J(i)}$, component $Q_{ij}(v_i)$ gives the probability assigned to action j. Intuitively, Q_i allows for arbitrary probabilistic mistakes in taking actions given the payoffs to each action, resulting perhaps from unmodeled costs of information processing.

To impose testable restrictions on data, we follow Goeree et al. [2005] by imposing the *regularity* axioms on the quantal response functions. Regularity defines a very important class of QRE models in that (1) regularity imposes testable restrictions on the data, and (2) all structural QRE[6] with i.i.d. errors (such as logit) are regular. In other words, the class of regular QRE is flexible enough to include the large majority of applications while still maintaining empirical content (i.e. the Haile et al. [2008] critique simply does not apply); so we impose the axioms throughout:

Assumption 1. *Quantal response function Q satisfies (A1)-(A4).*

(A1) **Interiority**: $Q_{ij}(v_i) \in (0, 1)$ for all $j \in 1, ..., J(i)$ and for all $v_i \in \mathbb{R}^{J(i)}$.

(A2) **Continuity**: $Q_{ij}(v_i)$ is a continuous and differentiable function for all $v_i \in \mathbb{R}^{J(i)}$.

(A3) **Responsiveness**: $\frac{\partial Q_{ij}(v_i)}{\partial v_{ij}} > 0$ for all $j \in 1, ..., J(i)$ and for all $v_i \in \mathbb{R}^{J(i)}$.

(A4) **Monotonicity**: $v_{ij} > v_{ik} \implies Q_{ij}(v_i) > Q_{ik}(v_i)$ for all $j, k \in 1, ..., J(i)$.

Responsiveness and *monotonicity* are the important behavioral axioms, and can be summarized as "sensitivity to payoff differences". These require that an all-else-equal increase in the payoff to some action increases the probability it is played, and that actions with higher payoffs are played with greater probability. The other axioms are technical in nature, ensuring existence and that all actions are played with positive probability.

A QRE is obtained when the distribution over all players' actions is consistent with their quantal response functions. Letting $Q = (Q_1, ..., Q_n)$ and $\bar{u} = (\bar{u}_1, ..., \bar{u}_n)$, QRE is any fixed point of the composite function $Q \circ \bar{u} : \Delta \to \Delta$.

Definition 1. Fix $\{\Gamma, Q\}$. A *QRE* is any $\sigma \in \Delta$ such that for all $i \in 1, ..., n$ and all $j \in 1, ..., J(i)$, $\sigma_{ij} = Q_{ij}(\bar{u}_i(\sigma_{-i}))$.

[6] In a structural QRE, player i chooses the action that maximizes the sum of expected utility and a random error, and thus $Q_{ij}(v_i) = \mathbb{P}(v_{ij} + \varepsilon_{ij} \geq v_{ik} + \varepsilon_{ik} \ \forall k)$.

2.3 Noisy Belief Equilibrium

In a *noisy belief equilibrium* (NBE), players draw beliefs about their opponents' actions probabilistically to which they best respond. This induces, for each player, an expected action. In equilibrium, the belief distributions are centered in some sense around the opponents' expected actions, which are similarly induced by best responding to realized beliefs.

Randomness in beliefs can be interpreted in several ways. It could result from mistakes in "solving" for an equilibrium or from noisy signals about opponents' behavior, possibly from observing random samples of their play. It could also be that there is a large population of subjects who form beliefs deterministically, and the distribution of beliefs reflects heterogeneity in the population.

2.3.1 An example

Before defining NBE for normal form games in the next section, we introduce it using our leading example: the family of generalized matching pennies games. Consisting of all 2×2 games with unique fully mixed NE, this family has been the subject of numerous experimental studies. The NBE of these games take a simple form, allowing us to introduce key ideas concisely.

Generalized matching pennies is defined by the payoff matrix in Figure 1.[7] We use the notation Γ^m for an arbitrary game in this family. The parameters a_L, a_R, b_U, and b_D give the base payoffs. The parameters c_L, c_R, d_U, and d_D are the payoff differences, which we assume are strictly positive to maintain the relevant features.[8] Since each player has only two actions in Γ^m, we identify Δ_i with $[0,1]$ and Δ with $[0,1]^2$. We also write σ_U and σ_L for the probabilities of playing U and L, respectively. As has been our convention, we refer to σ_U and σ_L as "actions" even though they are understood to be probabilities. Note that the Nash equilibrium $\{\sigma_U^{NE}, \sigma_L^{NE}\} = \{\frac{d_D}{d_U+d_D}, \frac{c_R}{c_L+c_R}\}$ depends only on the payoff differences.

More generally, for any game in which player k has two pure actions (i.e. $J(k) = 2$), we use $r \in [0,1]$ to refer to player k's action. This is simply to avoid using subscripts. In Γ^m, for example, r should be understood as one of σ_U or σ_L depending on the context.

In Γ^m, or any game in which $J(k) = 2$, player k's action is $r \in [0,1]$. We assume that player i's belief over k's action is drawn from a distribution that depends on r. In other words, player i's belief is a random variable that we denote $r^*(r)$, which is supported

[7] This notation is borrowed from Selten and Chmura [2008] with slight modification.

[8] Games in which the payoff differences are all strictly negative are equivalent up to the labelling of actions.

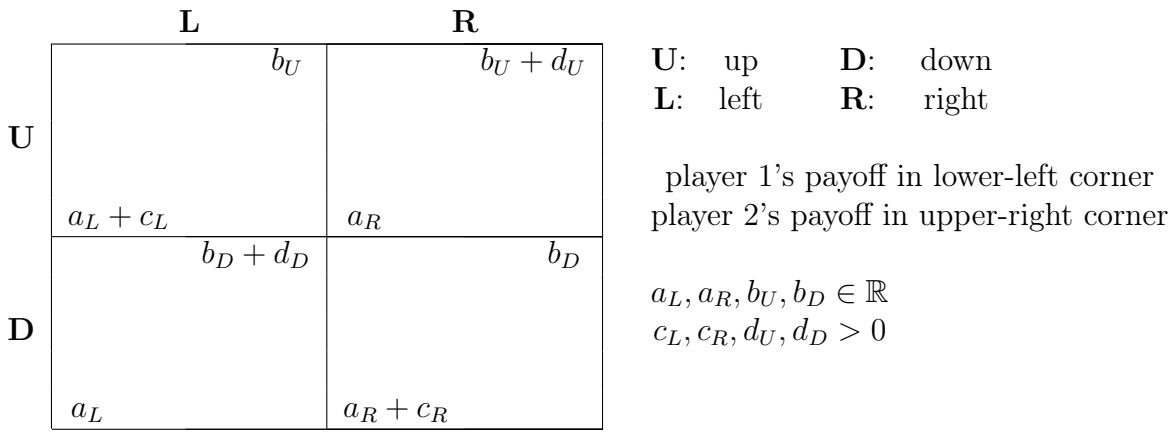

Figure 1: *Generalized matching pennies.*

on $[0,1]$. We call this family of random variables the *belief-map* (following Friedman and Mezzetti [2005]), and it is defined by a family of CDFs: for any potential belief $\bar{r} \in [0,1]$, $F_k^i(\bar{r}|r)$ is the probability of realizing a belief less than or equal to $\bar{r}$ given that player k is playing r.

After realizing belief r', player i takes an action. This is summarized by a *strategy* $s_i = (s_{i1}, s_{i2})$ in which component $s_{ij} : [0,1] \to [0,1]$ is a measurable function mapping realized beliefs to the probability of taking action a_{ij} (s_i must satisfy $s_{i1}(r') + s_{i2}(r') = 1$ for all $r' \in [0,1]$). Without loss, action a_{ij} is a best response to any belief in $[0, \bar{r}_i]$ and action a_{il} is a best response to any belief in $[\bar{r}_i, 1]$, where $\{\bar{r}_1, \bar{r}_2\} = \{\sigma_L^{NE}, \sigma_U^{NE}\}$ are the *indifferent beliefs* that correspond to the Nash equilibrium of Γ^m. We say that s_i is *rational* if it indicates a best response to any realized belief: $s_{ij}(r') = 1$ for $r' < \bar{r}_i$ and $s_{ij}(r') = 0$ for $r' > \bar{r}_i$. Since any $s_{ij}(\bar{r}_i) \in [0,1]$ is a best response, there are many rational strategies (all of which agree on $r' \neq \bar{r}_i$). We define player i's *expected best response correspondence* or *reaction correspondence* as the *set of* expected actions that can be induced by best response (i.e. integrating over any rational strategy) as a function of player k's action r:

$$\Psi_{ij}(r) = \{ \int_{[0,1]} s_{ij}(r') dF_k^i(r'|r) : \ s_i \text{ is rational} \}. \tag{1}$$

Before defining equilibrium, we restrict the belief-map to satisfy several axioms. We introduce them here for the binary-action case and will generalize them to allow for arbitrary numbers of actions in the next section. The axioms will simplify the form of the reaction (1) and be shown to impose restrictions both within and across games:

Assumption 2. *If $J(k) = 2$, the belief-map r^* satisfies (B1′)-(B4′).*

(B1′) **Interior full support**: For any $r \in (0,1)$, $F_k^i(\bar{r}|r)$ is strictly increasing and continuous in $\bar{r} \in [0,1]$; $r^*(0) = 0$ and $r^*(1) = 1$ with probability 1.[9]

(B2′) **Continuity**: For any $\bar{r} \in (0,1)$, $F_k^i(\bar{r}|r)$ is continuous in $r \in [0,1]$.

(B3′) **Belief-responsiveness**: For all $r < r' \in [0,1]$, $F_k^i(\bar{r}|r') < F_k^i(\bar{r}|r)$ for $\bar{r} \in (0,1)$.

(B4′) **Unbiasedness**: $F_k^i(r|r) = \frac{1}{2}$ for $r \in (0,1)$.

Axioms (B1′) and (B2′) are technical in nature and will be shown to ensure existence of equilibria and that the other axioms are well-defined. (B3′) restricts belief distributions to be *responsive* to changes in the opponent's behavior, (B4′) restricts belief distributions to be *unbiased* with respect to the opponent's action, and both axioms are required to meaningfully restrict the set of equilibrium outcomes. We explain each axiom in turn.

Interior full support (B1′) requires that belief distributions are atomless and have full support when the opponent's action is interior, i.e. for $r \in (0,1)$.[10] It further imposes that beliefs are correct with probability one (and therefore described by a single atom) when the opponent's action is on the boundary, i.e. for $r \in \{0,1\}$. Otherwise, beliefs would necessarily be biased.[11] (B1′) and the structure of Γ^m make the form of the reaction (1) particularly simple. Since the indifferent belief $\bar{r}_i$ is interior, it realizes with probability zero for all $r \in [0,1]$. Since all rational strategies agree on $r' \neq \bar{r}_i$, the reactions are single-valued functions indicating the probabilities with which U (for player 1) and L (for player 2) are best responses to realized beliefs:

$$\begin{aligned}\Psi_U(\sigma_L) \equiv& 1 - F_2^1(\sigma_L^{NE}|\sigma_L)\\ \Psi_L(\sigma_U) \equiv& F_1^2(\sigma_U^{NE}|\sigma_U).\end{aligned}$$

Continuity (B2′) implies that $\Psi = (\Psi_U, \Psi_L) : [0,1]^2 \to [0,1]^2$ is continuous in (σ_U, σ_L), which will ensure existence of equilibria. It is easy to show that (B1′) and (B2′) together imply that Ψ is jointly continuous in (σ_U, σ_L) and the payoff parameters, which will ensure that equilibria do not jump for small changes in the game. Importantly, despite that the reactions are continuous, we make the following remark:

Remark 1. There are *discontinuities in beliefs*: there exists a (Borel) subset of $[0,1]$ for which

[9] This is equivalent to having CDFs that satisfy $F_k^i(\bar{r}|0) = 1$ and $F_k^i(\bar{r}|1) = \mathbf{1}_{\{\bar{r}=1\}}$ for $\bar{r} \in [0,1]$.

[10] Even though beliefs have full support in this case, the probability that beliefs realize in any open subset of $[0,1]$ can still be made arbitrarily small; in this sense the axiom is very weak.

[11] If beliefs are not correct with probability one when $r \in \{0,1\}$, beliefs would be biased on *mean*, and if they are correct with probability less than one half, than they would be biased on *median*.

the probability that player i's beliefs realize in that set is *dis*continuous in k's action r.

For instance, by (B1$'$), the probability that belief $r' = 0$ realizes jumps from 0 to 1 as the opponent's action approaches 0, i.e. as $r \to 0^+$. More generally, there are other belief-discontinuities associated with other (Borel) sets of realized beliefs. From (B1$'$) and (B2$'$), it is easy to characterize all belief-discontinuities.[12] Intuitively, since beliefs have full support and are atomless for all interior r, but are correct with probability one when r is on the boundary, all discontinuities are related to sets of realized beliefs nearby one of the boundaries and r approaching that same boundary. In Γ^m, the reactions are continuous despite the existence of belief-discontinuities because the probability that beliefs realize in any of the "relevant" sets–the largest which induce unique best responses–are continuous in r. This issue of belief-discontinuities will require special care when generalizing NBE to arbitrary games, but we will still find that the reactions are continuous in generic games.

Belief-responsiveness (B3$'$) ensures that belief distributions shift in the same direction as changes in the opponent's action. To capture this idea, we use the notion of first-order stochastic dominance (FOSD). Importantly, (B3$'$) implies that $\Psi_U(\sigma_L)$ and $\Psi_L(\sigma_U)$ are strictly increasing and strictly decreasing respectively, which will imply a unique equilibrium in Γ^m.

Unbiasedness (B4$'$) imposes that beliefs are correct *on median*. An implication of (B4$'$) in Γ^m is that if player k is playing the *indifferent action* that equalizes the (objective) expected utility to both of player i's actions, the probability of taking either action is exactly one half. On the other hand, replacing (B4$'$) with *mean*-unbiasedness would place no restriction on player i's reaction when k plays the indifferent action. Nonetheless, it may still be of interest to impose mean-unbiasedness in some applications, so we note that it is consistent with (B4$'$) (and the other axioms) and therefore could be imposed in addition. In Section 2.4, we discuss how (B4$'$) can be microfounded via a model in which players sample their opponents' play.

Definition 2. Fix $\{\Gamma^m, \sigma^*\}$. An NBE is any $(\sigma_U, \sigma_L) \in [0,1]^2$ such that $\Psi_U(\sigma_L) = \sigma_U$ and $\Psi_L(\sigma_U) = \sigma_L$.

[12] *Characterizing belief-discontinuities* ($J(k) = 2$). We discuss the case that r is nearby 0, with the case of r nearby 1 being symmetric. Let $\mu_k^i(\cdot|r)$ be the probability measure on $[0,1]$ derived from $F_k^i(\cdot|r)$. From (B1$'$) and (B2$'$), it is easy to check that $\mu_k^i(\cdot|\cdot)$ satisfies (1) $\mu_k^i(\{0\}|r) = 0$ for $r > 0$, (2) $\mu_k^i(\{0\}|0) = 1$, (3) $\mu_k^i([0,\epsilon)|r)$ is continuous in $r \in [0,1]$, and (4) $\mu_k^i([0,\epsilon)|r) \to 1$ as $r \to 0^+$. Hence, there are discontinuities in $\mu_k^i(\{0\}|r)$ and $\mu_k^i((0,\epsilon)|r)$ as $r \to 0^+$, which jump from 0 to 1 and 1 to 0, respectively. More generally, letting $A, B_k \subset [0,1]$ be Borel subsets, there is a discontinuity in $\mu_k^i(B_k|r)$ as $r \to 0^+$ if and only if $B_k = \{0\} \cup A$ or $B_k = (0,\epsilon) \cup A$ where A is well-separated from 0 (i.e. $cl(A) \cap \{0\} = \emptyset$).

Theorem 1. *Fix* $\{\Gamma^m, \sigma^*\}$. *An NBE exists and is unique and interior.*

Proof. See Appendix 8.1. □

2.3.2 Normal Form Games

We generalize NBE to normal form games. To this end, we adapt the framework of random belief equilibrium (RBE) (Friedman and Mezzetti [2005]), but restrict the belief distributions to satisfy axioms in order to impose testable restrictions on the data. The general axioms nest their binary-action counterparts from the previous section.

Given player k's action $\sigma_k \in \Delta_k$, player i's beliefs over k's action are given by random vector $\sigma_k^{i*}(\sigma_k) = (\sigma_{k1}^{i*}(\sigma_k), ..., \sigma_{kJ(k)}^{i*}(\sigma_k))$ that is supported on Δ_k. We call this family of random vectors player i's *belief-map* over player k's action. For convenience, refer to all players' belief-maps by $\sigma^* \equiv (\sigma_k^{i*})_{i,k\neq i}$, and for all $\sigma_{-i} \in \Delta_{-i}$, define belief-maps over i's opponents' actions by $\sigma_{-i}^*(\sigma_{-i}) \equiv (\sigma_k^{i*}(\sigma_k))_{k\neq i}$.

For each $\sigma_k \in \Delta_k$, $\sigma_k^{i*}(\sigma_k)$ is defined by probability measure $\mu_k^i(\cdot|\sigma_k)$ on $\mathcal{B}(\Delta_k)$, the Borel σ-algebra on Δ_k, which gives the probability that beliefs are realized in any $B_k \in \mathcal{B}(\Delta_k)$. Assume that all of k's opponents draw their beliefs about k independently conditional on σ_k, and player i's beliefs about any two of his opponents are drawn independently conditional on their actions. Thus, for all $\sigma_{-i} \in \Delta_{-i}$, $\sigma_{-i}^*(\sigma_{-i})$ is associated with the product measure: $\mu_{-i}(B|\sigma_{-i}) \equiv \Pi_{k\neq i}\mu_k^i(B_k|\sigma_k)$ for any $B = \times_{k\neq i}B_k \in \otimes_{k\neq i}\mathcal{B}(\Delta_k) \equiv \mathcal{B}(\Delta_{-i})$.

Define the *ij-response set* $R_{ij} \subseteq \Delta_{-i}$ by

$$R_{ij} = \{\sigma_{-i}^{'} : \ \bar{u}_{ij}(\sigma_{-i}^{'}) \geq \bar{u}_{ik}(\sigma_{-i}^{'}) \ \forall k = 1, ..., J(i)\}. \tag{2}$$

This defines the set of beliefs about i's opponents for which action a_{ij} is a best response.[13] A *strategy* for player i is a measurable function $s_i = (s_{i1}, ..., s_{iJ(i)}) : \Delta_{-i} \to \Delta_i$ where for all $\sigma_{-i}^{'} \in \Delta_{-i}$, $s_{ij}(\sigma_{-i}^{'}) \geq 0$ and $\sum_{j=1}^{J(i)} s_{ij}(\sigma_{-i}^{'}) = 1$. As before, this maps any realized belief to an action. Strategy s_i is *rational* if it only puts positive probability on best responses: $s_{ij}(\sigma_{-i}^{'}) = 0$ if $\sigma_{-i}^{'} \notin R_{ij}$.

Given any $\sigma_{-i} \in \Delta_{-i}$, player i's belief-map σ_{-i}^* induces a distribution over his realized beliefs and thus also over his actions through his strategy s_i. Integrating i's strategy over his realized beliefs via the measure $\mu_{-i}(\cdot|\sigma_{-i})$ gives an expected action. Restricting attention to rational strategies in which player i best responds to realized beliefs, we define player i's

[13] Note that $R_{ij} \in \mathcal{B}(\Delta_{-i})$, i.e. that the response sets are measurable.

expected best response correspondence or *reaction correspondence* as

$$\Psi_i(\sigma_{-i};\sigma^*_{-i}) \equiv \{\int_{\Delta_{-i}} s_i(\sigma'_{-i})d\mu_{-i}(\sigma'_{-i}|\sigma_{-i}) : \ s_i \text{ is rational}\}. \tag{3}$$

This maps the opponents' action profile σ_{-i} to the set of i's expected actions that can be induced by best responding to realized beliefs.

Correspondence (3) generalizes the best response correspondence of NE, and analogous to NE, NBE will be defined as a fixed point of $(\Psi;\sigma^*) = ((\Psi_1;\sigma^*_{-1}),...,(\Psi_n;\sigma^*_{-n})) : \Delta \rightrightarrows \Delta$. Note that while the belief distributions depend on the opponents' expected actions, the dependence is arbitrary without additional restrictions on σ^*.

We generalize axioms (B1′)-(B4′) from the previous section to allow for arbitrary numbers of actions. Our general technical axioms *interior full support* and *continuity* require that belief distributions (1) are supported on the lower dimensional simplex over opponents' pure actions that are played with positive probability and (2) vary with the opponents' actions as continuously as possible, given the previous point. As in the binary-action case, these conditions are necessary to accommodate our behavioral axioms but imply that the belief distributions involve discontinuities associated with opponents' actions nearby the boundary. However, as in the matching pennies example, the reactions of which NBE is a fixed point will be continuous in generic games (and upper hemicontinuous for all games).

To state the axioms, we require additional notation. For any action $\sigma_k \in \Delta_k$, define $\Delta(\sigma_k) \equiv \{\sigma'_k \in \Delta_k : supp(\sigma'_k) = supp(\sigma_k)\}$ as the subset of Δ_k in which each element is a probability vector that has the same support as σ_k (i.e. has 0s in precisely the same components as σ_k). For example, if $\sigma_k = (0, \frac{1}{2}, \frac{1}{2})$, then $\Delta(\sigma_k) = \{(0, p, 1-p) : p \in (0,1)\}$. Let $\langle \Delta_k, \mathcal{B}(\Delta_k), \mathcal{L}_k \rangle$ be the Lebesgue probability space on Δ_k where $\mathcal{L}_k$ is the Lebesgue measure. For each σ_k, we also define the Lebesgue probability space $\langle \Delta(\sigma_k), \mathcal{B}(\Delta(\sigma_k)), \mathcal{L}_k^{\Delta(\sigma_k)} \rangle$, where $\mathcal{B}(\Delta(\sigma_k))$ is the Borel σ-algebra on $\Delta(\sigma_k)$ and $\mathcal{L}_k^{\Delta(\sigma_k)}$ is the Lebesgue measure on $\Delta(\sigma_k)$.[14] Note that if σ_k has 0 in some component and $A \in \mathcal{B}(\Delta(\sigma_k))$, then $\mathcal{L}_k(A) = 0$. For example, if $\sigma_k = (0, \frac{1}{2}, \frac{1}{2})$, $\mathcal{L}_k(\Delta(\sigma_k)) = 0$ even though $\mathcal{L}_k^{\Delta(\sigma_k)}(\Delta(\sigma_k)) = 1$. We now state our technical axioms:

(B1) **Interior full support**: $\mu^i_k(B_k|\sigma_k) > 0$ if and only if $\mathcal{L}_k^{\Delta(\sigma_k)}(B_k \cap \Delta(\sigma_k)) > 0$.

(B2) **Continuity**: Let $\{\sigma^t_k\} \subset \Delta_k$ be a sequence with $\sigma^t_k \to \sigma^\infty_k$ as $t \to \infty$. $\lim_{t\to\infty} \mu^i_k(B_k|\sigma^t_k) = \mu^i_k(B_k|\sigma^\infty_k)$ if, for sufficiently large t, either *(i)* $\{\sigma^t_k\} \subset \Delta(\sigma^\infty_k)$ or *(ii)* $cl(B_k \cap \Delta(\sigma^t_k)) \cap \Delta(\sigma^\infty_k) =$

[14] If $\sigma'_k \in \Delta(\sigma_k)$, then $\Delta(\sigma'_k) = \Delta(\sigma_k)$, so this defines only finite probability spaces.

$B_k \cap \Delta(\sigma_k^\infty)$ up to $\mathcal{L}_k^{\Delta(\sigma_k^\infty)}$-measure 0.[15]

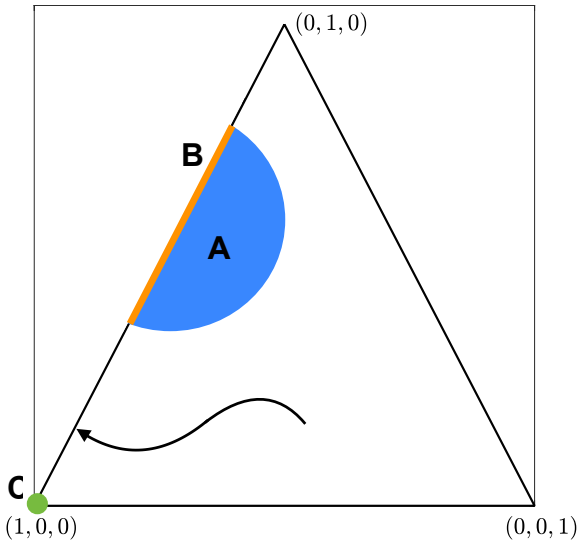

Figure 2: *Technical axioms: an example with 3 pure actions.* This figure plots the simplex that defines player k's action when he has 3 pure actions. Consider the sets A, B, and C. A is entirely in the interior of the simplex, but $cl(A) \cap B = B$, where B is a subset of $\{(p, 1-p, 0) : p \in (0,1)\}$, and $C = \{(1,0,0)\}$. Now consider the sequence $\{\sigma_k^t\}_t$ drawn as the black arrow, which starts from the interior and limits to $\sigma_k^\infty \in \{(p, 1-p, 0) : p \in (0,1)\}$. By (B1), for $t < \infty$, $\mu_k^i(A|\sigma_k^t) > 0$ and $\mu_k^i(B|\sigma_k^t) = \mu_k^i(C|\sigma_k^t) = 0$. Also by (B1), $\mu_k^i(A|\sigma_k^\infty) = 0$, $\mu_k^i(B|\sigma_k^\infty) > 0$, and $\mu_k^i(C|\sigma_k^t) = 0$. Thus, there is a discontinuity: $\mu_k^i(B|\sigma_k^t) = 0$ for $t < \infty$, but $\mu_k^i(B|\sigma_k^\infty) > 0$. However, by (B2), there is not a discontinuity in $\mu_k^i(A \cup B|\sigma_k^t)$ as $t \to \infty$ because A and B overlap in the sense of (B2)-*(ii)*.

Interior full support says that (1) the support of belief distributions is the subset of the simplex (over the opponent's pure actions) whose elements put positive probability on the opponent's pure actions that are played with positive probability, and (2) the belief measure is absolutely continuous with respect to the relevant Lebesgue measure. Suppose the opponent's action is σ_k. Then, $\Delta(\sigma_k)$ is the set of beliefs over k's action that put positive probability precisely on the pure actions that k plays with positive probability. The axiom says that the probability beliefs realize in B_k will be positive if and only if there is a nontrivial intersection of B_k with $\Delta(\sigma_k)$. The "only if" part implies that there are no atoms, unless the

[15] That is, set equality may only hold as $\{cl(B_k \cap \Delta(\sigma_k^t)) \cap \Delta(\sigma_k^\infty)\} \cup c_1 = \{B_k \cap \Delta(\sigma_k^\infty)\} \cup c_2$ for some $c_1, c_2 \in \Delta(\sigma_k^\infty)$ with $\mathcal{L}_k^{\Delta(\sigma_k^\infty)}(c_1) = \mathcal{L}_k^{\Delta(\sigma_k^\infty)}(c_2) = 0$.

opponent is taking a pure action with probability one in which case beliefs are correct with probability one.

Continuity is best understood by contrast with a more standard notion. It is similar to requiring that, for any sequence $\sigma_k^t \to \sigma_k^\infty$ and Borel set B_k, $\lim_{t\to\infty} \mu_k^i(B_k|\sigma_k^t) = \mu_k^i(B_k|\sigma_k^\infty)$, which is simply convergence of measures $\mu_k^i(\cdot|\sigma_k^t)$ to $\mu_k^i(\cdot|\sigma_k^\infty)$ in the total variation distance of probability measures. This is the technical condition assumed in Friedman and Mezzetti [2005]. However, this is incompatible with *interior full support*, which we require for the behavioral axioms. Hence, we weaken this condition by allowing for discontinuities associated with some $\{\{\sigma_k^t\},B_k\}$-pairs. In the one-dimensional case, *interior full support* only implies discontinuities when the opponent's action approaches the boundary (see discussion following Remark 1). With higher dimensions, the analogue is when the opponent's action "gains zeros" in the limit, i.e. puts positive probability on fewer pure actions. If $\{\sigma_k^t\} \subset \Delta(\sigma_k^\infty)$ for sufficiently large t, then σ_k^t does not gain zeros in the limit and so there will not be discontinuities for any B_k ((B2)-*(i)*). If σ_k^t does gain zeros in the limit, then there necessarily will be discontinuities for some B_k since the probability that beliefs realize in $\Delta(\sigma_k^\infty)$ goes from 0 to 1 by *interior full support*. However, we require continuity if $cl(B_k \cap \Delta(\sigma_k^t)) \cap \Delta(\sigma_k^\infty) = B_k \cap \Delta(\sigma_k^\infty)$, which means that the part of B_k in $\Delta(\sigma_k^t)$ (the subset of Δ_k with 0s in the same components as σ_k^t) "overlaps" with the part of B_k in $\Delta(\sigma_k^\infty)$ (the subset of Δ_k with 0s in the same components as the limit σ_k^∞) ((B2)-*(ii)*).

By construction, belief-discontinuities can only arise when the overlapping condition (B2)-*(ii)* fails.[16] In the one-dimensional, binary-action case, failures of the overlapping condition are equivalent to belief-discontinuities (see footnote 12), and it is easy to rewrite (B2) for this special case.[17] To give more intuition for (B1) and (B2) in higher dimensions, we provide some examples for the case of three pure actions in Figure 2.

To state our general behavioral axioms, we introduce the *marginal belief distribution* (CDF) defined by $F_{kj}^i(\bar{\sigma}_{k0}|\sigma_k) \equiv \mu_k^i(\{\sigma_k' \in \Delta_k : \sigma_{kj}' \in [0,\bar{\sigma}_{k0}]\}|\sigma_k)$ for all i,k, and j. This gives the probability that player i believes player k plays action a_{kj} with probability less

[16] (B2)-*(i)* is actually redundant since it implies (B2)-*(ii)*, but we include it separately as a natural sufficient condition

[17] *Continuity* (B2) *in the binary-action case.* Consider a sequence with $r \to 0^+$. For $B_k \in \{\{0\},(0,\epsilon)\}$, the overlapping condition fails: for $B_k = \{0\}$, $cl(\{0\} \cap (0,1)) \cap \{0\} = \emptyset$ and $\{0\} \cap \{0\} = \{0\}$, and for $B_k = (0,\epsilon)$, $cl((0,\epsilon) \cap (0,1)) \cap \{0\} = \{0\}$ and $(0,\epsilon) \cap \{0\} = \emptyset$. For $B_k = [0,\epsilon)$, the overlapping condition is satisfied: $cl([0,\epsilon) \cap (0,1)) \cap \{0\} = \{0\}$ and $[0,\epsilon) \cap \{0\} = \{0\}$. Given these results, it is easy to show that (B2) becomes: (1) $\mu_k^i(B_k|r)$ is continuous for all $r \in (0,1)$, (2) $\lim_{r\to 0^+} \mu_k^i(B_k|r) = \mu_k^i(B_k|0)$ for any $B_k = [0,\epsilon) \cup A$, (3) $\lim_{r\to 0^+} \mu_k^i(B_k|r) = \mu_k^i(B_k|0)$ if $B_k \cap [0,\epsilon) = \emptyset$ for some $\epsilon > 0$, (4) $\lim_{r\to 1^-} \mu_k^i(B_k|r) = \mu_k^i(B_k|1)$ for any $B_k = A \cup (\epsilon,1]$, and (5) $\lim_{r\to 1^-} \mu_k^i(B_k|r) = \mu_k^i(B_k|1)$ if $B_k \cap (1-\epsilon,1] = \emptyset$ for some $\epsilon > 0$.

than or equal to $\bar{\sigma}_{k0} \in [0,1]$ as a function of $\sigma_k \in \Delta_k$. *Belief-responsiveness* requires that the jth marginal belief distribution increases in the sense of FOSD as the probability the opponent plays the corresponding action increases. *Unbiasedness* requires that the marginal belief distributions are correct on median:

(B3) **Belief-responsiveness**: If, for some j, σ_k and $\sigma_k^{'}$ satisfy $\sigma_{kj} < \sigma_{kj}^{'}$ and $\sigma_{kl} \geq \sigma_{kl}^{'}$ for all $l \neq j$, then $F_{kj}^{i}(\bar{\sigma}_{k0}|\sigma_k^{'}) < F_{kj}^{i}(\bar{\sigma}_{k0}|\sigma_k)$ for $\bar{\sigma}_{k0} \in (0,1)$.

(B4) **Unbiasedness**: $F_{kj}^{i}(\sigma_{kj}|\sigma_k) = \frac{1}{2}$ for σ_k with $\sigma_{kj} \in (0,1)$.

The general axioms nest their binary-action counterparts. When $J(k) = 2$, it is immediate that (B1), (B3), and (B4) collapse to (B1$^{'}$), (B3$^{'}$), and (B4$^{'}$), respectively. That (B2) collapses to (B2$^{'}$) is less obvious, but becomes clear once (B2) is rewritten for the binary-action case (see Footnote 17).

Remark 2. If $J(k) = 2$, (B1)-(B4) are equivalent to (B1$^{'}$)-(B4$^{'}$).

Several other axioms come to mind as natural, and in fact will be satisfied by our parametric model.[18] However, we only impose (B1)-(B4) as they are sufficient to impose testable restrictions, and the resulting NBE will have a similar degree of flexibility as QRE:

Assumption 3. *The belief-map σ^* satisfies (B1)-(B4).*

Definition 3. Fix $\{\Gamma, \sigma^*\}$. An *NBE* is any $\sigma \in \Delta$ such that for all $i \in 1, ..., n$, $\sigma_i \in \Psi_i(\sigma_{-i}; \sigma^*_{-i})$.

From *continuity* and the fact that that the R_{ij} sets are closed in Δ, it is easy to show that $\Psi : \Delta \rightrightarrows \Delta$ is upper hemicontinuous (as well as non-empty and convex-valued). Existence of NBE then follows from standard arguments.

Theorem 2. *Fix $\{\Gamma, \sigma^*\}$. An NBE exists.*

Proof. See Appendix 8.1. □

In general, Ψ may not be single-valued, which is the case if and only if a player can be indifferent between two best responses with positive probability, i.e. if $\mu_{-i}(R_{ij} \cap R_{il}|\sigma_{-i}) > 0$ for some $\sigma_{-i} \in \Delta_{-i}$. Since *interior full support* requires that beliefs are correct with

[18] One is *belief-monotonicity* in which the distribution of $\sigma_{kj}^{i*}(\sigma_k)$ first-order stochastically dominates the distribution of $\sigma_{kl}^{i*}(\sigma_k)$ if $\sigma_{kj} > \sigma_{kl}$. Another is *label independence* in which $\sigma_{kl}^{i*}(\sigma_k)$ and $\sigma_{kj}^{i*}(\sigma_k)$ have the same distribution if $\sigma_{kj} = \sigma_{kl}$; and if σ_k and $\sigma_k^{'}$ are the same up to permutation of components, then $\sigma_{kj}^{i*}(\sigma_k)$ has the same distribution as $\sigma_{k\iota(j)}^{i*}(\sigma_k^{'})$ where $\iota : \{1, ..., J(k)\} \to \{1, ..., J(k)\}$ is the permutation mapping.

probability one when the opponents take pure actions, this occurs when a_{ij} and a_{il} are best responses to some pure action profile a_{-i}. In games without such actions, however, *interior full support* implies that Ψ is single-valued.

Lemma 1. *Fix* $\{\Gamma, \sigma^*\}$. *If* $u_i(a_{ij}, a_{-i}) \neq u_i(a_{il}, a_{-i})$ *for all* i, a_{ij}, a_{il}, *and* a_{-i}, *then* Ψ *is single-valued.*[19]

Proof. See Appendix 8.1. □

Since Ψ is upper hemicontinuous, the lemma implies that Ψ is a continuous function for generic games.

2.4 Relationship to other concepts

A refinement of rationalizability. The theory of rationalizability (Bernheim [1984] and Pearce [1984]) finds strategy profiles that cannot be ruled out on the basis of rationality and common knowledge of rationality alone, recognizing that these are not enough to form correct beliefs as required in an NE. One view is that NE predictions are too precise whereas rationalizability may be too permissive.[20] NBE is a compromise between NE and rationalizability in that it acknowledges the difficulty in forming correct beliefs and yet pins down distributions over beliefs and actions. What's more, NBE respects rationalizability in the following sense.

Lemma 2. *If* $\sigma \in \Delta$ *is an NBE, then* $a_{ij} \in supp(\sigma_i)$ *is rationalizable for all* i *and* j.

Proof. $a_{ij} \in supp(\sigma_i)$ is 1-rationalizable (a best response to some beliefs) by the assumption of rationality embedded in Ψ. Suppose every $a_{ij} \in supp(\sigma_i)$ is k-rationalizable. It must be the case that $a_{ij} \in supp(\sigma_i)$ is $(k+1)$-rationalizable. In particular, it is a best response to some k-rationalizable profile $a_{-i} \in supp(\sigma_{-i})$ because every $a_{ij} \in supp(\sigma_i)$ is a best response to some belief realization σ'_{-i}, and $supp(\sigma'_{-i}) = supp(\sigma_{-i})$ with probability one by (B1). This completes the induction. □

On the other hand, QRE does not respect rationalizability as *interiority* (A1) requires all pure actions are played with positive probability.[21]

[19] Friedman and Mezzetti [2005] have a similar result, and in fact, we invoke theirs as the final step in our proof.

[20] In generalized matching pennies, for example, NE makes a unique prediction that is sensitive to every non-affine transformation of the payoff matrix, whereas every action is rationalizable independent of payoffs (as long as the game retains the matching pennies structure).

[21] In the prisoner's dilemma for example, NBE will always predict the unique NE, whereas QRE allows for any distribution of play in which both players play the dominant strategy with probability greater than $\frac{1}{2}$.

Random belief equilibrium. NBE adopts the basic structures of RBE (Friedman and Mezzetti [2005])–belief distributions that depend on the opponent's expected actions in equilibrium. The difference between the models lies in the restrictions imposed on the belief distributions, which are tailored for different purposes. Whereas we introduce NBE as a tool for understanding the testable restrictions of equilibria with belief-noise that is "bounded away from zero", Friedman and Mezzetti [2005] use RBE for equilibrium selection and hence study the limiting case as belief-noise "goes to zero". Specifically, they consider belief measures that converge weakly to the opponents' expected action profile. Along the sequence, the restrictions they impose on belief distributions are (1) full support on the simplex and absolute continuity with respect to Lebesgue measure ($\mu_k^i(B_k|\sigma_k) > 0$ if and only if $\mathcal{L}_k(B_k) > 0$) and (2) the natural notion of continuity ($\mu_k^i(B_k|\sigma_k)$ is continuous in $\sigma_k \in \Delta_k$).[22] The only restrictions imposed by these conditions are that weakly dominated actions are played with zero probability and undominated actions are played with positive probability. In particular, RBE does not respect rationalizability as players must expect (incorrectly) that their opponents play never-best-responses. NBE's technical axioms (B1) and (B2) neither nest or are nested in the RBE conditions. In particular, the RBE conditions imply that the belief-map cannot be unbiased.[23] However, NBE can be approximated arbitrarily well by RBE in generic games,[24] so NBE can be considered as a refinement of RBE in this sense.

A model of sampling. Rubinstein and Osborne [2003] introduces sampling equilibrium in which players are frequentists who best respond to an m-length sample of their opponents' pure actions.[25] As in NBE, this sampling procedure induces a mapping from opponents' actions to distributions over beliefs. Since no sample involves actions that are not played with positive probability, and the variance of the sampling distribution goes to zero as the opponents' actions put increasing probability on a pure action profile, NBE's technical axioms capture belief formation that has a sampling flavor. Moreover, even though the sampling belief distributions are discrete, it is easy to show that they respect *belief-responsiveness*,

[22] Though unimportant for this discussion, Friedman and Mezzetti [2005] also allow for belief distributions to have finite atoms, so these restrictions only apply to the absolutely continuous part of the belief measures. That the belief distributions converge weakly implies that all atoms must vanish in the limit, except possibly for an atom on the opponents' expected action profile.

[23] From RBE's full support condition, a belief distribution cannot be unbiased on median or mean unless the opponent's action is interior. In particular, belief distributions will necessarily be biased in equilibrium when the opponent has a dominated action.

[24] Not all games: a pure strategy NE in weakly dominated strategies is an NBE (see *Multiplicity* discussion below), but no RBE can put positive probability on weakly dominated actions.

[25] If the opponent's action is $\sigma_k \in \Delta_k$, the sampling distribution follows a multinomial distribution with parameters $\sigma_k = (\sigma_{k1}, ..., \sigma_{kJ(k)})$ and player i's realized belief is the count data divided by sample size m.

and, in large samples, are approximately *unbiased* on both median and mean.[26] Hence, one can regard NBE as a generalized and "smoothed" sampling model, so our results will have implications for the empirical content of sampling models and their relationship to QRE.

Multiplicity. By *interior full support*, beliefs are correct with probability one when opponents play a pure action. Thus, any pure strategy NE is also an NBE. It is obvious therefore that for some games, there are multiple NBE for any belief-map σ^* satisfying the axioms. QRE, on the other hand, always predicts a unique equilibrium for *some* quantal response function. This is well-known as logit QRE predicts a unique equilibrium for sufficiently small λ (McKelvey and Palfrey [1995]). In this sense, NBE is more like NE, and could be paired with standard refinements. For instance, a pure strategy NE in weakly dominated strategies is an NBE, but would not survive a trembling hand.

3 Within-game restrictions and the Own Payoff Effect

Contrary to the predictions of Nash equilibrium (NE) in fully mixed games, experimental studies report two regularities. First, whereas NE predicts that each player's choice probabilities keep their opponents indifferent, there are systematic deviations *within a game*: the empirical frequency of actions typically leads to a ranking of actions for each player by expected utility to which they noisily best respond. Second, subjects exhibit the "own payoff effect" *across games*: whereas NE predicts that a change in a player's own payoffs does not affect his equilibrium choice probabilities since the other players have to be kept indifferent, subjects' behavior is systematically affected by *non-affine* transformations of their payoffs.

The best evidence for these regularities comes from generalized matching pennies (see for example Ochs [1995], McKelvey et al. [2000], and Goeree and Holt [2001]), and QRE is well-known for capturing both effects in this context (Goeree et al. [2005]). In this section, we show that NBE also captures both effects, and thus these empirical patterns can be explained equally well by adding noise to actions or adding noise to beliefs without relying on any specific functional form.

We first show that NBE imposes the same testable restrictions as QRE for any individual matching pennies game, and hence captures deviations from NE equally well.

[26] The jth marginal of the sampling distribution is a binomial with parameter σ_{kj}; dividing the count data by m gives the distribution of realized beliefs (see footnote 25). From results on the binomial distribution (Kaas and Buhrman [1980]): (1) If $m\sigma_{kj}$ is an integer, then the unique (strong) median belief is $M = \sigma_{kj}$. (2) If $m\sigma_{kj}$ is not an integer, then any (weak) median belief M satisfies $\frac{\lfloor m\sigma_{kj} \rfloor}{m} \leq M \leq \frac{\lceil m\sigma_{kj} \rceil}{m}$ due to discreteness; the bounds contain σ_{kj} and get arbitrarily tight as $m \to \infty$.

Theorem 3. *Fix Γ^m. The set of attainable NBE is equal to the set of attainable QRE.*

Proof. See Appendix 8.1. □

The theorem states that, for any individual matching pennies game, (1) any NBE is a QRE for some quantal response function, and (2) any QRE is an NBE for some belief-map. Remember that these objects must satisfy (A1)-(A4) and (B1′)-(B4′), respectively, so the result requires careful construction. The intuition, however, is simple.

When player k takes the indifferent action $r = \bar{r}_i$ that equates the expected utilities to player i's actions, player i will take each action with one half probability under both models. This follows from *monotonicity* (A4) in a QRE and *unbiasedness* (B4′) in an NBE (beliefs are equally likely to realize on either side of the indifferent belief). As player k increases his action to $r = r' > \bar{r}_i$, then one of player i's actions increases in expected utility (while decreasing for the other). This action will now be played with probability greater than one half in a QRE by *responsiveness* (A3) as well as in a NBE by *belief-responsiveness* (B3′) (as the belief distribution shifts up, the probability that this action is *subjectively* better increases).

In Appendix 8.2, we derive the set of attainable NBE for any given matching pennies game, which corresponds to the set of QRE by Theorem 3. The following example, which illustrates such a set, was derived in Goeree et al. [2005] for QRE; we re-derive the set using NBE.

	L	**R**
U	$X, 0$	$0, 1$
D	$0, 1$	$1, 0$

Figure 3: *Matching Pennies X.*

Example 1. Let $X > 0$. In the game of Figure 3, $\{\sigma_U, \sigma_L\}$ is an NBE (QRE) if and only if

$$\begin{cases} \sigma_U \leq \frac{1}{2} & \text{if } \sigma_L \leq \frac{1}{1+X} \\ \sigma_U \geq \frac{1}{2} & \text{if } \sigma_L \geq \frac{1}{1+X} \end{cases} \text{ and } \begin{cases} \sigma_L \geq \frac{1}{2} & \text{if } \sigma_U \leq \frac{1}{2} \\ \sigma_L \leq \frac{1}{2} & \text{if } \sigma_U \geq \frac{1}{2}. \end{cases}$$

Proof. Suppose $\{\sigma_U, \sigma_L\}$ is an NBE. By (B4′), the probability player 1 plays U when player 2 is playing $\sigma_L = \frac{1}{1+X}$ (the action that makes player 1 indifferent) is exactly $\sigma_U = \frac{1}{2}$. By (B3′), if $\sigma_L < \frac{1}{1+X}$, the probability player 1 plays U is strictly less than $\frac{1}{2}$. The other inequalities are similar. Conversely, any $\{\sigma_U, \sigma_L\}$ satisfying the inequalities can be attained as an NBE, which we explain below. □

For any $X > 0$, the set of attainable NBE (QRE) is given by the inequalities in Example 1. The left panel of Figure 4 plots this set when $X = 4$ as a gray rectangle, in which case only 15% of all possible outcomes are consistent with the model. A representative NBE is plotted as the intersection of reaction functions. Player 1's reaction must be strictly increasing in σ_L (*belief-responsiveness*) and pass through the point $\{\sigma_U, \sigma_L\} = \{\frac{1}{2}, \frac{1}{1+X}\}$ (*unbiasedness*) as well as the corners of the square.[27] Similarly, player 2's reaction must be strictly decreasing in σ_U and pass through the point $\{\sigma_U, \sigma_L\} = \{\frac{1}{2}, \frac{1}{2}\}$. These are the *only* restrictions on the reaction functions,[28] and hence any $\{\sigma_U, \sigma_L\}$ satisfying the inequalities can be attained in an NBE.

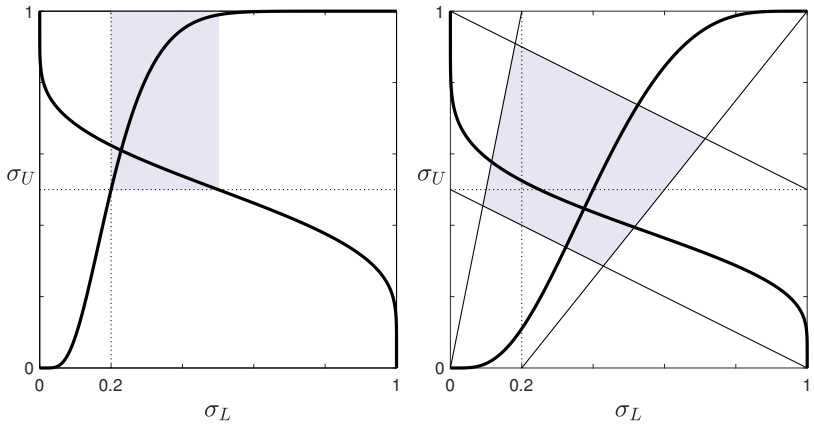

Figure 4: *NBE (QRE) in Matching Pennies X.* The left panel plots the set of attainable NBE (QRE) in the game of Figure 3 ($X = 4$) as a gray region. The NE is given as the intersection of the best response correspondences (dotted lines), and a representative NBE is given as the intersection of reaction functions (black curves). The right panel plots the set of attainable NBE in which *unbiasedness* is modified so that beliefs are correct on mean instead of median.

In the right panel of Figure 4, we illustrate the set of attainable NBE in which the *unbiasedness* axiom is modified so that beliefs are correct on mean instead of median (which we generalize to any matching pennies game in Appendix 8.2.1). The reaction function for player 1 must be increasing, fall between the upward sloping lines, and include the corners of the square, with a similar condition for player 2. Note that the reactions are unrestricted

[27] By *interior full support*, beliefs are correct with probability one when the opponent is playing a pure action to which a pure action is the unique best response. The QRE reaction functions would look qualitatively similar, except would be bounded away from the corners by *interiority*.

[28] This is implicit in our proof of Theorem 3 in which, for any QRE reaction function and sufficiently small $\epsilon > 0$, we construct a belief-map that induces an NBE reaction that agrees with the QRE reaction on $r \in [\epsilon, 1 - \epsilon]$.

at the indifferent action, but there are still testable restrictions on equilibria. If beliefs are correct on both median *and* mean, then the set of attainable NBE would be the intersection of the gray regions from the two panels and account for less than 10% of possible outcomes.

The next example illustrates the own payoff effect, which in this case is a simple comparative static in player 1's payoff parameter X. NE predicts that player 1's action does not change with X as he must mix to keep his opponent indifferent, but empirical evidence suggests a different pattern that is well-known to be explained by QRE (Goeree et al. [2005]). We now show that NBE makes the same prediction (this is *not* a corollary of Theorem 3 which only concerns individual games).

Example 2. Let $X > 0$. In the NBE (QRE) of the game in Figure 3, σ_U is strictly increasing in X and σ_L is strictly decreasing in X.

Proof. Fix σ^*. The NBE of this game is given as the unique fixed point

$$\sigma_U = \underset{+}{\Psi_U}(\sigma_L, \underset{+}{X}) \tag{4}$$

$$\sigma_L = \underset{-}{\Psi_L}(\sigma_U), \tag{5}$$

where the reactions can be written as $\Psi_U(\sigma_L, X) = 1 - F_2^1(\frac{1}{1+X}|\sigma_L)$ and $\Psi_L(\sigma_U) = F_1^2(\frac{1}{2}|\sigma_U)$. From (B1′) and (B3′), $\Psi_U(\sigma_L, X)$ is strictly increasing in both arguments, and $\Psi_L(\sigma_U)$ is strictly decreasing in σ_U. From (5), as X increases, it must be that either σ_U increases and σ_L decreases, σ_U decreases and σ_L increases, or that both σ_U and σ_L remain constant. The latter two cases are impossible since (4) implies that as X increases, σ_U increases if σ_L is constant or increases. Thus, as X increases, σ_U must strictly increase and σ_L must strictly decrease. □

Our next example combines previous results to make the simple point that, while NBE and QRE can make similar predictions, this depends crucially on the structure of the game. In particular, NBE's relationship to non-rationalizable actions is very different.

	L	**R**	**R′**
U	$4, 0$	$0, 1$	$0, -1$
D	$0, 1$	$1, 0$	$0, -1$
D′	$-1, 0$	$-1, 0$	$Z, -1$

Figure 5: *A 3×3 game with a matching pennies "core".*

Example 3. Figure 5 shows a 3×3 game with a matching pennies "core". It is constructed from the game in Figure 3 ($X = 4$) by giving each player one additional action, labelled D' and R'. R' is strictly dominated, and D' is either strictly dominated (for $Z < 0$) or iteratively dominated after deleting R' (for $Z > 0$). After removing R' and D', the *reduced game* is a standard matching pennies game. NBE respects rationalizability, so it is immediate that the set of attainable NBE for this game equals the set of NBE in the reduced game as depicted in the left panel of Figure 4, a result that holds for all values of Z. On the other hand, QRE predicts that both D' and R' are played with positive probability and thus that behavior is sensitive to changes in Z.

The types of predictions from Examples 1 and 2 find strong support in data, and we have shown they are explained equally well by noise in actions (QRE) or noise in beliefs (NBE). By contrast, Example 3 suggests an experiment (varying Z in the game of Figure 5) in which the two types of noise imply starkly different predictions. Which model would be a better description of the data, however, is an open question. Certainly, some subjects would take dominated actions, but it is unclear if the tendency to take them is sensitive to Z in the manner prescribed by QRE. Furthermore, we conjecture that, for any value of Z and with sufficient opportunity to learn, subjects would play the non-rationalizable actions with a probability that diminishes to zero long before play converges to NE.

4 The effects of payoff magnitude

It is important for the external validity of experiments to understand the effects of payoff magnitude in games. Indeed, games played in the lab are often meant to emulate their real-world counterparts, but are typically played at much lower stakes.

In applications of QRE, it is common to assume the quantal response function takes the familiar logit form. When parameter λ is chosen to best explain data from individual games, the fit is often very good. However, it is well-known that logit implies considerable sensitivity to changes in the payoff magnitude of games: for fixed λ, equilibrium predictions are sensitive to scaling one or more players' payoffs by positive constants. Such predictions have been tested experimentally by McKelvey et al. [2000] using generalized matching pennies. They find that subjects' behavior within a game is qualitatively consistent with logit, but the scaling predictions across games find little support as subjects' behavior is unaffected by scale. Importantly, since equilibria vary continuously with payoffs, this "scaling issue" implies a more general difficulty in explaining behavior across games.

In this section, we establish that the scaling issue of logit is general to all QRE in the sense that, if QRE is to explain the empirical regularities discussed in Section 3–systematic deviations from NE and the own payoff effect–it must be non-trivially sensitive to scaling and/or translating payoffs. For the class of translation invariant regular QRE, which includes logit and more generally all structural QRE with i.i.d. errors, sensitivity to scale is inescapable. By contrast, we show that NBE is invariant to both scaling and translation, and as we have already shown, explains both empirical regularities. We discuss the economics of scale and translation invariance at the end of this section.

To study QRE's properties, we begin with an analysis of quantal response functions directly before extending the results to games. This is the approach taken by Goeree et al. [2005], who define a notion of translation invariance for quantal response functions. We present their definition, along with an analogous notion of scale invariance. For technical reasons, scale invariance can only be defined for strictly positive utility vectors.[29]

Translation invariance: $Q_i(v_i) = Q_i(v_i + \gamma e_{J(i)})$ for all $v_i \in \mathbb{R}^{J(i)}$ and $\gamma \in \mathbb{R}$.[30]

Scale invariance: $Q_i(v_i) = Q_i(\beta v_i)$ for all $v_i \in \mathbb{R}^{J(i)}_{++}$ and $\beta > 0$.

For some results, we introduce an additional condition, requiring that an action is played more often when the payoffs to *all* other actions are weakly lowered. Though not implied by regularity alone, the condition is extremely weak: satisfied by all structural QRE[31] and implied by *responsiveness* when $J(i) = 2$.

Weak substitutability: $Q_{ij}(v_i) > Q_{ij}(v_i')$ whenever $v_{ij} \geq v_{ij}'$ and $v_{ik} \leq v_{ik}'$ for all $k \neq j$ with strict for some k.

An example of a quantal response function that is translation invariant, but not scale invariant, is *logit*:

$$Q_{ij}(v_i; \lambda) = \frac{e^{\lambda v_{ij}}}{\sum_{k=1}^{J(i)} e^{\lambda v_{ik}}}, \ \lambda \in [0, \infty), \tag{6}$$

where parameter λ controls the sensitivity to payoff differences. More generally, Goeree et al. [2005] show that all structural quantal response functions are translation invariant.

An example of a quantal response function that is scale invariant, but not translation

[29] To see why, consider the utility vector $v_i = (1, 0, ..., 0) \in \mathbb{R}^{J(i)}$. *Responsiveness* implies that $Q_{i1}(v_i) < Q_{i1}(\beta v_i)$ for $\beta > 1$, and hence no quantal response function can be truly scale invariant over the entire domain $\mathbb{R}^{J(i)}$.

[30] $e_{J(i)} = (1, ..., 1)$ is the vector of ones with length $J(i)$.

[31] It is a weakening of the *strong substitutability* condition of Goeree et al. [2005], which is satisfied by all structural QRE.

invariant, is the *Luce* model (Luce [1959]) for strictly positive payoffs:

$$Q_{ij}(v_i;\mu) = \frac{(v_{ij})^{\frac{1}{\mu}}}{\sum_{k=1}^{J(i)}(v_{ik})^{\frac{1}{\mu}}},\ \mu \in (0,\infty). \tag{7}$$

Hence, there exist quantal response functions that are translation invariant and those that are scale invariant. However, we show in Lemma 3 that no quantal response function satisfies both properties. In particular, for translation invariant Q_i, scale increases lead to *increasing sensitivity*: the high payoff actions are played with greater probability. For scale invariant Q_i, translation increases lead to *diminishing sensitivity*: the high payoff actions are played with smaller probability. We also characterize the limiting choice probabilities (as β and γ tend to infinity, respectively).

For simplicity, we give the result in the binary-action case, whose proof has a simple geometry which we plot in Figure 6. In Appendix 8.3, we generalize the result to arbitrary numbers of actions with the additional assumption of weak substitutability.

Lemma 3. *Fix $J(i) = 2$ and let $v_i \in \mathbb{R}^2_{++}$ be such that $v_{i1} > v_{i2}$.*

(i) Let Q_i be translation invariant and $\beta > 1$:

(a) $Q_{i1}(\beta v_i) = Q_{i1}(v_{i1} + \delta(\beta), v_{i2})$ where $\delta(\beta) > 0$ is strictly increasing in β and $\lim_{\beta\to\infty} \delta(\beta) = \infty$.

(b) $Q_{i1}(\beta v_i) > Q_{i1}(v_i)$.

(c) $\lim_{\beta\to\infty} Q_{i1}(\beta v_i) = \lim_{x\to\infty} Q_{i1}(x, v_{i2})$.[32]

(d) $|Q_{i1}(\beta v_i) - Q_{i1}(v_i)| < \epsilon$ for all $\beta \in (1, \bar{\beta}]$ if and only if $|Q_{i1}(v_{i1} + \delta, v_{i2}) - Q_{i1}(v_i)| < \epsilon$ for all $\delta \in (0, \delta(\bar{\beta})]$.

(ii) Let Q_i be scale invariant and $\gamma > 0$:

(a) $Q_{i1}(v_i + \gamma e_2) = Q_{i1}(v_{i1}, v_{i2} + \delta(\gamma))$ where $\delta(\gamma) > 0$ is strictly increasing in γ and $\lim_{\gamma\to\infty} \delta(\gamma) = v_{i1} - v_{i2} > 0$.

(b) $Q_{i1}(v_i + \gamma e_2) < Q_{i1}(v_i)$.

(c) $\lim_{\gamma\to\infty} Q_{i1}(v_i + \gamma e_2) = \frac{1}{2}$.

(d) $|Q_{i1}(v_i+\gamma e_2) - Q_{i1}(v_i)| < \epsilon$ for all $\gamma \in (0, \bar{\gamma}]$ if and only if $|Q_{i1}(v_{i1}, v_{i2}+\delta) - Q_{i1}(v_i)| < \epsilon$ for all $\delta \in (0, \delta(\bar{\gamma})]$.

Proof. (i): Take any $v_i \in \mathbb{R}_{++}$ and translation invariant Q_i. Referring to the left panel of Figure 6, scaling by $\beta > 1$ causes a shift along the dashed line to $v_i' = \beta v_i$. By translation

[32] $\lim_{x\to\infty} Q_{i1}(x, v_{i2}) = 1$ for all structural QRE as well as for all parametric models of which we are aware, but (A1)-(A4) only require that $\lim_{x\to\infty} Q_{i1}(x, v_{i2}) > \frac{1}{2}$.

invariance of Q_i, v_i and $v_i^{'}$ are on different iso-quantal response curves (dotted 45°-lines that pass through them). Define $v_i^{''}$ as the projection of $v_i^{'}$ along its iso-quantal response curve onto the horizontal line passing through v_i. This point is $v_i^{''} = (v_{i1} + \delta(\beta), v_{i2})$ where $\delta(\beta) \equiv (\beta - 1)(v_{i1} - v_{i2}) > 0$ is strictly increasing in β and $\lim_{\beta \to \infty} \delta(\beta) = \infty$. By construction, $v_i^{''}$ is on the same iso-quantal response curve as $v_i^{'}$ and directly to the right of v_i. Thus, $Q_{i1}(\underbrace{\beta v_i}_{=v_i^{'}}) = Q_{i1}(\underbrace{v_{i1} + \delta(\beta), v_{i2}}_{=v_i^{''}}) > Q_{i1}(v_i)$, where the inequality follows from *responsiveness* (A3). This shows (a) and (b); (c) follows from properties of $\delta(\beta)$; and (d) is immediate from part (a).

(ii): Take any $v_i \in \mathbb{R}_{++}$ and scale invariant Q_i. Referring to the right panel of Figure 6, translating by $\gamma > 0$ causes a shift along the dashed line to $v_i^{'} = v_i + \gamma e_2$. By scale invariance of Q_i, v_i and $v_i^{'}$ are on different iso-quantal response curves (dotted rays that pass through them and the origin). Define $v_i^{''}$ as the projection of $v_i^{'}$ along its iso-quantal response curve onto the vertical line passing through v_i. This point is $v_i^{''} = (v_{i1}, v_{i2} + \delta(\gamma))$ where $\delta(\gamma) \equiv \frac{v_{i1}}{v_{i1}+\gamma}(v_{i2} + \gamma) - v_{i2} > 0$ is strictly increasing in γ and $\lim_{\gamma \to \infty} \delta(\gamma) = v_{i1} - v_{i2} > 0$. By construction, $v_i^{''}$ is on the same iso-quantal response curve as $v_i^{'}$ and directly above v_i. Thus, $Q_{i1}(\underbrace{v_i + \gamma e_2}_{=v_i^{'}}) = Q_{i1}(\underbrace{v_{i1}, v_{i2} + \delta(\gamma)}_{=v_i^{''}}) < Q_{i1}(v_i)$, where the inequality follows from *responsiveness* (A3). This shows (a) and (b); (c) follows from properties of $\delta(\gamma)$ and *monotonicity* (A4); and (d) is immediate from part (a). □

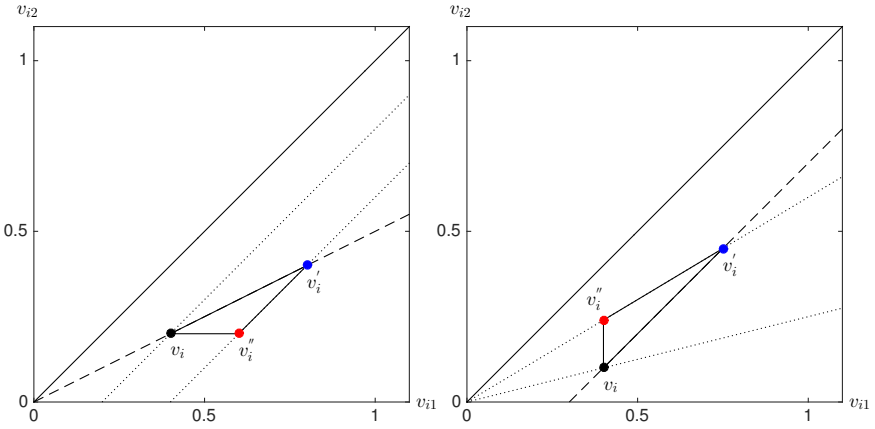

Figure 6: *Proof of Lemma 3.* The left panel plots some iso-quantal response curves (dotted lines) for a translation invariant model and illustrates the method of projection used in part *(i)* of the lemma. The right panel gives the analogous plot for scale invariant quantal response that is used in part (ii).

Lemma 3 establishes that quantal response cannot be invariant to both scale and translation, but it does not rule out translation invariant quantal response functions with very weak scale effects and vice versa. However, parts *(i)-(d)* and *(ii)-(d)* establish that translation (scale) invariant quantal response functions are *nearly* insensitive to scale (translation) if and only if they are *nearly* insensitive to payoff differences between actions. In particular, this implies that, in the limit as translation (scale) invariant quantal response becomes insensitive to scale (translation), it must assign uniform probabilities to all actions, independent of their payoffs. This generalizes what is known of the logit model (6), where λ controls both sensitivity to payoff differences *and* sensitivity to scale, and at one extreme ($\lambda = 0$) assigns uniform probabilities to all actions.[33] An important implication is that in order to explain the two empirical regularities discussed in Section 3, QRE must be non-trivially sensitive to affine transformations.

We now extend our results to games. To this end, we define families of games that differ only in *affine transformations* of payoffs.

Definition 4. Fix $\Gamma = \{N, A, u\}$.

- The *scale family* $\mathcal{S}(\Gamma)$ consists of all games Γ' such that $N' = N$, $A' = A$, and for all i, there exists $\beta_i > 0$ such that $u_i' = \beta_i u_i$.
- The *translation family* $\mathcal{T}(\Gamma)$ consists of all games Γ' such that $N' = N$, $A' = A$, and for all i and $a_{-i} \in A_{-i}$, there exists $\gamma_i(a_{-i}) \in \mathbb{R}$ such that $\bar{u}_{ij}'(a_{-i}) = \bar{u}_{ij}(a_{-i}) + \gamma_i(a_{-i})$ for all j.
- The *affine family*[34] $\mathcal{A}(\Gamma)$ consists of all games Γ' such that $N' = N$, $A' = A$, for all i and $a_{-i} \in A_{-i}$, there exists $\beta_i > 0$ and $\gamma_i(a_{-i}) \in \mathbb{R}$ such that $\bar{u}_{ij}'(a_{-i}) = \beta_i \bar{u}_{ij}(a_{-i}) + \gamma_i(a_{-i})$ for all j.

Theorem 4, which is immediate, extends the generalization of Lemma 3 (Appendix 8.3) to the QRE of games.

Theorem 4. *Fix* $\{\Gamma, Q\}$.
(i) If Q is translation (scale) invariant, the set of QRE is the same for all $\Gamma' \in \mathcal{T}(\Gamma)$ $(\Gamma' \in \mathcal{S}(\Gamma))$.

[33] Similarly, in the Luce model (7), μ controls both sensitivity to payoff differences *and* sensitivity to translation, and at one extreme ($\mu = \infty$) assigns uniform probabilities to all actions, independent of their payoffs.

[34] Note that the affine family is generically a strict superset of the others $\mathcal{S}(\Gamma) \cup \mathcal{T}(\Gamma) \subsetneq \mathcal{A}(\Gamma)$ and that the scale and translation families only overlap at the generating game $\mathcal{S}(\Gamma) \cap \mathcal{T}(\Gamma) = \Gamma$.

(ii) Let Q be weakly substitutable, and suppose $\sigma \in \Delta A$ is a QRE of Γ in which $\sigma_{ij} \neq \sigma_{ik}$ for some player i and actions j, k:

(a) If Q is translation invariant, σ is not a QRE of $\Gamma' \in \mathcal{S}(\Gamma)$.[35]

(b) If Q is scale invariant, σ is not a QRE of $\Gamma' \in \mathcal{T}(\Gamma)$.[36]

Note that in part *(ii)* of Theorem 4, we must rule out the case in which the expected utility to each of player i's actions is the same, for in that case scaling and translation coincide. This occurs in a QRE if and only if $\sigma_{ij} = \frac{1}{J(i)}$ for all j, which is clearly non-generic.

Our results suggest that by augmenting regular QRE with translation or scale invariance, we may sharpen out-of-sample predictions, i.e. that hold across games for a given quantal response function. To this end, in Appendix 8.4, we extend the method of projection used in Lemma 3 to derive necessary conditions for a dataset from *sets of binary-action games* to be consistent with QRE for some regular quantal response function under the additional maintained assumptions of translation or scale invariance, respectively. For each of translation or scale invariance, our result takes the form of inequalities that the empirical choice probabilities must satisfy. Melo et al. [2017] derive a similar result for *structural* QRE in arbitrary games under additional maintained assumptions using results from convex analysis, and when the games have binary actions, our translation invariant inequalities coincide with theirs. Though beyond the scope of the current paper, it may be interesting to extend our result to arbitrary games augmented with additional conditions.

Unlike QRE, NBE is invariant to affine transformations of the game, which is no more than a simple observation. Given choice from lotteries, the (expected utility-maximizing) best response does not depend on affine transformations of Bernoulli utilities, and this extends to games.

Theorem 5. *Fix $\{\Gamma, \sigma^*\}$. The set of NBE is the same for all $\Gamma' \in \mathcal{A}(\Gamma)$.*

Proof. See Appendix 8.1. □

The next example uses results from this section to show that the comparative static predictions of NBE and QRE may differ. In the game of Figure 7, parameter Y scales player 2's payoffs and hence indexes games in the same scale family. For any fixed Y, the sets of attainable NBE and QRE are identical (Theorem 3), but the models' comparative static predictions in Y may differ. Here, NBE makes an unambiguous prediction, while QRE does not. If QRE is augmented with scale invariance, the QRE prediction coincides with that of

[35] That is, if player i's payoffs are (non-degenerately) scaled ($\beta \neq 1$).

[36] That is, if player i's payoffs are (non-degenerately) translated ($\gamma \neq 0$).

NBE trivially. If QRE is augmented with translation invariance, the predictions diverge. Of course, by Lemma 3, one cannot impose both scale and translation invariance.

	L	**R**
U	$9, 0$	$0, Y$
D	$0, Y$	$1, 0$

Figure 7: *Matching Pennies* Y.

Example 4. Let $Y > 0$ and consider the game in Figure 7.

(i) Fix σ^*. In the NBE, σ_U and σ_L are constant in Y.

(ii) Fix scale invariant Q. In the QRE, σ_U and σ_L are constant in Y.

(iii) Fix translation invariant Q. In the QRE, σ_U and σ_L are strictly decreasing in Y.

Proof.

(i) and *(ii)*: These follow directly from Theorems 4 and 5.

(iii): Suppose Q is translation invariant. Any QRE of this game is given as the unique fixed point

$$\sigma_U = Q_U(\underset{+}{9\sigma_L}, \underset{-}{1-\sigma_L}) \tag{8}$$

$$\sigma_L = Q_L(\underset{+}{(1-\sigma_U)Y}, \underset{-}{\sigma_U Y}). \tag{9}$$

As Y increases, it must be from (8) that either σ_U and σ_L remain constant, σ_U and σ_L increase, or σ_U and σ_L decrease. The first case is impossible since if σ_U were constant, an increase in Y (a scale increase) would change σ_L (by Lemma 3) from (9). The second case is also impossible since $\sigma_U > \frac{1}{2}$ for all $Y > 0$ (as is easy to show along the lines of Example 1), and thus an increase in σ_U and Y must increase $\sigma_U Y$ by more than $(1-\sigma_U)Y$ increases, which implies a decrease in σ_L from (9) by translation invariance. □

The economics of translation and scale invariance. We view invariance to translation as an appealing normative property. Without it, models would predict that giving players "side payments" independent of the game's outcome would systematically change their behavior through a channel entirely distinct from a wealth effect, which could be embedded into the utility function directly. On the other hand, many popular models predict that behavior will change systematically with scale, so we view the question of whether behavior in games is sensitive to scale as best left to the data.

That QRE is sensitive to scale finds justification in the literature on control costs, which models errors as resulting from "trembles" that can be reduced through costly effort. Both Stahl [1990] and Mattsson and Weibull [2002] derive the multinomial logit from such an optimization, which is sensitive to scale.[37]

As stakes get higher, we may also expect subjects to expend more effort in forming their beliefs. However, by appealing to a particular theory of belief formation, we argue that in many types of games, there is no reason to expect all subjects to revise their beliefs in a similar direction. As a result, it may be that the aggregate distribution of beliefs, and hence actions, is unaffected by scale. The "endogenous depth of reasoning" theory of Alaoui and Penta [2015] uses a level k-type framework in which the "cognitive bound" depends on the payoffs of the game. Increasing the stakes induces subjects to incur additional cognitive costs to walk through more steps of higher ordered thinking, which has the effect of increasing their levels. However, if levels "cycle",[38] as is typically the case in completely mixed games, then the assumption that levels increase with scale does not provide any explanatory power for the aggregate distribution of beliefs. In other types of games, we do not necessarily expect scale invariance to hold, though the extent to which it does may give insight into subjects' strategic considerations.

5 Logit transform NBE

For applications, we construct a parametric model based on the *logit transform*. In this section, we consider the case of binary actions, and we give its generalization to normal form games in Appendix 8.5.[39] When actions are binary, player k's action is $r \in [0,1]$, and we derive player i's belief-map through the following procedure:

[37] It is tempting to interpret QRE as a model of rational inattention, as it is well-known from Matejka and McKay [2015] that if the state is a vector of payoffs (i.e. the payoff to each action), then the solution to the rational inattention problem with mutual information costs is a generalized multinomial logit that depends on the prior. Such an interpretation does not readily extend to QRE, however, since the vector of expected utilities is deterministic in equilibrium.

[38] For example, in generalized matching pennies (see Figure 1), the best response to L is U to which the best response is R to which the best response is D to which the best response is L. If level 0 is taken to uniformly mix (a common assumption), assuming U is the unique best response to $\sigma_L = \frac{1}{2}$ and R is the unique best response to $\sigma_U = \frac{1}{2}$, then the sequences of best responses indexed by levels for players 1 and 2 are $\sigma_U = \frac{1}{2}, U, D, D, U, U, D, D, ...$ and $\sigma_L = \frac{1}{2}, R, R, L, L, R, R, ...$

[39] The binary action model satisfies the axioms exactly, while the general model satisfies *unbiasedness* (B4) only approximately.

1. Map $r \in [0,1]$ to the extended real line via the logit transform

$$\mathcal{L}(r) = ln\left(\frac{r}{1-r}\right),$$

using the convention that $\mathcal{L}(0) = -\infty$ and $\mathcal{L}(1) = \infty$.

2. Add $\tau\varepsilon_i$ to $\mathcal{L}(r)$ where $\varepsilon_i \sim_{iid} \mathcal{N}(0,1)$ and $\tau \in (0,\infty)$.

3. Map $\mathcal{L}(r) + \tau\varepsilon_i$ back to $[0,1]$ via the inverse logit transform

$$r^*(r;\tau) \equiv \mathcal{L}^{-1}(\mathcal{L}(r) + \tau\varepsilon_i) = \frac{exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)}{1 + exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)}. \tag{10}$$

The parameter τ is the standard deviation of the noise added to the logit transformed action, and thus can be interpreted as the "noisiness" of beliefs. This belief-map admits closed form CDFs (and PDFs).

Fact 1. *$r^*(r;\tau)$ has CDF*[40]

$$F_k^i(\bar{r}|r;\tau) = \Phi\left(\tfrac{1}{\tau}\left[ln\left(\tfrac{\bar{r}}{1-\bar{r}}\right) - ln\left(\tfrac{r}{1-r}\right)\right]\right). \tag{11}$$

Proof. See Appendix 8.1. □

Figure 8 plots the CDF and PDF[41] of logit transform belief distributions for different values of r. Visually, it appears that the belief distributions increase in the sense of FOSD (*belief-responsiveness*) as r increases and that the median belief is correct (*unbiasedness*).[42] Since the noisy beliefs have closed form CDFs, these and the other technical axioms are easily verified.

Fact 2. *$r^*(r;\tau)$ satisfies (B1′)-(B4′).*

Proof. See Appendix 8.1. □

[40] To make the CDF well-defined, we resolve indeterminacies as follows: $-\infty-(-\infty) = \infty$ and $\infty-\infty = \infty$. As is standard, we also need $\Phi(-\infty) = 0$ and $\Phi(\infty) = 1$.

[41] Taking a derivative yields: $f_k^i(\bar{r}|r;\tau) = \phi\left(\tfrac{1}{\tau}\left[ln\left(\tfrac{\bar{r}}{1-\bar{r}}\right) - ln\left(\tfrac{r}{1-r}\right)\right]\right)\tfrac{1}{\tau}\left(\tfrac{1}{\bar{r}(1-\bar{r})}\right)$.

[42] The distributions also have skewness: when the opponent is playing $r > 0.5$, the skew is toward the left and when $r < 0.5$, the skew is toward the right. This is a consequence of the logit transform, but also a reasonable property given the boundedness of the space.

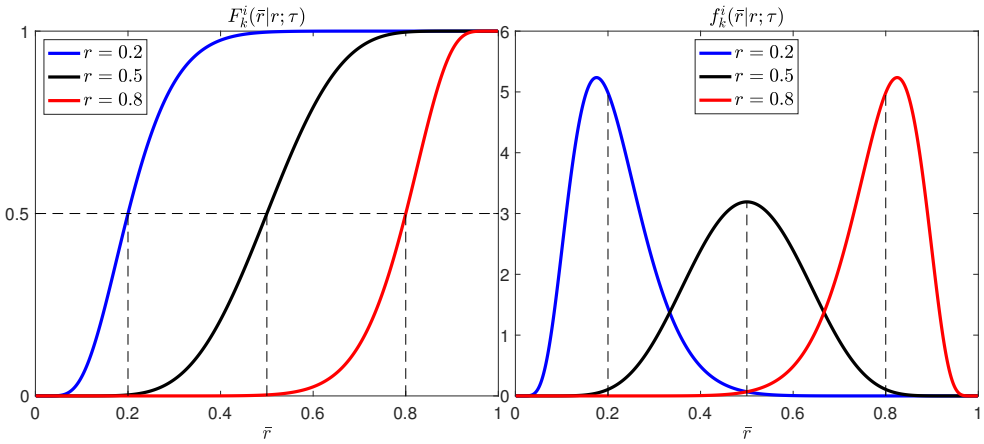

Figure 8: *Distribution of logit transform belief distributions.* This figure plots the CDFs and PDFs of player i's logit transform belief distributions for noise parameter $\tau = 0.5$ and player k's action $r \in \{0.2, 0.5, 0.8\}$.

Since the logit transform $\mathcal{L} : [0,1] \to \mathbb{R} \cup \{-\infty, \infty\}$ is a strictly increasing homeomorphism, and the normal distribution is symmetric, it is clear from the construction of the belief-map (10) that they satisfy *belief-responsiveness* and *unbiasedness*. This suggests, more generally, that *any* such homeomorphism and symmetric distribution with full support can be used to generate valid noisy beliefs. The logit transform and normal distribution are chosen only for convenience.

6 Analysis of experimental data

We consider data from several studies to test specific qualitative predictions as well as for quantitative measures of fit. We focus on three studies, McKelvey et al. [2000], Selten and Chmura [2008], and Melo et al. [2017], whose inclusion we motivate on specific grounds. We also use additional datasets for a "survey" exercise in which we fit parametric models to many datasets pooled together. First, we briefly explain the methodology we use for fitting parametric models.

6.1 Methodology

Best-fit parameters for logit transform NBE and logit QRE are chosen to minimize the *squared distance* between theoretical and observed values, as in Selten and Chmura [2008], Nyarko and Schotter [2002], and others. We focus on minimizing squared distance instead of maximizing likelihood because when fitting models to several games pooled together from

different studies, we would like to equally weight each game to get an overall measure of fit, despite the different sample sizes used.[43] Also, squared distance is more naturally extended to a measure of out-of-sample prediction error, which we make extensive use of.

For model M with parameter θ, the *squared distance* for game x is

$$\mathcal{D}_x(M,\theta) = \|\sigma_x - \sigma_x^M(\theta)\|^2,$$

where σ_x is the empirical frequency of actions, $\sigma_x^M(\theta)$ is the model prediction, and $\|\cdot\|$ is the Euclidean distance. We also define the best-fit parameter and resulting squared distance:

$$\begin{aligned} \hat{\theta}_x &= argmin_\theta\ \mathcal{D}_x(M,\theta) \\ \hat{\mathcal{D}}_x(M) &= \mathcal{D}_x(M,\hat{\theta}_x), \end{aligned} \tag{12}$$

as well as their counterparts in fitting a single parameter value to a set of K games *pooled together*:

$$\begin{aligned} \hat{\theta} &= argmin_\theta\ \frac{1}{K}\sum_x \mathcal{D}_x(M,\theta) \\ \hat{\mathcal{D}}(M) &= \frac{1}{K}\sum_x \mathcal{D}_x(M,\hat{\theta}). \end{aligned}$$

To assess if NBE significantly outperforms QRE in a set of games, we analyze the *difference in pooled squared distance*

$$\Delta\hat{\mathcal{D}} = \hat{\mathcal{D}}(QRE) - \hat{\mathcal{D}}(NBE),$$

so that $\Delta\hat{\mathcal{D}} > 0$ if and only if NBE outperforms QRE.

To determine if NBE significantly outperforms QRE, we bootstrap the distribution of $\Delta\hat{\mathcal{D}}$ and estimate the probability $\mathbb{P}(\Delta\hat{\mathcal{D}} > 0)$, which gives the maximum confidence level with which NBE outperforms QRE. We explain the bootstrap procedure in Appendix 8.7.[44] We report the bootstrap estimate $\mathcal{P} \equiv 1 - \mathbb{P}(\Delta\hat{\mathcal{D}} > 0)$, which is conceptually similar to the p-value of the one-sided hypothesis test that $\Delta\hat{\mathcal{D}} > 0$ against the null that $\Delta\hat{\mathcal{D}} = 0$.[45] We say that NBE significantly outperforms QRE if $\mathcal{P}$ is below the conventional levels for

[43] Importantly, we still make use of the sample sizes for inference.

[44] Our baseline method ignores within-subject correlation and hence overstates significance, but we show that significance is generally robust to "throwing away" a large percentage of the data and argue that this proxies for within-subject correlation in the data-generating process.

[45] Since $\mathbb{P}(\Delta\hat{\mathcal{D}} > 0)$ gives the confidence level of the largest one-sided confidence interval $[\underline{b}, \infty)$ of $\Delta\hat{\mathcal{D}}$ that excludes 0, if the conditions for standard asymptotics were met, $\mathcal{P}$ would coincide with the p-value.

significance.

6.2 McKelvey et al. 2000

We include McKelvey et al. [2000] in our analysis, and give it special attention, because their study was designed to test the payoff magnitude predictions of QRE.

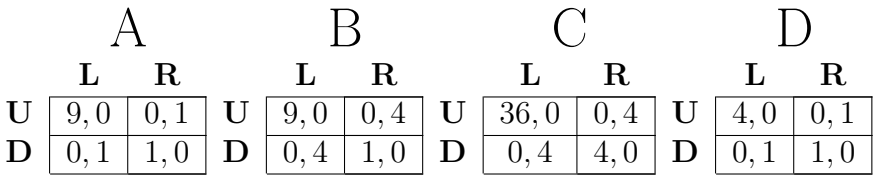

Figure 9: *Matching Pennies from* McKelvey et al. [2000].

Statistical evidence for scale effects. McKelvey et al. [2000] played the generalized matching pennies games in Figure 9. Games A-C are part of the same scale family. Relative to A, player 2's payoffs are scaled by 4 in B and both players' payoffs are scaled by 4 in C. Game D, though similar in form, is not part of this family. The action frequencies from these games are given in Table 1 and plotted in Figure 10. As is clear from the figure, the data from games A-C are very similar, with the data from game D standing out from the rest. This seems consistent with a hypothesis of scale invariance, which requires equilibria to be the same in A-C but allows for differences between D and the others.

Game	Data		
	σ_U	σ_L	N
A	0.643	0.241	1800
B	0.630	0.244	1200
C	0.594	0.257	1200
D	0.550	0.328	600

Table 1: *Data from* McKelvey et al. [2000].

Table 2 reports the results of t-tests to determine whether scale invariance can be rejected statistically. Separate tests are run for each pair of games in A-C. Since each subject in the study played a game 50 times, we cluster standard errors at the subject level to account for within-subject correlation between observed actions.[46] In all cases, scale invariance cannot

[46] *McKelvey et al. [2000]* run similar tests without clustering, though note that they are under-estimating the standard errors for exactly this reason. Even so, without clustering, only 1 of the 6 tests is significant at the 5% level.

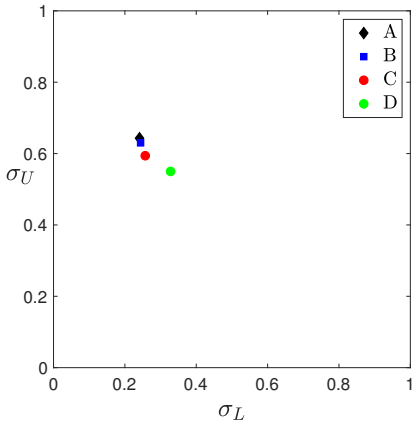

Figure 10: *Data from McKelvey et al. [2000].*

H_o	H_a	Actual	$\lvert t\rvert$	p-val.
$\sigma_U^A = \sigma_U^B$	$\sigma_U^A \neq \sigma_U^B$	$0.643 > 0.630$	0.241	0.811
$\sigma_U^A = \sigma_U^C$	$\sigma_U^A \neq \sigma_U^C$	$0.643 > 0.594$	0.893	0.376
$\sigma_U^B = \sigma_U^C$	$\sigma_U^B \neq \sigma_U^C$	$0.630 > 0.594$	0.572	0.570
$\sigma_L^A = \sigma_L^B$	$\sigma_L^A \neq \sigma_L^B$	$0.241 < 0.244$	0.062	0.951
$\sigma_L^A = \sigma_L^C$	$\sigma_L^A \neq \sigma_L^C$	$0.241 < 0.257$	0.271	0.788
$\sigma_L^B = \sigma_L^C$	$\sigma_L^B \neq \sigma_L^C$	$0.244 < 0.257$	0.212	0.833

Table 2: *Statistical Tests of Scale Effects.* This table reports the results of t-tests to determine if scale invariance can be rejected. Standard errors are clustered at the subject level.

be rejected with very large p-values ranging from 0.376 to 0.951. In the words of *McKelvey et al. [2000]*, there is an "apparent absence of payoff magnitude effects".

Qualitative predictions. We now statistically explore other qualitative predictions of NBE and QRE. Table 3 reports the results of standard t-tests of the models' predictions, and is adapted from Table 6 of *McKelvey et al. [2000]*. These predictions come in several forms. Some predictions are about the relative action frequencies *across games* and some are predictions about action frequencies *within a game* relative to the NE benchmark. We label these two kinds of predictions as "*OOS*" for *out-of-sample* and "*IS*" for *in-sample*. We mark the out-of-sample predictions across games *A*-*C* with an "*S*" since they are related to changes in scale. We also label in-sample predictions relative to the NE prediction with an "*NE*".

The NBE predictions in Table 3 hold generally under the axioms. QRE, on the other hand, makes ambiguous comparative static predictions across games A and B and across B and C without additional structure, so we augment axioms (A1)-(A4) with translation

invariance following our discussion in Section 4.[47] We have already derived several of the predictions in the table as games A and D correspond to $X = 9$ and $X = 4$ in Figure 3 and games A and B correspond to $Y = 1$ and $Y = 4$ in Figure 7. The remaining predictions can be similarly derived. Importantly, the out-of-sample predictions in the table constitute *every* such prediction that can be made unambiguously, and the in-sample predictions are the selection chosen by *McKelvey et al. [2000]*.[48]

Since the predictions hold for all NBE and all translation invariant QRE, they can be visualized in Figure 11 which plots the set of logit transform NBE and logit QRE (which is translation invariant), indexed by parameters τ and λ respectively. The out-of-sample predictions in the table correspond to all those that can be made unambiguously from the figure, i.e. those that hold for any parameter value (held fixed across the pair of games).

The results are clear. All predictions shared by NBE and QRE are in the correct direction, with most of the in-sample predictions highly significant and the out-of-sample predictions marginally significant. All of the NBE-only predictions are in the correct direction and only 1 out of 5 QRE-only predictions are in the correct direction. While none of the NBE-only or QRE-only predictions are significant at conventional levels, the p-values of the NBE-only predictions (0.164-0.296) are uniformly lower than those of the QRE-only predictions (0.405-0.715). Furthermore, as is clear from Figure 11, even when the models make unambiguous sign predictions across games, the theory allows for the differences to be arbitrarily small. Hence, marginal significance may be the best one can hope for in finite data. In any case, the qualitative patterns in the data clearly favor NBE over QRE, especially in light of the absence of scale effects documented in Table 2.

Fitting the data. We have established the qualitative patterns in the McKelvey et al. [2000] data using statistical tests, which seem to favor NBE over translation invariant QRE. So far, we have only used the structure provided by the models' axioms. We now study their parametric forms for quantitative measures of fit.

In Table 4, we compare the performance of the parametric models in the McKelvey et al. [2000] data. NBE outperforms QRE in 3 of 4 games individually by a small margin as well as when the individual squared distances are averaged together. However, since axiomatic NBE and QRE cannot be distinguished by looking at any game in isolation (Theorem 3), these

[47] *McKelvey et al. [2000]* show that these predictions hold for logit QRE, but our results establish that they hold generally for all translation invariant QRE.

[48] Additional such predictions are shared by both NBE and QRE, follow from transitivity of predictions already in the table, and are supported.

[49] These predictions require that QRE axioms (A1)-(A4) be augmented with translation invariance.

Model	Predict.	Type	H_o	H_a	Actual	t	p-val.
NBE	$\sigma_U^B > \sigma_U^D$	*OOS*	$\sigma_U^B = \sigma_U^D$	$\sigma_U^B > \sigma_U^D$	$0.630 > 0.550$	0.993	0.164
	$\sigma_U^C > \sigma_U^D$	*OOS*	$\sigma_U^C = \sigma_U^D$	$\sigma_U^C > \sigma_U^D$	$0.594 > 0.550$	0.542	0.296
QRE[49]	$\sigma_U^A > \sigma_U^B$	*OOS (S)*	$\sigma_U^A = \sigma_U^B$	$\sigma_U^A > \sigma_U^B$	$0.643 > 0.630$	0.241	0.405
	$\sigma_U^B < \sigma_U^C$	*OOS (S)*	$\sigma_U^B = \sigma_U^C$	$\sigma_U^B < \sigma_U^C$	$0.630 > 0.594$	-0.572	0.715
	$\sigma_L^A > \sigma_L^B$	*OOS (S)*	$\sigma_L^A = \sigma_L^B$	$\sigma_L^A > \sigma_L^B$	$0.241 < 0.244$	-0.062	0.523
	$\sigma_L^A > \sigma_L^C$	*OOS (S)*	$\sigma_L^A = \sigma_L^C$	$\sigma_L^A > \sigma_L^C$	$0.241 < 0.257$	-0.271	0.606
	$\sigma_L^B > \sigma_L^C$	*OOS (S)*	$\sigma_L^B = \sigma_L^C$	$\sigma_L^B > \sigma_L^C$	$0.244 < 0.257$	-0.212	0.584
NBE, QRE	$\sigma_U^A > \sigma_U^D$	*OOS*	$\sigma_U^A = \sigma_U^D$	$\sigma_U^A > \sigma_U^D$	$0.643 > 0.550$	1.252	0.108
	$\sigma_L^A < \sigma_L^D$	*OOS*	$\sigma_L^A = \sigma_L^D$	$\sigma_L^A < \sigma_L^D$	$0.241 < 0.328$	1.033	0.153
	$\sigma_L^B < \sigma_L^D$	*OOS*	$\sigma_L^B = \sigma_L^D$	$\sigma_L^B < \sigma_L^D$	$0.244 < 0.328$	0.993	0.164
	$\sigma_L^C < \sigma_L^D$	*OOS*	$\sigma_L^C = \sigma_L^D$	$\sigma_L^C < \sigma_L^D$	$0.257 < 0.328$	0.837	0.204
	$\sigma_U^B > 0.5$	*IS (NE)*	$\sigma_U^B = 0.5$	$\sigma_U^B > 0.5$	$0.630 > 0.500$	2.958	0.004
	$\sigma_U^D > 0.5$	*IS (NE)*	$\sigma_U^D = 0.5$	$\sigma_U^D > 0.5$	$0.550 > 0.500$	0.717	0.244
	$\sigma_L^A < 0.5$	*IS (NE)*	$\sigma_L^A = 0.5$	$\sigma_L^A < 0.5$	$0.241 < 0.500$	6.173	0.000
	$\sigma_L^A > 0.1$	*IS (NE)*	$\sigma_L^A = 0.1$	$\sigma_L^A > 0.1$	$0.241 > 0.100$	3.344	0.001
	$\sigma_L^C > 0.1$	*IS (NE)*	$\sigma_L^C = 0.1$	$\sigma_L^C > 0.1$	$0.257 > 0.100$	3.648	0.001
	$\sigma_L^D < 0.5$	*IS (NE)*	$\sigma_L^D = 0.5$	$\sigma_L^D < 0.5$	$0.328 < 0.500$	2.247	0.023

Table 3: *Summary of Predictions vs. Actual Behavior.* This table reports the results of t-tests of model predictions. Standard errors are clustered at the subject level. "*IS*" and "*OOS*" mark predictions *within* and *across* games respectively. "*S*" refers to a prediction across games related by scale, and "*NE*" refers to a prediction relative to the corresponding NE prediction. Positive (negative) t-values indicate that the predicted direction of the effect is correct (incorrect).

differences in performance must be related to model structures and should not be interpreted as fundamental. To distinguish the two models, we favor the pooled squared distance for all four games as a measure of overall goodness of fit. We find that NBE significantly outperforms QRE, with NBE's pooled squared distance (0.0159) 72% that of QRE (0.0218).

That NBE outperforms QRE in the pooled data is a finding we attribute to scale effects. Games *A*-*C* belong to the same scale family, and the data from these games are very similar. NBE with a single value of τ predicts the same behavior in these games (scale invariance), whereas QRE with a single value of λ predicts widely diverging behavior in *A*-*C* (sensitivity to scale). Hence, while NBE and QRE perform similarly when fit to each game individually, QRE's performance suffers much more by restricting its parameter to be the same across games. For each model, we take the ratio of the average individual squared distance to the pooled squared distance, which heuristically measures the loss in reducing the number

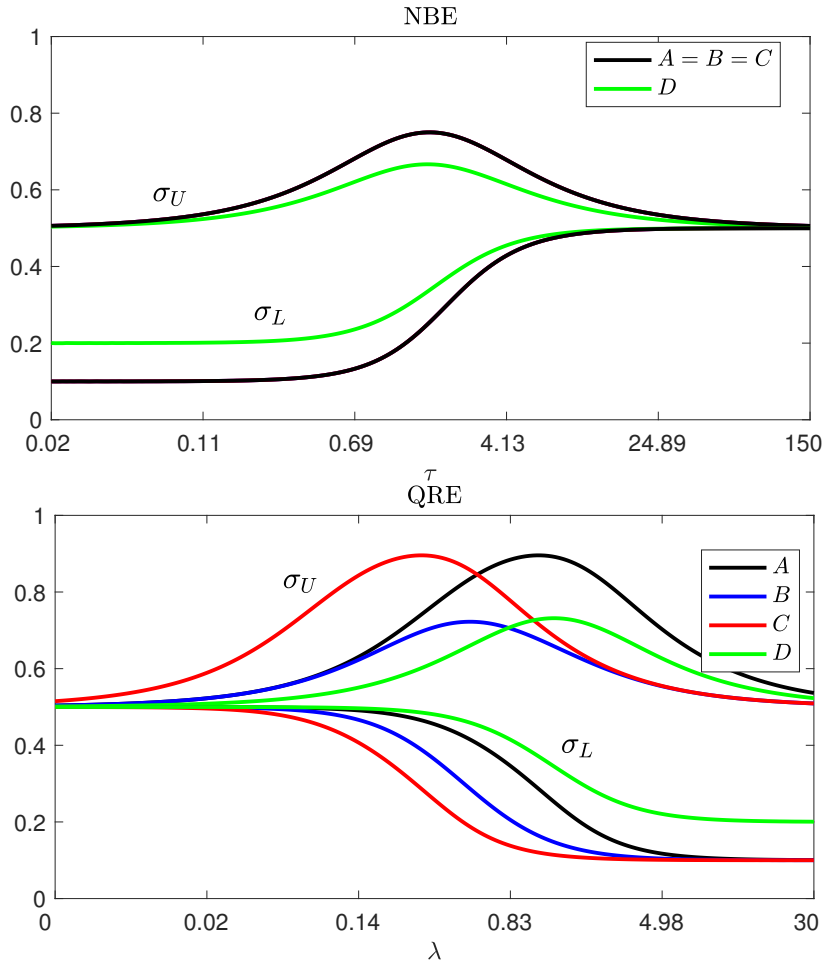

Figure 11: *NBE and QRE Correspondences.* This figure plots the entire set of logit transform NBE and logit QRE (i.e. indexed by their parameters) for games *A*-*D* from *McKelvey et al. [2000].*

of parameters. For NBE and QRE, the figures are 0.95 and 0.72 respectively, indicating a sense in which NBE-τ is more stable than QRE-λ. This is also clear from inspecting the λ estimates. For instance, since game C is the same as A up to a scale factor of 4, QRE makes the same prediction in these games when $\lambda_A = 4\lambda_C$, and indeed we see that $\hat{\lambda}_A$ is much larger than $\hat{\lambda}_C$. The pooled estimate satisfies $\hat{\lambda}_C < \hat{\lambda}_{pooled} < \hat{\lambda}_A$, implying over-sensitivity to payoff differences in C and under-sensitivity in A.

We have established that, qualitatively, NBE and QRE make similar predictions *in-sample*, but very different predictions *out-of-sample* across games that differ in scale. We now quantify these effects by examining the *prediction error* of the parametric models in making out-of-sample predictions *across* games. For each game $x \in \{A, B, C, D\}$, we fit logit transform NBE and logit QRE as in (12), resulting in estimates $\hat{\tau}_x$ and $\hat{\lambda}_x$. We then

Game	Data			NBE		QRE		NBE	QRE	NBE	QRE
	σ_U	σ_L	N	σ_U	σ_L	σ_U	σ_L	$\hat{\tau}$	$\hat{\lambda}$	$\hat{\mathcal{D}}$	
A	0.643	0.241	1800	0.747	0.221	0.662	0.110	1.411	6.459	0.0112	0.0174
B	0.630	0.244	1200	0.748	0.227	0.707	0.210	1.451	0.800	0.0142	0.0071
C	0.594	0.257	1200	0.629	0.114	0.607	0.104	0.436	2.513	0.0217	0.0234
D	0.550	0.328	600	0.661	0.295	0.570	0.206	1.239	9.579	0.0134	0.0153
Pooled	–	–	–	–	–	–	–	1.456	5.365	**0.0159**	**0.0218**
($\mathcal{P}$)	–	–	–	–	–	–	–	–	–	–	(0.000)

Table 4: *Summary of Estimates from* McKelvey et al. [2000].

$\Delta\hat{\mathcal{E}}_{xy}$

	A	B	C	D
A	0.0062	0.0136	0.0024	0.0025
B	0.0582	−0.0072	0.0249	0.0189
C	0.0236	0.0002	0.0017	0.0124
D	0.0081	0.0163	0.0038	0.0019

$\mathcal{P}$

	A	B	C	D
A	0.000	0.019	0.325	0.261
B	0.000	0.996	0.000	0.002
C	0.016	0.389	0.000	0.019
D	0.174	0.097	0.245	0.000

Table 5: *Out-of-sample Differences in Prediction Error (QRE* minus *NBE).* The xy-th entry corresponds to $\Delta\hat{\mathcal{E}}_{xy}$ as in (13) for games $x, y \in \{A, B, C, D\}$ and gives the difference in prediction error between the two models using the data from game x (row) to make predictions about game y (column). Positive (negative) entries indicate that NBE performs better than (worse than) QRE.

use these parameter estimates to make *out-of-sample* predictions for game $y \in \{A, B, C, D\}$. We define the *xy-difference in prediction error* as

$$\Delta\hat{\mathcal{E}}_{xy} \equiv \mathcal{D}_y(QRE, \hat{\lambda}_x) - \mathcal{D}_y(NBE, \hat{\tau}_x), \tag{13}$$

which we use to populate the matrix in Table 5. The diagonal entries are *in-sample*, the off-diagonal entries are *out-of-sample*, and positive entries indicate that NBE outperforms QRE. From the table, it is clear that NBE outperforms QRE in 3 of 4 games in-sample and in all 12 out-of-sample comparisons, though not all differences are significantly positive. Note also that the average (absolute) in-sample difference in prediction error of 0.0043 is small compared to the out-of-sample average of 0.0154.

Risk aversion. Several hypotheses have been proposed to account for QRE's apparent over-sensitivity to scale. In our view, the most successful is based on risk aversion and explored in Goeree et al. [2003] who fit logit QRE to games A-D by jointly estimating λ and a risk aversion parameter. Incorporating curvature into the utility function reduces the effect of scaling games' monetary payoffs in the following sense. Holding fixed the opponent's action, scaling a game's monetary payoffs by a factor of 4 (say) increases expected utility

differences by a factor less than 4. Hence, risk aversion is a natural candidate for reconciling the QRE hypothesis with data. It is also the case that risk aversion implies that A-C are no longer in the same scale family once expressed in utiles, and hence NBE may give different predictions in these games for the same value of τ. In Appendix 8.8, we replicate the exercise from Goeree et al. [2003] by fitting both NBE and QRE with constant relative risk averse (CRRA) utility to the data. We show that (1) NBE significantly outperforms QRE and (2) NBE *is* invariant to scaling monetary payoffs under CRRA utility.

6.3 Selten and Chmura 2008

Selten and Chmura [2008] collect data from 12 generalized matching pennies games, to which they fit several "stationary concepts" including logit QRE and sampling equilibrium (Rubinstein and Osborne [2003]).[50] They state that "It is not easy to understand why the predictions...are not very far apart, in spite of the fact that they are based on very different principles". However, due to the structure of the games and the fact that sampling equilibrium approximately satisfies the NBE axioms (see Section 2.4), our results (Theorem 3) suggest that the predictions should be very similar in individual games. The models should also make similar predictions in sets of games as long as they are of similar payoff magnitude, which is the case for the 12 games considered. Hence, we revisit their study to test this hypothesis and shed light on the puzzle they put forth.

The 12 games are reported in Appendix 8.6, and Table 6 reports the fits of logit transform NBE and logit QRE. Unsurprisingly, the individual game performance between the two models is virtually identical; averaging the squared distances across all 12 game gives 0.0029 for NBE and 0.0028 for QRE. However, the pooled fit does favor QRE by a reasonable margin, which has a squared distance (0.0049) that is 83% that of NBE (0.0059). We have no intuition for this finding because the games are of similar payoff magnitude and there is no obvious relationship between them.

6.4 Melo et al. 2018

We fit logit transform NBE and logit QRE to the 3×3 "Joker" games from Melo et al. [2017]. We chose these games because they are among the simplest games with unique regular QRE in which each player has more than 2 actions. Our hypothesis is that NBE will behave

[50] Selten and Chmura [2008] misreport the fits of both QRE and sampling equilibrium, which is pointed out in a comment by Brunner et al. [2011], who report the correct fits and remark that the models are "about equally accurate".

Game	Data			NBE		QRE		NBE	QRE	NBE	QRE
	σ_U	σ_L	N	σ_U	σ_L	σ_U	σ_L	$\hat{\tau}$	$\hat{\lambda}$	$\hat{\mathcal{D}}$	
1	0.079	0.690	9600	0.059	0.689	0.037	0.688	0.960	1.342	0.0004	0.0017
2	0.217	0.527	9600	0.177	0.514	0.179	0.505	0.998	0.625	0.0018	0.0019
3	0.163	0.793	9600	0.150	0.796	0.172	0.790	0.911	1.199	0.0002	0.0001
4	0.286	0.736	9600	0.266	0.745	0.278	0.740	0.689	1.112	0.0005	0.0001
5	0.327	0.664	9600	0.310	0.671	0.312	0.670	0.542	1.257	0.0003	0.0002
6	0.445	0.596	9600	0.423	0.616	0.423	0.616	0.436	1.371	0.0009	0.0009
7	0.141	0.564	4800	0.077	0.550	0.073	0.530	1.473	0.609	0.0043	0.0058
8	0.250	0.586	4800	0.163	0.576	0.156	0.567	0.687	0.955	0.0077	0.0092
9	0.254	0.827	4800	0.226	0.896	0.211	0.865	0.199	2.731	0.0055	0.0033
10	0.366	0.699	4800	0.266	0.745	0.278	0.740	0.689	1.112	0.0121	0.0094
11	0.331	0.652	4800	0.306	0.660	0.309	0.660	0.627	1.094	0.0007	0.0006
12	0.439	0.604	4800	0.425	0.618	0.425	0.617	0.406	1.455	0.0004	0.0004
Pooled	–	–	–	–	–	–	–	0.966	1.049	**0.0059**	**0.0049**
($\mathcal{P}$)	–	–	–	–	–	–	–	–	–	–	(1.000)

Table 6: *Summary of Estimates from* Selten and Chmura [2008].

similarly to QRE, and this is indeed the case. The games and data are in Figure 12, and the estimates are in Table 7. Not only is every NBE prediction also a regular QRE,[51] the NBE predictions are very similar to the logit QRE predictions. All three games are of similar scale, and hence the pooled fits are very similar also. These results suggest that it is difficult to distinguish noise-in-actions from noise-in-beliefs, even in more general fully mixed games.

Game	Player	NBE			QRE			NBE	QRE	NBE	QRE
		1	*2*	*J*	*1*	*2*	*J*	$\hat{\tau}$	$\hat{\lambda}$	$\hat{\mathcal{D}}$	
2	1	0.279	0.279	0.441	0.286	0.286	0.428	1.358	0.206	0.0045	0.0052
	2	0.399	0.399	0.201	0.391	0.391	0.217				
3	1	0.368	0.368	0.265	0.368	0.368	0.263	8.840	0.149	0.0241	0.0168
	2	0.325	0.325	0.350	0.297	0.297	0.407				
4	1	0.389	0.306	0.306	0.384	0.308	0.308	0.934	0.363	0.0315	0.0320
	2	0.386	0.228	0.386	0.388	0.224	0.388				
Pooled	–	–	–	–	–	–	–	1.173	0.243	**0.0204**	**0.0197**
($\mathcal{P}$)	–	–	–	–	–	–	–	–	–	–	(0.811)

Table 7: *Summary of Estimates from* Melo et al. [2017].

[51] As shown in Melo et al. [2017], all regular QRE satisfy: $\sigma_{11} = \sigma_{12} \in (0, \frac{1}{3})$ and $\sigma_{21} = \sigma_{22} \in (\frac{1}{3}, \frac{9}{22}]$ (game 2); $\sigma_{11} = \sigma_{12} \in (\frac{1}{3}, 1)$ and $\sigma_{21} = \sigma_{22} \in [\frac{4}{15}, \frac{1}{3})$ (game 3); and $\sigma_{12} = \sigma_{1J} \in (0, \frac{1}{3})$ and $\sigma_{21} = \sigma_{2J} \in (\frac{1}{3}, \frac{2}{5}]$ (game 4).

$2(n = 1620)$

	1 (0.359)	**2** (0.439)	**J** (0.202)
1 (0.253)	10, 30	30, 10	10, 30
2 (0.304)	30, 10	10, 30	10, 30
J (0.442)	10, 30	10, 30	55, 10

$3(n = 940)$

	1 (0.258)	**2** (0.323)	**J** (0.419)
1 (0.340)	25, 30	30, 10	10, 30
2 (0.464)	30, 10	25, 30	10, 30
J (0.196)	10, 30	10, 30	30, 10

$4(n = 300)$

	1 (0.487)	**2** (0.147)	**J** (0.366)
1 (0.473)	20, 30	30, 10	10, 30
2 (0.220)	30, 10	10, 30	10, 30
J (0.307)	10, 30	10, 30	30, 10

Figure 12: *Joker games from* Melo et al. [2017].

6.5 A survey

The *McKelvey et al. [2000]* study was specifically concerned with the effects of scale, and hence ran games that were constructed to be in the same scale family. While there are no formal results akin to Theorem 5 for games that are not precisely in the same scale family, it is obvious "by continuity" that QRE is sensitive to the general scale of games in a way that NBE is not. Hence, we add to the evidence of the previous sections by fitting parametric models to a dataset of 21 generalized matching pennies games from 5 studies: MPW 2000 (McKelvey et al. [2000]), SC 2008 (Selten and Chmura [2008]), Ochs 1995 (Ochs [1995]), GH 2001 (Goeree and Holt [2001]), and NS 2002 (Nyarko and Schotter [2002]). The dataset includes games with much larger scale differences than can be found within any one study.

Logit transform NBE (or any parametric NBE) has the property that the estimates are insensitive to the "exchange rate" between utility and money. That is, *holding the data fixed* but scaling the arbitrary utilities in the payoff matrix for one or more players does not effect the estimated τ or the resulting prediction. Logit QRE on the other hand, or any translation invariant parameterization of QRE, is sensitive to these scalings. If all players' utility numbers are scaled up by some factor $c > 0$, the estimate $\hat{\lambda}$ is simply replaced with $\hat{\lambda}' = (\frac{1}{c})\hat{\lambda}$ and the predicted equilibria remain unchanged. If, however, not all players' utilities are scaled by the same factor, the λ estimate will change in unpredictable ways and lead to different predicted equilibria as well.

Hence, to fit QRE to data pooled together from different studies (or to make λ estimates

comparable across studies) requires that the utility payoffs are adjusted for inflation and currency-to-currency exchange rates. We take McKelvey et al. [2000] as the base-study to derive the rate of 1 utile per 0.10 year-2000 US dollars. Using this rate, we multiply the utility payoffs from each game in our dataset by a study-specific conversion factor prior to estimation. We explain the exact procedure and give the conversion factors in Appendix 8.9. The 21 pre-transformed games are given in Appendix 8.6 along with details of the experimental procedures, though it is the transformed payoffs that are used in all estimations.

We fit logit transform NBE, logit QRE, and Luce QRE (i.e. with quantal response function (7)) to the data. Recall that logit is translation invariant, Luce is scale invariant, and NBE is both. Table 13 of Appendix 8.10 gives the parameter estimates and resulting squared distances. Unsurprisingly, the individual-game performance of all three models is very similar for most games. In any case, the differences should not be interpreted as fundamental (due to Theorem 3) and we favor a test based on all 21 game pooled together.

Comparing the pooled squared distances of the three parametric models yields a clear ordering from best to worst: NBE (0.0133), Luce QRE (0.0148), and Logit QRE (0.0248). The NBE distance is 54% that of logit and 90% that of Luce; and that NBE outperforms the others is highly significant. Since the models perform similarly when each game is fit individually, we interpret this finding as suggesting the value of both scale and translation invariance in explaining the data. It is unsurprising given the results of Section 6.2 that scale invariance is powerful in explaining the data, which accounts for the large difference in performance between NBE and logit QRE. Interestingly, translation invariance also seems somewhat valuable as NBE outperforms the scale invariant (and translation sensitive) Luce QRE.

It is a valid concern that the pooled estimate of λ may be very sensitive to the utility-money exchange rate conversions. If that were the case, even small errors in conversions could effect the logit QRE squared distance (in either direction). Even holding fixed the year and country across studies might not be sufficient, as different subject pools might value the same amount of real money differently, resulting in different utility valuations. As a robustness check, we show in Appendix 8.9 that the result of this section is not sensitive to even very large perturbations of the conversion factors.

7 Conclusion

It is well-known that Nash equilibrium (NE) fails to explain the richness of experimental data. Many models have been proposed as a result. One prominent example is *quantal response equilibrium* (QRE), which relaxes the rationality requirement of NE by allowing for "noise in actions". We introduce *noisy belief equilibrium* (NBE), which relaxes the other condition of NE by allowing for "noise in beliefs". In an NBE, axioms restrict belief distributions to be *unbiased with respect to* and *responsive to changes in* the opponents' behavior. We study the testable restrictions imposed by NBE, which we compare to those of *regular* QRE in which axioms restrict the primitive quantal response function.

We find that NBE explains, just as QRE does, some commonly observed deviations from NE and the own payoff effect. The mechanism whereby QRE achieves this is a sensitivity to payoff differences, which we show is linked inextricably to a sensitivity to affine transformations of payoffs. By contrast, beliefs-based models such as NBE are generally invariant to affine transformations, which we show is valuable in explaining experimental data. Unlike QRE, NBE respects rationalizability, and hence has a fundamentally different relationship to dominated and iteratively deleted actions, which we believe merits further experimental study. It would also be interesting to test the NBE axioms directly by comparing elicited beliefs to the empirical distribution of actions.

References

Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *Review of Economic Studies*, 2015. 1, 4

Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 1984. 1, 2.4

Christoph Brunner, Colin Camerer, and Jacob Goeree. Stationary concepts for experimental 2x2 games: Comment. *American Economic Review*, 2011. 50, 58

Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 2004. 1

Hsiao-Chi Chen, James W. Friedman, and Jacques-Francois Thisse. Boundedly rational nash equilibrium: A probabilistic choice approach. *Games and Economic Behavior*, 1997. 1

Evan Friedman. Endogenous quantal response equilibrium. *Working Paper*, 2018. 4

James Friedman and Claudio Mezzetti. Random belief equilibrium in normal form games. *Games and Economic Behavior*, 2005. 1, 2.3.1, 2.3.2, 2.3.2, 19, 2.4, 22, 8.1, 8.1

Jacob Goeree and Charles Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 2001. 3, 6.5, 8, 8.6.4, 8.7, 13

Jacob Goeree and Charles Holt. A model of noisy introspection. *Games and Economic Behavior*, 2004. 5

Jacob Goeree, Charles Holt, and Thomas Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 2003. 3, 4, 6.2, 8.8, 8.8, 55, 15, 11

Jacob Goeree, Charles Holt, and Thomas Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 2005. 1, 2, 2.2, 3, 3, 3, 4, 4, 31

Jacob Goeree, Charles Holt, and Thomas Palfrey. Quantal response equilibrium. *Princeton University Press*, 2016. 1

Duarte Goncalves and Teresa Esteban Casanelles. Costly introspection equilibrium. *working paper*, 2018. 1

Philip Haile, Ali Hortacsu, and Grigory Kosenok. On the empirical content of quantal response equilibrium. *American Economic Review*, 2008. 1, 2, 2.2

Yuval Heller and Eyal Winter. Biased-belief equilibrium. *working paper*, 2018. 1

Rob Kaas and J.M. Buhrman. Mean, median, and mode in binomial distributions. *Statistica Neerlandia*, 1980. 26

Robert Luce. Individual choice behavior. *Wiley*, 1959. 4

Filip Matejka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit. *American Economic Review*, 2015. 37

Lars-Goran Mattsson and Jorgen Weibull. Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*, 2002. 4

Felix Mauersberger. Thompson sampling: Endogenously random behavior in games and markets. *Working Paper*, 2018. 1

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 1995. 1, 2.4, 8.6.3

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1998. 1

Richard McKelvey, Thomas Palfrey, and Roberto Weber. Endogenous rationality equilibrium. *Working Paper*, 1997. 4

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. 1, 4, 3, 4, 6, 6.2, 9, 6.2, 1, 46, 10, 6.2, 6.2, 47, 11, 4, 6.5, 8, 8.6.1, 10, 8.8, 11, 8.9.1, 13

Emerson Melo, Kiril Pogorelskiy, and Matthew Shum. Testing the quantal response hypothesis. *Working Paper*, 2017. 4, 6, 6.4, 7, 51, 12, 8.5, 8, 8.6.6

Rosemarie Nagel. Unraveling in guessing games: an experimental study. *American Economic Review*, 1995. 1

Yaw Nyarko and Andrew Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 2002. 6.1, 6.5, 8, 8.6.5, 13

Jack Ochs. An experimental study of games with unique mixed strategy equilibria. *Games and Economic Behavior*, 1995. 3, 3, 6.5, 8, 8.6.3, 13

David Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 1984. 1, 2.4

Ariel Rubinstein and Martin Osborne. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 2003. 1, 2.4, 6.3

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2 games. *American Economic Review*, 2008. 1, 7, 6, 6.1, 6.3, 50, 6, 6.5, 8, 8.6.2, 13, 58

Dale Stahl. Entropy control costs and entropic equilibrium. *International Journal of Game Theory*, 1990. 4

Dale Stahl and Paul Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 1995. 1

8 Appendix

8.1 Proofs

Theorem 1. *Fix* $\{\Gamma^m, \sigma^*\}$. *An NBE exists and is unique and interior.*

Proof. $\Psi : [0,1]^2 \to [0,1]^2$ is a continuous function mapping from a compact and convex set to itself (from (B2$'$) as already shown). By Brouwer's fixed point theorem, there exists a fixed point of Ψ. To show interiority of any fixed points, suppose for purposes of contradiction that some player k is playing $r \in \{0,1\}$ in an NBE. But then by (B1$'$), player i forms belief $r^*(r) = r$ to which a pure action $s \in \{0,1\}$ is the only best response. (r,s) cannot be an NBE, since if it were, it would also be an NE, and the game has no pure strategy NE. Thus, all fixed points of Ψ are interior, and we only need to check $(\sigma_U, \sigma_L) \in (0,1)^2$. That the fixed point is unique follows from the fact that $\Psi_U(\sigma_L)$ is strictly increasing in $\sigma_L \in (0,1)$ and $\Psi_L(\sigma_U)$ is strictly decreasing in $\sigma_U \in (0,1)$ by (B3$'$). □

Theorem 2. *Fix* $\{\Gamma, \sigma^*\}$. *An NBE exists.*

Proof. An NBE is a fixed point of $\Psi : \Delta \rightrightarrows \Delta$. It is trivial to show that Ψ is non-empty and convex-valued; and Δ is non-empty, compact, and convex. Existence of NBE follows from Kakutani's fixed point theorem after showing that Ψ_i (and thus Ψ) is upper hemicontinuous. To this end, let $z_i \subset \{1,2,...,J(i)\} \equiv [J(i)]$ be an arbitrary, possibly empty, subset of action indices, and define

$$e_i(z_i) = \left\{ \sigma_{-i}^{'} \in \Delta_{-i} |\ \sigma_{-i}^{'} \in \bigcap_{j \in z_i} R_{ij},\ \sigma_{-i}^{'} \notin R_{ik} \text{ for } k \notin z_i \right\} \tag{14}$$

as the set of beliefs for which actions indexed in z_i, and only those actions, are best responses (let $e_i(\emptyset) = \emptyset$). Friedman and Mezzetti [2005] previously defined this object which is used in some of their results. Note that the collection $\mathcal{E} \equiv \{e_i(z_i)\}_{z_i \subset [J(i)]}$ defines a partition of Δ_{-i}. Let $\{\sigma_{-i}^t\} \subset \Delta_{-i}$ be an arbitrary convergent sequence with $\sigma_{-i}^t \to \sigma_{-i}^\infty$ as $t \to \infty$. Ψ_i is upper hemicontinuous if, for all such sequences, there exists a rational strategy $s_i : \Delta_{-i} \to \Delta_i$ for which $\int_{\Delta_{-i}} s_i(\sigma_{-i}^{'}) d\mu_{-i}(\sigma_{-i}^{'} | \sigma_{-i}^t) \to \int_{\Delta_{-i}} s_i(\sigma_{-i}^{'}) d\mu_{-i}(\sigma_{-i}^{'} | \sigma_{-i}^\infty)$. There are two cases to consider (any convergent sequence will fall into one of the cases for sufficiently high t), and for each, we construct such a strategy s_i.

Case 1: Let $\{\sigma_{-i}^t\} \subset \prod_{k \neq i} \Delta(\sigma_k^\infty)$. Let $E \in \mathcal{E}$ be an arbitrary partition element. From (B2)-*(i)*, $\mu_{-i}(E | \sigma_{-i}^t)$ is continuous for all t, and hence we can set s_i to be any that is rational and constant within each partition element.

Case 2: Let $\{\sigma_{-i}^t\} \not\subset \prod_{k \neq i} \Delta(\sigma_k^\infty)$, meaning σ_k^t "gains zeros" in the limit for some k. Further suppose t is sufficiently high such that, for all k, $\Delta(\sigma_k^t) = \Delta(\sigma_k^{t'})$ for all $t < t' < \infty$ so that along the remaining sequence σ_k^t does not gain or lose zeros except in the limit. We must modify the proof from that of case 1 because $\mu_{-i}(E | \sigma_{-i}^t)$ may be discontinuous for some $E \in \mathcal{E}$ as $t \to \infty$ by (B1).[52] Define $C_{ij} \equiv cl(R_{ij} \cap \prod_{k \neq i} \Delta(\sigma_k^t))$ for all j. Since R_{ij} is closed, $C_{ij} \subset R_{ij}$. By construction,

[52] Consider a 2×2 game in which player i (row) has 2 actions U and D in which U weakly dominates D. Only when player i's realized belief is $r^{'} = 0$ is D a best response, and hence $e_i(U, D) = \{0\}$. (B1) implies that $\mu_k^i(\{0\} | r) = 0$ for all $r \in (0,1)$ and $\mu_k^i(\{0\} | 0) = 1$, and hence $u_k^i(\{0\} | r)$ is discontinuous as $r \to 0^+$.

$\mu_{-i}(C_{ij}|\sigma^t_{-i}) = \mu_{-i}(R_{ij}|\sigma^t_{-i})$ for $t < \infty$ and by the second bullet point of (B2) $\mu_{-i}(C_{ij}|\sigma^t_{-i}) \to \mu_{-i}(C_{ij}|\sigma^\infty_{-i})$. The proof proceeds exactly as in case 1, with C_{ij} replacing R_{ij}. That is, define $c_i(\cdot)$ as in (14) except with C_{ij} replacing R_{ij}. $\mathcal{C} \equiv \{c_i(z_i)\}_{z_i \subset [J(i)]}$ defines a partition of Δ_{-i} with $C \in \mathcal{C}$ an arbitrary element. By (B2)-*(ii)*, $\mu_{-i}(C|\sigma^t_{-i})$ is continuous for all t, and hence we can set s_i to be any that is rational and constant within each partition element.[53] □

Lemma 1. *Fix* $\{\Gamma, \sigma^*\}$. *If* $u_i(a_{ij}, a_{-i}) \neq u_i(a_{il}, a_{-i})$ *for all* i, a_{ij}, a_{il}, *and* a_{-i}, *then* Ψ *is single-valued.*

Proof. Fix player i and any $\sigma_{-i} \in \Delta_{-i}$. If σ_{-i} is a pure action profile, then it is immediate from (B1) that player i will have correct beliefs with probability one, to which one of his pure actions is a strict best response by assumptions on u_i, making Ψ_i single-valued. So assume not, i.e. at least one $k \neq i$ has an action j such that $\sigma_{kj} \in (0,1)$. By (B1), with probability one, player i's beliefs only put positive probability on pure actions in the support of σ_{-i}, and so we show that it is as if player i is playing a restricted game $\Gamma' = \{N', A', u'\}$ in which his opponents take a fully mixed profile. Specifically, define $N'_{-i}(\sigma_{-i}) \equiv \{k \in N | k \neq i,\ \sigma_{kj} \in (0,1) \text{ for some } j\}$ as the set of i's opponents who are mixing (over at least 2 pure actions) under σ_{-i} and for each $k \in N'_{-i}(\sigma_{-i})$, define $A'_k(\sigma_{-i}) \equiv \{a_{kj} \in A_k | \sigma_k(a_{kj}) \in (0,1)\}$ as the set of pure actions played by k with interior probability under σ_k. Define $A'_{-i}(\sigma_{-i}) \equiv \times_{k \in N'_{-i}(\sigma_{-i})} A'_k(\sigma_{-i})$ and $A'(\sigma_{-i}) \equiv A_i \times A'_{-i}(\sigma_{-i})$ as restricted action spaces. Define $u'(\cdot; \sigma_{-i}) : A'(\sigma_{-i}) \to \mathbb{R}$ in the natural way by $u'_i(a_{ij}, a'_{-i}; \sigma_{-i}) = u_i(a_{ij}, b(a'_{-i}, \sigma_{-i}))$ where $b(a'_{-i}, \sigma_{-i}) \in A_{-i}$ records the pure actions taken by i's opponents $k \notin N'_{-i}(\sigma_{-i})$ and otherwise agrees with $a'_{-i} \in A'_{-i}(\sigma_{-i})$ (define u'_k for players $k \neq i$ arbitrarily). Finally, defining $\sigma'_{-i}(\sigma_{-i}) \in \Delta'_{-i} \equiv \times_{k \in N'_{-i}(\sigma_{-i})} \Delta A'_k(\sigma_{-i})$ as the natural projection of $\sigma_{-i} \in \Delta_{-i}$ onto Δ'_{-i} after dropping players $k \notin N'_{-i}(\sigma_{-i})$ and zeros corresponding to $\sigma_{kj} = 0$ for $k \in N'_{-i}(\sigma_{-i})$, it is as if player i faces $\Gamma'(\sigma_{-i}) = \{N'(\sigma_{-i}), A'(\sigma_{-i}), u'(\cdot; \sigma_{-i})\}$ with opponents who are playing a fully mixed profile $\sigma'_{-i}(\sigma_{-i}) \in \Delta'_{-i}$. By (B1), player i's beliefs do not realize with positive probability in any subset of Δ'_{-i} with zero Lebesgue measure. By assumption, $u_i(a_{ij}, a_{-i}) \neq u_i(a_{il}, a_{-i})$ for all a_{ij} a_{il}, and $a_{-i} \in A_{-i}$, and thus for no two actions a_{ij} and a_{il} is it the case that $u'_i(a_{ij}, a'_{-i}) = u'_i(a_{il}, a'_{-i})$ for all $a'_{-i} \in A'_{-i}(\sigma_{-i})$. By Lemma 8 of Friedman and Mezzetti [2005], the event that player i is indifferent between any two pure actions has zero Lebesgue measure, and hence Ψ_i is single-valued. □

Theorem 3. *Fix* Γ^m. *The set of attainable NBE is equal to the set of attainable QRE.*

Proof. To simplify the proof, we additionally assume that NBE axiom (B2′) contains a differentiability condition: for any $\bar{r} \in (0,1)$, $F^i_k(\bar{r}|r)$ is differentiable in $r \in (0,1)$. In particular, this implies that $\frac{\partial F^i_k(\bar{r}|r)}{\partial r}|_{\bar{r}, r \in (0,1)} < 0$ by (B3′). Including differentiability has no effect on the result, as it does not effect the set of attainable NBE; it simplifies the proof because of an analogous differentiability condition assumed in QRE axiom (A2), which does not effect the set of QRE.

The proof proceeds by construction; for every NBE (satisfying (B1′)-(B4′)), we construct the corresponding QRE (satisfying (A1)-(A4)) and vice versa.

[53] In the example from Footnote 52, this construction implies the strategy $s_i(r') = 1$ (corresponding to U) for all $r' \in [0, \epsilon)$.

Step 1: *Every NBE is a QRE.*

Fix $\{\Gamma^m, \sigma^*\}$. Player i's belief-map r^* induces NBE reaction function $\Psi_{ij} : [0,1] \to [0,1]$. By Theorem 1, all NBE are interior, so the unique NBE must be a fixed point of $\Psi = (\Psi_{1j}, \Psi_{2l}) : [\epsilon, 1-\epsilon]^2 \to [\epsilon, 1-\epsilon]^2$ for sufficiently small $\epsilon > 0$. For convenience, define $U_i(\epsilon) \equiv \bar{u}_i([\epsilon, 1-\epsilon]) = (\bar{u}_{i1}(r), \bar{u}_{i2}(r))_{r\in[\epsilon,1-\epsilon]} \subset \mathbb{R}^2$ as the set of utility vectors associated with any belief $r \in [\epsilon, 1-\epsilon]$.

Step 1a: Construct a *pre-quantal response function* $\tilde{Q}_{ij} : U_i(\epsilon) \to [0,1]$ such that $\tilde{Q}_{ij} \circ \bar{u}_i|_{[\epsilon,1-\epsilon]} = \Psi_{ij}|_{[\epsilon,1-\epsilon]}$ and $\tilde{Q}_{ij}$ satisfies analogues of (A1)-(A4):

(A1°): $\tilde{Q}_{ij} \circ \bar{u}_i(r) \in (0,1)$ for all $r \in [\epsilon, 1-\epsilon]$.
(A2°): $\tilde{Q}_{ij} \circ \bar{u}_i(r)$ is a continuous and differentiable function for all $r \in [\epsilon, 1-\epsilon]$.
(A3°): $\frac{\partial \tilde{Q}_{i1} \circ \bar{u}_i(r)}{\partial r} > 0$, $\frac{\partial \tilde{Q}_{i2} \circ \bar{u}_i(r)}{\partial r} < 0$ for all $r \in [\epsilon, 1-\epsilon]$ (given $\frac{\partial \bar{u}_{i1}(r)}{\partial r} > 0$, $\frac{\partial \bar{u}_{i2}(r)}{\partial r} < 0$ without loss).
(A4°): For $r \in [\epsilon, 1-\epsilon]$ such that $\bar{u}_{ij}(r) > \bar{u}_{il}(r)$, $\tilde{Q}_{ij} \circ \bar{u}_i(r) > \tilde{Q}_{il} \circ \bar{u}_i(r)$.

From this the result almost follows. Intuitively, $\tilde{Q}_{ij}$ is very much like a quantal response function but is restricted to the subset of $\mathbb{R}^2$ that is relevant for equilibrium in this game, $\tilde{Q}_{ij} \circ \bar{u}_i$ is a more convenient reparameterization, and (A1°)-(A4°) are just (A1)-(A4) restricted to the relevant space. Once $\tilde{Q}_{ij}$ is constructed for both players $i \in \{1,2\}$, the fixed point of Ψ representing the NBE is also the fixed point of $(\tilde{Q}_{1j} \circ \bar{u}_1, \tilde{Q}_{2l} \circ \bar{u}_2) : [\epsilon, 1-\epsilon]^2 \to (0,1)^2$ representing the corresponding QRE. All that remains is to extend $\tilde{Q}_{ij}$ to a proper quantal response function defined over $\mathbb{R}^2$ that satisfies (A1)-(A4), which we do in *step 1b*.

Take $\tilde{Q}_{ij} : U_i(\epsilon) \to [0,1]$ defined by $\tilde{Q}_{ij}(v_i) \equiv \Psi_{ij}(\bar{u}_i^{-1}(v_i))$ as the pre-quantal response function, which satisfies $\tilde{Q}_{ij} \circ \bar{u}_i|_{[\epsilon,1-\epsilon]} = \Psi_{ij}|_{[\epsilon,1-\epsilon]}$ by construction. We now show that $\tilde{Q}_{ij}$ satisfies (A1°)-(A4°). We make extensive use of the fact that (without loss) $\tilde{Q}_{i1} \circ \bar{u}_i(r) = \Psi_{i1}(r) = 1 - F_k^i(\bar{r}|r)$ and $\tilde{Q}_{i2} \circ \bar{u}_i(r) = \Psi_{i2}(r) = F_k^i(\bar{r}|r)$ where $\bar{r} \in (\epsilon, 1-\epsilon)$ is the unique value that satisfies $\bar{u}_{i1}(\bar{r}) = \bar{u}_{i2}(\bar{r})$.

(A1°): $\tilde{Q}_{ij}$ satisfies (A1°) because $\tilde{Q}_{i1} \circ \bar{u}_i(r) = 1 - F_k^i(\bar{r}|r) \in (0,1)$ for $r \in [\epsilon, 1-\epsilon]$ by (B1′).
(A2°): $\tilde{Q}_{ij}$ satisfies (A2°) because $\tilde{Q}_{i1} \circ \bar{u}_i(r) = 1 - F_k^i(\bar{r}|r)$ is continuous and differentiable for all $r \in [\epsilon, 1-\epsilon]$ by (B2′).
(A3°): Without loss of generality, for all $r \in [\epsilon, 1-\epsilon]$: $\frac{\partial \bar{u}_{i1}(r)}{\partial r} > 0$, $\frac{\partial \bar{u}_{i2}(r)}{\partial r} < 0$, $\tilde{Q}_{i1} \circ \bar{u}_i(r) = 1 - F_k^i(\bar{r}|r)$, and $\tilde{Q}_{i2} \circ \bar{u}_i(r) = F_k^i(\bar{r}|r)$. That $\tilde{Q}_{ij}$ satisfies (A3°) follows because $\frac{\partial F_k^i(\bar{r}|r)}{\partial r}|_{\bar{r},r\in[\epsilon,1-\epsilon]} < 0$ from (B3′).
(A4)°: Recall that $\bar{u}_{i1}(r) = \bar{u}_{i2}(r)$ if and only if $r = \bar{r}$. Notice that by (B4′), $\tilde{Q}_{i1} \circ \bar{u}_i(\bar{r}) = 1 - F_k^i(\bar{r}|\bar{r}) = 1 - \frac{1}{2} = \frac{1}{2}$ and $\tilde{Q}_{i2} \circ \bar{u}_i(\bar{r}) = F_k^i(\bar{r}|\bar{r}) = \frac{1}{2}$. Hence, by (B3′), $\tilde{Q}_{i1} \circ \bar{u}_i(r) = \tilde{Q}_{i2} \circ \bar{u}_i(r)$ if and only if $r = \bar{r}$. (A4°) then follows from (A3°).

Step 1b: Extend $\tilde{Q}_{ij} : U_i(\epsilon) \to [0,1]$ to a proper quantal response function $Q_{ij} : \mathbb{R}^2 \to [0,1]$ that satisfies (A1)-(A4):

We now construct the extension, which we illustrate in Figure 13. First, assume without loss that $\bar{u}_{i1}$ is strictly increasing (and $\bar{u}_{i2}$ is strictly decreasing) in r. Define $U_i(-\infty) \equiv (\bar{u}_{i1}(r), \bar{u}_{i2}(r))_{r\in(-\infty,\infty)}$ as the one-dimensional affine plane that results from evaluating the expected utility vector for any r on the real line. Choose some function $Q_{i1} : U_i(-\infty) \to (0,1)$ such that $Q_{i1} \circ \bar{u}_i : (-\infty, \infty) \to (0,1)$ agrees with $\tilde{Q}_{i1} \circ \bar{u}_i(r)$ on $r \in [\epsilon, 1-\epsilon]$ and is strictly increasing and differentiable on $r \in (-\infty, \infty)$,

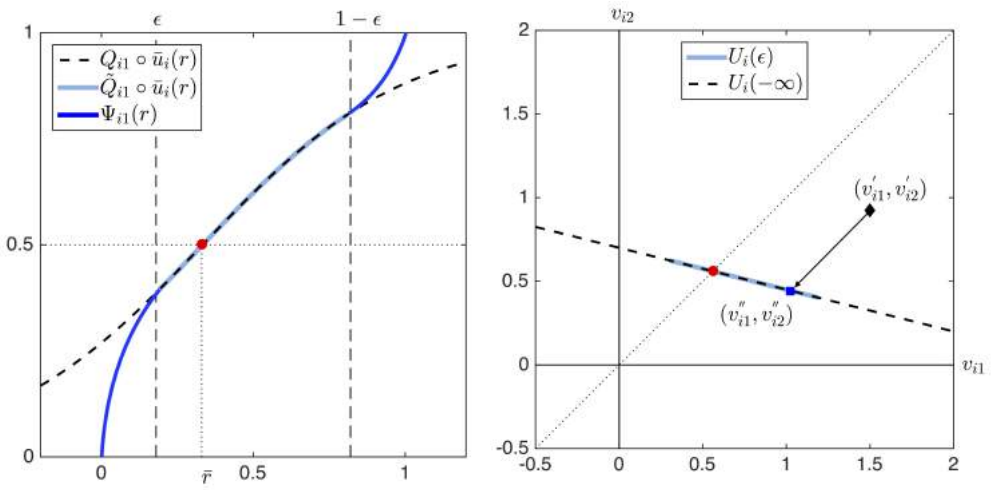

Figure 13: *Construction of the Quantal Response Function.*

which is possible because $\tilde{Q}_i$ satisfies (A1°)-(A3°) (see the left panel of Figure 13). Now extend Q_{i1} to $\mathbb{R}^2$ as follows. For any $(v'_{i1}, v'_{i2}) \in \mathbb{R}^2$, define $Q_{i1}(v'_{i1}, v'_{i2}) \equiv Q_{i1}(v''_{i1}, v''_{i2})$ where (v''_{i1}, v''_{i2}) is the projection of (v'_{i1}, v'_{i2}) along the 45°-line onto subspace $U_i(-\infty)$ (see right panel of Figure 13). It is easy to verify that Q_{ij} satisfies (A1)-(A4).

Step 2: *Every QRE is an NBE.*

We are now given quantal response function $Q_{ij} : \mathbb{R}^2 \to [0,1]$. First, we construct a family of CDFs $F^i_k(\cdot|r)$ representing belief-map $r^*(r)$. We then show that $r^*(r)$ induces a reaction function $\Psi_{ij} : [0,1] \to [0,1]$ such that $\Psi_{ij}|_{[\epsilon,1-\epsilon]} = Q_{ij} \circ \bar{u}_i|_{[\epsilon,1-\epsilon]}$ and that $r^*(r)$ satisfies (B1′)-(B4′), from which the result follows.

We may assume $\bar{u}_{i1}(r)$ and $Q_{i1} \circ \bar{u}_i(r)$ are strictly decreasing in $r \in [0,1]$ without loss by (A3). For the unique $\bar{r} \in (\epsilon, 1-\epsilon)$ such that $\bar{u}_{i1}(\bar{r}) = \bar{u}_{i2}(\bar{r})$, define

$$F^i_k(\bar{r}|r) \equiv \begin{cases} g(r) & r \in [0,\epsilon) \\ Q_{i1} \circ \bar{u}_i(r) & r \in [\epsilon, 1-\epsilon] \\ h(r) & r \in (1-\epsilon, 1], \end{cases}$$

where $g(r)$ and $h(r)$ are *any* functions chosen so that the whole function is strictly decreasing, continuous, differentiable, and $F^i_k(\bar{r}|0) = 1$ and $F^i_k(\bar{r}|1) = 0$. That this is possible relies on (A1)-(A3). Notice that $F^i_k(\bar{r}|r)|_{[\epsilon,1-\epsilon]} = Q_{i1} \circ \bar{u}_i(r)|_{[\epsilon,1-\epsilon]}$ and $F^i_k(\bar{r}|r)$ goes from the top left corner of the unit square to the bottom right corner. Also, by (A4), we have that $F^i_k(\bar{r}|\bar{r}) = Q_{i1} \circ \bar{u}_i(\bar{r}) = \frac{1}{2}$. Now choose positive number n_0 sufficiently large such that $(1-r)^{n_0} < F^i_k(\bar{r}|r) < 1 - r^{n_0}$ for all $r \in (0,1)$, which exists since as $n \to \infty$, $(1-r)^n \to 0$ and $(1-r^n) \to 1$ pointwise on $r \in (0,1)$. Figure 14 gives an illustration of the functions defined so far, and is a useful reference for the whole construction.

Define $r_\circ \equiv \{r : (1-r)^{n_0} = \frac{1}{2}\}$ and $r^\circ \equiv \{r : 1 - r^{n_0} = \frac{1}{2}\}$, and notice that $r_\circ < \epsilon < \bar{r} < 1 - \epsilon < r^\circ$. For all $\tilde{r} \in [r_\circ, \bar{r}]$, define $\alpha(\tilde{r}) \equiv \{\alpha \in [0,1] : \alpha F^i_k(\bar{r}|\tilde{r}) + (1-\alpha)(1-\tilde{r})^{n_0} = \frac{1}{2}\}$ and

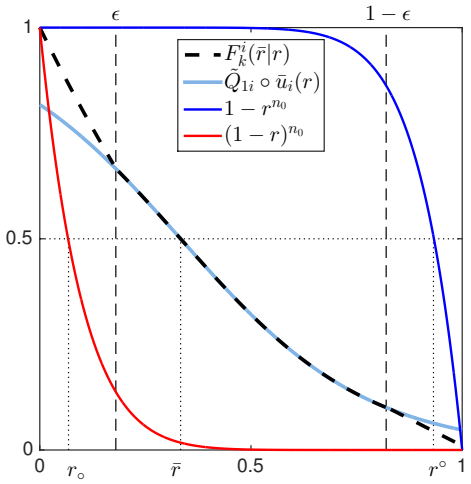

Figure 14: *Construction of belief-map.*

$F_k^i(\tilde{r}|r) \equiv \alpha(\tilde{r})F_k^i(\bar{r}|r) + (1-\alpha(\tilde{r}))(1-r)^{n_0}$. Similarly, for all $\tilde{r} \in [\bar{r}, r^\circ]$, define $\beta(\tilde{r}) \equiv \{\beta \in [0,1] : \beta F_k^i(\bar{r}|\tilde{r}) + (1-\beta)(1-(\tilde{r})^{n_0}) = \frac{1}{2}\}$ and $F_k^i(\tilde{r}|r) \equiv \beta(\tilde{r})F_k^i(\bar{r}|r) + (1-\beta(\tilde{r}))(1-r^{n_0})$. Now for $\tilde{r} \in (0, r_\circ)$, define $F_k^i(\tilde{r}|r) \equiv (1-r)^{m(\tilde{r})}$ where $m(\tilde{r}) \equiv \{m \in [n_0, \infty) : (1-\tilde{r})^m = \frac{1}{2}\}$. Similarly, for $\tilde{r} \in (r^\circ, 1)$, define $F_k^i(\tilde{r}|r) \equiv 1 - r^{n(\tilde{r})}$ where $n(\tilde{r}) \equiv \{n \in [n_0, \infty) : 1-\tilde{r}^n = \frac{1}{2}\}$. Finally, set $F_k^i(0|r) = 0$ and $F_k^i(1|r) = 1$ for $r \in (0,1)$. We have defined a family of CDFs $F_k^i(\tilde{r}|r)_{\tilde{r}\in[0,1], r\in(0,1)}$, pinning down belief-map $r^*(r)$ for all $r \in (0,1)$. We may also impose that $r^*(0) = 0$ and $r^*(1) = 1$, which gives $F_k^i(\tilde{r}|1) = \mathbf{1}_{\{\tilde{r}=1\}}$ and $F_k^i(\tilde{r}|0) = 1$, and thus we have constructed the entire family $F_k^i(\tilde{r}|r)_{\tilde{r}\in[0,1], r\in[0,1]}$. The NBE reaction is now given by $\Psi_{i1}(r) = F_k^i(\bar{r}|r)$ and $\Psi_{i2}(r) = 1 - F_k^i(\bar{r}|r)$, which by construction satisfies $\Psi_{ij}|_{[\epsilon,1-\epsilon]} = Q_{ij} \circ \bar{u}_i|_{[\epsilon,1-\epsilon]}$. Finally, that $r^*(r)$ satisfies (B1′)-(B4′) is immediate from construction of $F_k^i(\cdot|r)$. □

Theorem 5. *Fix* $\{\Gamma, \sigma^*\}$. *The set of NBE is the same for all* $\Gamma' \in \mathcal{A}(\Gamma)$.

Proof. Fix $\{\Gamma, \sigma^*\}$. First, we show that response set R_{ij} is the same for all $\Gamma' \in \mathcal{A}(\Gamma)$. By definition of $\mathcal{A}$, for all i and a_{-i}, there exists β_i and $\gamma_i(a_{-i})$ such that $\bar{u}'_{ij}(a_{-i}) = \beta_i \bar{u}_{ij}(a_{-i}) + \gamma_i(a_{-i})$ for all j. By linearity of expected utility, for all i and $\tilde{\sigma}_{-i}$, β_i and $\gamma_i(\tilde{\sigma}_{-i}) \equiv \sum_{a_{-i}} \tilde{\sigma}_{-i}(a_{-i})\gamma_i(a_{-i})$ satisfy $\bar{u}'_{ij}(\tilde{\sigma}_{-i}) = \beta_i \bar{u}_{ij}(\tilde{\sigma}_{-i}) + \gamma_i(\tilde{\sigma}_{-i})$ for all j. Thus, we have that

$$\begin{aligned}
R'_{ij} &= \{\tilde{\sigma}_{-i} : \ \bar{u}'_{ij}(\tilde{\sigma}_{-i}) \geq \bar{u}'_{ik}(\tilde{\sigma}_{-i}) \ \forall k = 1, ..., J(i)\} \\
&= \{\tilde{\sigma}_{-i} : \ \beta_i \bar{u}_{ij}(\tilde{\sigma}_{-i}) + \bar{\gamma}_i(\tilde{\sigma}_{-i}) \geq \beta_i \bar{u}_{ik}(\tilde{\sigma}_{-i}) + \bar{\gamma}_i(\tilde{\sigma}_{-i}) \ \forall k = 1, ..., J(i)\} \\
&= \{\tilde{\sigma}_{-i} : \ \beta_i \bar{u}_{ij}(\tilde{\sigma}_{-i}) \geq \beta_i \bar{u}_{ik}(\tilde{\sigma}_{-i}) \ \forall k = 1, ..., J(i)\} \\
&= \{\tilde{\sigma}_{-i} : \ \bar{u}_{ij}(\tilde{\sigma}_{-i}) \geq \bar{u}_{ik}(\tilde{\sigma}_{-i}) \ \forall k = 1, ..., J(i)\} \\
&= R_{ij}.
\end{aligned}$$

It is immediate that, for any belief-map σ^*, NBE reaction Ψ, and thus any NBE, is the same for all $\Gamma' \in \mathcal{A}(\Gamma)$. □

Fact 1. *$r^*(r;\tau)$ has CDF*

$$F_k^i(\bar{r}|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right).$$

Proof.

$$\begin{aligned} F_k^i(\bar{r}|r;\tau) \equiv \mathbb{P}(r^*(r;\tau) \leq \bar{r}) &= \mathbb{P}\left(\frac{exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)}{1 + exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)} \leq \bar{r}\right) \\ &= \mathbb{P}\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i \leq ln\left(\frac{\bar{r}}{1-\bar{r}}\right)\right) \\ &= \mathbb{P}\left(\varepsilon_i \leq \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) \\ &= \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) \end{aligned}$$

Notice that F_k^i is differentiable in $\bar{r}$ for $\bar{r}, r \in (0,1)$. Hence, the PDF is easily derived as $f_k^i(\bar{r}|r;\tau) \equiv \frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial \bar{r}}|_{\bar{r},r\in(0,1)}$ using the chain rule. □

Fact 2. *$r^*(r;\tau)$ satisfies (B1′)-(B4′).*

Proof.
(B1) For any $r \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is strictly increasing and continuous in $\bar{r} \in [0,1]$; $r^*(0;\tau) = 0$ and $r^*(1;\tau) = 1$:

That $r^*(0;\tau) = 0$ and $r^*(1;\tau) = 1$ is obvious from the definition of $r^*(\cdot;\tau)$ and the convention that $\mathcal{L}(0) = -\infty$ and $\mathcal{L}(1) = \infty$ where $\mathcal{L}(r) = ln\left(\frac{r}{1-r}\right)$. It is also obvious that, for any $r \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is continuous in $\bar{r} \in [0,1]$. For all $r \in (0,1)$, $F_k^i(0|r;\tau) = 0$ and $F_k^i(1|r;\tau) = 1$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$). All we need to show is that $F_k^i(\bar{r}|r;\tau)$ is strictly increasing in $\bar{r} \in (0,1)$ for all $r \in (0,1)$. Notice that $\frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial \bar{r}}|_{\bar{r},r\in(0,1)} = \phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)\frac{1}{\tau}\left(\frac{1}{\bar{r}(1-\bar{r})}\right) > 0$ since $\phi(\cdot) > 0$, and we are done.

(B2′) For any $\bar{r} \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is continuous in $r \in [0,1]$:

We show something stronger, that $F_k^i(\bar{r}|r;\tau)$ is jointly continuous in $(\bar{r},r) \in (0,1)\times[0,1]$. $F_k^i(\bar{r}|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)$ is obviously continuous for every $(\bar{r},r) \in (0,1)\times(0,1)$. $F_k^i(\bar{r}|r;\tau)$ is also continuous at all points $(\bar{r},r) \in (0,1)\times\{0,1\}$. To see this, notice that $F_k^i(\bar{r}|0;\tau) = 1$ for all $\bar{r} \in (0,1)$ and $lim_{r\to 0^+} F_k^i(\bar{r}|r;\tau) = 1$ for all $\bar{r} \in (0,1)$, showing continuity at $(\bar{r},r) \in (0,1)\times\{0\}$. A similar argument shows continuity at $(\bar{r},r) \in (0,1)\times\{1\}$.

(B3′) For all $r < r' \in [0,1]$, $F_k^i(\bar{r}|r';\tau) \leq F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in [0,1]$ and $F_k^i(\bar{r}|r';\tau) < F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in (0,1)$:

(i) If $r' > r \in (0,1)$:

(a) If $\bar{r} \in (0,1)$,

$$\begin{aligned} F_k^i(\bar{r}|r';\tau) <& F_k^i(\bar{r}|r;\tau) \iff \\ \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r'}{1-r'}\right)\right]\right) <& \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) \iff \\ \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r'}{1-r'}\right)\right] <& \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right] \iff \\ r' >& r. \end{aligned}$$

(b) If $\bar{r} = 0$, $F_k^i(\bar{r}|r;\tau) = F_k^i(\bar{r}|r';\tau) = 0$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$).
(c) If $\bar{r} = 1$, $F_k^i(\bar{r}|r;\tau) = F_k^i(\bar{r}|r';\tau) = 1$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$).
(ii) If $1 = r' > r > 0$, $F_k^i(\bar{r}|r';\tau) = \mathbf{1}_{\{\bar{r}=1\}} \leq F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in [0,1]$ (using $r^*(1;\tau) = 1$).
(iii) If $1 > r' > r = 0$, $F_k^i(\bar{r}|r';\tau) \leq F_k^i(\bar{r}|r;\tau) = 1$ for $\bar{r} \in [0,1]$ (using $r^*(0;\tau) = 0$).
Finally, $\frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial r}|_{\bar{r},r\in(0,1)} = -\phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)\frac{1}{\tau}\left(\frac{1}{r(1-r)}\right) < 0$ since $\phi(\cdot) > 0$.

(B4′) $F_k^i(r|r;\tau) = \frac{1}{2}$ for $r \in (0,1)$:
For $r \in (0,1)$, $F_k^i(r|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{r}{1-r}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) = \Phi(0) = \frac{1}{2}$. □

8.2 The NBE of generalized matching pennies

We derive the set of NBE (and hence QRE) attainable for arbitrary Γ^m.

Along the lines of example 1, it is easy to show that the reactions in Γ^m depend only on the Nash equilibrium $\{\sigma_U^{NE}, \sigma_L^{NE}\}$ and satisfy:

$$\Psi_U(\sigma_L) \in \begin{cases} \{0\} & \sigma_L = 0 \\ (0, \frac{1}{2}) & \sigma_L \in (0, \sigma_L^{NE}) \\ \{\frac{1}{2}\} & \sigma_L = \sigma_L^{NE} \\ (\frac{1}{2}, 1) & \sigma_L \in (\sigma_L^{NE}, 1) \\ \{1\} & \sigma_L = 1 \end{cases} \qquad \Psi_L(\sigma_U) \in \begin{cases} \{1\} & \sigma_U = 0 \\ (\frac{1}{2}, 1) & \sigma_U \in (0, \sigma_U^{NE}) \\ \{\frac{1}{2}\} & \sigma_U = \sigma_U^{NE} \\ (0, \frac{1}{2}) & \sigma_U \in (\sigma_U^{NE}, 1) \\ \{0\} & \sigma_U = 1. \end{cases}$$

The set of attainable NBE is given by $\{\{\sigma_U, \sigma_L\} | \sigma_U \in \Psi_U(\sigma_L),\ \sigma_L \in \Psi_L(\sigma_U)\}$ and consists of the union of one or more rectangles of positive measure, except when $\{\sigma_U^{NE}, \sigma_L^{NE}\} = \{\frac{1}{2}, \frac{1}{2}\}$ in which case the unique NBE is $\{\sigma_U, \sigma_L\} = \{\frac{1}{2}, \frac{1}{2}\}$.

8.2.1 Mean-unbiasedness

We derive the set of attainable NBE when *unbiasedness* (B4′) is replaced with *mean*-unbiasedness.

We first derive the upper and lower bounds on player 1's reaction function $\Psi_U(\sigma_L)$ under the restriction that belief distributions are correct on mean. That is, we find $\bar{\Psi}_U(\sigma_L) \equiv \underset{\sigma_L^* |\ \mathbb{E}(\sigma_L^*) = \sigma_L}{sup} \mathbb{P}(\sigma_L^*(\sigma_L) \geq$

$\sigma_L^{NE})$ and $\underline{\Psi}_U(\sigma_L) \equiv \inf_{\sigma_L^*|\ \mathbb{E}(\sigma_L^*)=\sigma_L} \mathbb{P}(\sigma_L^*(\sigma_L) \geq \sigma_L^{NE})$. These bounds can be achieved through the following family of two-atom belief distributions:

$$\hat{\sigma}_L^*(\sigma_L) = \begin{cases} \underline{\sigma}_L(\sigma_L) & \text{with probability } 1-\alpha(\sigma_L) \\ \bar{\sigma}_L(\sigma_L) & \text{with probability } \alpha(\sigma_L) \end{cases}.$$

This violates *continuity* (B1′), but it is clear that a continuous version can approximate arbitrarily well the reactions they induce, and hence it is sufficient to find $\bar{\Psi}_U(\sigma_L) \equiv \sup_{\underline{\sigma}_L,\bar{\sigma}_L,\alpha|(1-\alpha)\underline{\sigma}_L+\alpha\bar{\sigma}_L=\sigma_L} \mathbb{P}(\hat{\sigma}_L^*(\sigma_L) \geq \sigma_L^{NE})$ and $\underline{\Psi}_U(\sigma_L) \equiv \inf_{\underline{\sigma}_L,\bar{\sigma}_L,\alpha|(1-\alpha)\underline{\sigma}_L+\alpha\bar{\sigma}_L=\sigma_L} \mathbb{P}(\hat{\sigma}_L^*(\sigma_L) \geq \sigma_L^{NE})$.

Case 1: $\sigma_L \in (0, \sigma_L^{NE})$. It is obvious that $\underline{\Psi}_U(\sigma_L) = 0$. It is easy to check that $\bar{\Psi}_U(\sigma_L)$ is achieved when $\underline{\sigma}_L = 0$ and $\bar{\sigma}_L = \sigma_L^{NE}$, and thus α is determined by the constraint $(1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = \alpha\sigma_L^{NE} = \sigma_L$, which implies $\bar{\Psi}_U(\sigma_L) = \alpha = \frac{\sigma_L}{\sigma_L^{NE}}$.

Case 2: $\sigma_L \in (\sigma_L^{NE}, 1)$. It is obvious that $\bar{\Psi}_U(\sigma_L) = 1$. It is easy to check that $\underline{\Psi}_U(\sigma_L)$ is achieved when $\underline{\sigma}_L = \sigma_L^{NE}$ and $\bar{\sigma}_L = 1$, and thus α is determined by the constraint $(1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = (1-\alpha)\sigma_L^{NE} + \alpha = \sigma_L$, which implies $\underline{\Psi}_U(\sigma_L) = \alpha = \frac{\sigma_L - \sigma_L^{NE}}{1-\sigma_L^{NE}}$.

The knife's edge cases $\sigma_L \in \{0, \sigma_L^{NE}, 1\}$ are obvious. Summarizing, and giving the analogue for player 2:

$$\Psi_U^{mean}(\sigma_L) \in \begin{cases} \{0\} & \sigma_L = 0 \\ (0, \frac{\sigma_L}{\sigma_L^{NE}}) & \sigma_L \in (0, \sigma_L^{NE}) \\ (0,1) & \sigma_L = \sigma_L^{NE} \\ (\frac{\sigma_L - \sigma_L^{NE}}{1-\sigma_L^{NE}}, 1) & \sigma_L \in (\sigma_L^{NE}, 1) \\ \{1\} & \sigma_L = 1 \end{cases} \qquad \Psi_L^{mean}(\sigma_U) \in \begin{cases} \{1\} & \sigma_U = 0 \\ (1 - \frac{\sigma_U}{\sigma_U^{NE}}, 1) & \sigma_U \in (0, \sigma_U^{NE}) \\ (0,1) & \sigma_U = \sigma_U^{NE} \\ (0, \frac{1-\sigma_U}{1-\sigma_U^{NE}}) & \sigma_U \in (\sigma_U^{NE}, 1) \\ \{0\} & \sigma_U = 1. \end{cases}$$

The set of attainable "mean NBE" is given by $\{\{\sigma_U, \sigma_L\} | \sigma_U \in \Psi_U^{mean}(\sigma_L),\ \sigma_L \in \Psi_L^{mean}(\sigma_U)\}$ and consists of a single diamond region which has positive measure for all $\{\sigma_U^{NE}, \sigma_L^{NE}\}$ including the case that $\{\sigma_U^{NE}, \sigma_L^{NE}\} = \{\frac{1}{2}, \frac{1}{2}\}$. The set always contains an open ball around the NE, and hence has a non-trivial intersection with the set of NBE/QRE.

8.3 Generalizing Lemma 3

For the statement of the lemma and its proof, we let $v_i \in \mathbb{R}_{++}^{J(i)}$ be a utility vector with strictly positive components, where, without loss, $v_{i1} \geq v_{i2} \geq \cdots \geq v_{iJ(i)}$. Let $J^+(v_i) \equiv \{j :\ v_{ij} \geq v_{ik}\ \forall k\}$ and $J^-(v_i) \equiv \{j :\ v_{ij} \leq v_{ik}\ \forall k\}$ be the indices corresponding to the highest and lowest payoff components respectively. $J^+(v_i) \cap J^-(v_i) = \emptyset$ if and only if $v_{ij} \neq v_{ik}$ for some j and k.

Lemma. *Let $v_i \in \mathbb{R}_{++}^{J(i)}$ be such that $J^+(u_i') \cap J^-(u_i') = \emptyset$.*
(i) Let Q_i be translation invariant (and weakly substitutable) and $\beta > 1$:
(a) $Q_i(\beta v_i) = Q_i(\tilde{v}_i(\beta))$ for some $\tilde{v}_i(\beta)$ such that $\tilde{v}_{il}(\beta) = v_{il} + \delta_l(\beta)$ where $\delta_l(\beta) = 0$

if $l \in J^-(v_i)$, $\delta_l(\beta) > 0$ *and* $\delta_l(\beta) \to \infty$ *if* $l \notin J^-(v_i)$, *and* $\delta_j(\beta) - \delta_k(\beta) \to \infty$ *if* $v_{ij} > v_{ik}$.

(b) $Q_{ij}(\beta v_i) > Q_{ij}(v_i)$ *for all* $j \in J^+(v_i)$ *and* $Q_{ik}(\beta v_i) < Q_{ik}(v_i)$ *for all* $k \in J^-(v_i)$.

(c) $\lim_{\beta\to\infty} Q_{ij}(\beta v_i) \geq \lim_{x\to\infty} Q_{ij}(x, ..., x, 0, ..., 0)$ *if* $j \in J^+(v_i)$ *(*x *in first* $|J^+(v_i)|$ *entries);*
$\lim_{\beta\to\infty} Q_{ik}(\beta v_i) \leq \lim_{x\to\infty} Q_{ik}(0, .., 0, -x, ..., -x)$ *if* $k \in J^-(v_i)$ *(*$-x$ *in last* $|J^-(v_i)|$ *entries).*

(ii) Let Q_i *be scale invariant (and weakly substitutable) and* $\gamma > 0$:

(a) $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\tilde{v}_i(\gamma))$ *for some* $\tilde{v}_i(\gamma)$ *such that* $\tilde{v}_{il}(\gamma) = v_{il} + \delta_l(\gamma)$ *where* $\delta_l(\gamma) = 0$ *if* $l \in J^+(v_i)$, $\delta_l(\gamma) > 0$ *and* $\delta_l(\gamma) \to v_{i1} - v_{il}$ *if* $l \notin J^+(v_i)$.

(b) $Q_{ij}(v_i+\gamma e_{J(i)}) < Q_{ij}(v_i)$ *for all* $j \in J^+(v_i)$ *and* $Q_{ik}(v_i+\gamma e_{J(i)}) > Q_{ik}(v_i)$ *for all* $k \in J^-(v_i)$.

(c) $\lim_{\gamma\to\infty} Q_{il}(v_i + \gamma e_{J(i)}) = \frac{1}{J(i)}$ *for all* l.

Proof. (i): Fix v_i with $J^+(v_i) \cap J^-(v_i) = \emptyset$ and let $\beta > 1$. It is easy to show that $\bar{\gamma}(\beta) \equiv (\beta-1)v_{i1} > 0$ and $\underline{\gamma}(\beta) \equiv (\beta-1)v_{iJ(i)} > 0$ satisfy $\beta v_{ij} - \bar{\gamma}(\beta) = v_{ij}$ for all $j \in J^+(v_i)$ and $\beta v_{ik} - \underline{\gamma}(\beta) = v_{ik}$ for all $k \in J^-(v_i)$. Notice that $v_{i1} > v_{ik} \iff \beta v_{ik} - \bar{\gamma}(\beta) < v_{ik}$ and $v_{iJ(i)} < v_{ij} \iff \beta v_{ij} - \underline{\gamma}(\beta) > v_{ij}$ and thus $\beta v_{ik} - \bar{\gamma}(\beta) < v_{ik}$ for all $k \notin J^+(v_i)$ and $\beta v_{ij} - \underline{\gamma}(\beta) > v_{ij}$ for all $j \notin J^-(v_i)$.

By translation invariance, $Q_i(\beta v_i) = Q_i(\beta v_i - \underline{\gamma}(\beta)e_{J(i)}) = Q_i(\beta v_i - \bar{\gamma}(\beta)e_{J(i)})$. Since $Q_i(\beta v_i) = Q_i(\beta v_i - \underline{\gamma}(\beta)e_{J(i)})$, we have that $Q_i(\beta v_i) = Q_i(\tilde{v}(\beta))$ where $\tilde{v}_{il}(\beta) \equiv \beta v_{il} - \underline{\gamma}(\beta) = \beta v_{il} - (\beta - 1)v_{iJ(i)} = v_{il} + \delta_l(\beta)$ where $\delta_l(\beta) = (\beta-1)(v_{il} - v_{iJ(i)})$. This shows *(a)*.

By weak substitutability, $Q_{ij}(\beta v_i - \bar{\gamma}(\beta)e_{J(i)}) > Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\beta v_i - \underline{\gamma}(\beta)e_{J(i)}) < Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. Therefore, $Q_{ij}(\beta v_i) > Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\beta v_i) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. This shows *(b)*.

By translation invariance, $Q_i(\tilde{v}_i(\beta)) = Q_i(\tilde{v}_i(\beta) - \tilde{\gamma}(\beta)e_{J(i)})$ for any $\tilde{\gamma}(\beta)$. Since $\delta_j(\beta) - \delta_k(\beta) \to \infty$ if $v_{ij} > v_{ik}$, we can set $\tilde{\gamma}(\beta)$ such that $\tilde{v}_{ij}(\beta) - \tilde{\gamma}(\beta) \to \infty$ for $j \in J^+(v_i)$ and $\tilde{v}_{ik}(\beta) - \tilde{\gamma}(\beta) \to -\infty$ for all $k \notin J^+(v_i)$. By weak substitutability $\lim_{\beta\to\infty} Q_{ij}(\tilde{v}_i(\beta) - \tilde{\gamma}(\beta)e_{J(i)}) \geq \lim_{x\to\infty} Q_{ij}(x, ..., x, 0, ..., 0)$ (x in first $|J^+(v_i)|$ entries) and thus $\lim_{\beta\to\infty} Q_{ij}(\beta v_i) \geq \lim_{x\to\infty} Q_{ij}(x, ..., x, 0, ..., 0)$ if $j \in J^+(v_i)$. Similarly, one can show that $\lim_{\beta\to\infty} Q_{ik}(\beta v_i) \leq \lim_{x\to\infty} Q_{ik}(0, .., 0, -x, ..., -x)$ if $k \in J^-(v_i)$, which shows *(c)*.

(ii): Fix v_i with $J^+(v_i) \cap J^-(v_i) = \emptyset$ and let $\gamma > 0$. It is easy to show that $\bar{\beta}(\gamma) \equiv \frac{v_{i1}}{v_{i1}+\gamma} \in (0, 1)$ and $\underline{\beta}(\gamma) \equiv \frac{v_{iJ(i)}}{v_{iJ(i)}+\gamma} \in (0, 1)$ satisfy $\bar{\beta}(\gamma)(v_{ij}+\gamma) = v_{ij}$ for all $j \in J^+(v_i)$ and $\underline{\beta}(\gamma)(v_{ik}+\gamma) = v_{ik}$ for all $k \in J^-(v_i)$. Notice that $v_{i1} > v_{ik} \iff \bar{\beta}(\gamma)(v_{ik}+\gamma) > v_{ik}$ and $v_{iJ(i)} < v_{ij} \iff \underline{\beta}(\gamma)(v_{ij}+\gamma) < v_{ij}$ and thus $\bar{\beta}(\gamma)(v_{ik} + \gamma) > v_{ik}$ for all $k \notin J^+(v_i)$ and $\underline{\beta}(\gamma)(v_{ij} + \gamma) < v_{ij}$ for all $j \notin J^-(v_i)$.

By scale invariance, $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)})) = Q_i(\bar{\beta}(\gamma)(v_i + \gamma e_{J(i)}))$. Since $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)}))$, we have that $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\tilde{v}(\gamma))$ where $\tilde{v}_{il}(\gamma) \equiv \bar{\beta}(\gamma)(v_{il} + \gamma) = \frac{v_{i1}}{v_{i1}+\gamma}(v_{il} + \gamma) = v_{il} + \delta_l(\gamma)$ where $\delta_l(\gamma) = \frac{\gamma(v_{i1}-v_{il})}{v_{i1}+\gamma}$. This shows *(a)*.

By weak substitutability, $Q_{ij}(\bar{\beta}(\gamma)(v_i + \gamma e_{J(i)})) < Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)})) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. Therefore, $Q_{ij}(v_i + \gamma e_{J(i)}) < Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(v_i + \gamma e_{J(i)}) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. This shows *(b)*.

As $\gamma \to \infty$, $\tilde{v}_{il}(\gamma) \to v_{i1}$ for all l. Thus $Q_{il}(v_i + \gamma e_{J(i)}) = Q_{il}(\tilde{v}_i(\gamma)) \to Q_{il}(v_{i1}, ..., v_{i1}) = \frac{1}{J(i)}$, which shows *(c)*. □

8.4 QRE in sets of binary action games: necessary conditions

Theorem. *Fix dataset* $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ *where* $\mathcal{G} = \{g^1, ...g^m, ..., g^M\}$ *is a set of games that differ only in payoffs with* $J(i) = 2$ *for all* i, $\hat{\sigma} = \{\hat{\sigma}_{ij}^m\}_{ijm}$ *are action frequencies, and* $\hat{u} = \{\hat{u}_{ij}^m\}_{ijm}$ *are expected utilities* ($\hat{u}_{ij}^m \equiv \bar{u}_{ij}^m(\hat{\sigma}_{-i}^m)$). *Without loss, relabel all actions so that* $\hat{u}_{i1}^m \geq \hat{u}_{i2}^m$ *for all* m *and* i.

(i) $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ *is consistent with translation invariant QRE only if, for all* i*:*

$$\hat{u}_{i1}^m - \hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'} \iff \hat{\sigma}_{i1}^m \geq \hat{\sigma}_{i1}^{m'} \ \forall m, m'$$

$$\text{and } \hat{\sigma}_{i1}^m \geq \frac{1}{2} \ \forall m.$$

(ii) $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ *is consistent with scale invariant QRE*[54] *only if, for all* i*:*

$$\hat{u}_{i1}^m / \hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'} / \hat{u}_{i2}^{m'} \iff \hat{\sigma}_{i1}^m \geq \hat{\sigma}_{i1}^{m'} \ \forall m, m'$$

$$\text{and } \hat{\sigma}_{i1}^m \geq \frac{1}{2} \ \forall m.$$

Proof. (i): By translation invariance, $Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m - \hat{u}_{i2}^m, 0)$ and $Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = Q_{i1}(\hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}, 0)$ for all m and m'. Therefore, by *responsiveness* (A3), $\hat{\sigma}_{i1}^m \equiv Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m - \hat{u}_{i2}^m, 0) \geq Q_{i1}(\hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}, 0) = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) \equiv \hat{\sigma}_{i1}^{m'} \iff \hat{u}_{i1}^m - \hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}$ for all m and m'. $\hat{u}_{i1}^m \geq \hat{u}_{i2}^m$, and hence $\hat{\sigma}_{i1}^m \equiv Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) \geq \frac{1}{2}$ for all m by *monotonicity* (A4).

(ii): By scale invariance, $Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m / \hat{u}_{i2}^m, 1)$ and $Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = Q_{i1}(\hat{u}_{i1}^{m'} / \hat{u}_{i2}^{m'}, 1)$ for all m and m'. Therefore, by *responsiveness* (A3), $\hat{\sigma}_{i1}^m \equiv Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m / \hat{u}_{i2}^m, 1) \geq Q_{i1}(\hat{u}_{i1}^{m'} / \hat{u}_{i2}^{m'}, 1) = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) \equiv \hat{\sigma}_{i1}^{m'} \iff \hat{u}_{i1}^m / \hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'} / \hat{u}_{i2}^{m'}$ for all m and m'. $\hat{u}_{i1}^m \geq \hat{u}_{i2}^m$, and hence $\hat{\sigma}_{i1}^m \equiv Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) \geq \frac{1}{2}$ for all m by *monotonicity* (A4). □

8.5 Logit transform NBE in normal form games

For arbitrary normal form games, we generalize (10) by parametrizing player i's belief-map over action j of player k as

$$\sigma_{kj}^{i*}(p_k; \tau) = \frac{exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon_{kj}^i\right)}{1 + exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon_{kj}^i\right)} \cdot \left(\sum_{l=1}^{J(k)} \frac{exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon_{kl}^i\right)}{1 + exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon_{kl}^i\right)}\right)^{-1}, \tag{15}$$

where $\varepsilon_{kj}^i \sim_{iid} \mathcal{N}(0,1)$, and $\tau \in (0, \infty)$ determines the noisiness of beliefs. This belief-map is derived through the following procedure:

[54] Here, we require also that $\hat{u}_{ij}^m > 0$ for all i, j, and m.

1. Map each $\sigma_{kj} \in [0,1]$ to the extended real line via the logit transform

$$\mathcal{L}(\sigma_{kj}) = ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right),$$

   using the convention that $\mathcal{L}(0) = -\infty$ and $\mathcal{L}(1) = \infty$.

2. Add $\tau\varepsilon^i_{kj}$ to each $\mathcal{L}(\sigma_{kj})$.

3. Map each $\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj}$ back to $[0,1]$ via the inverse logit transform

$$\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj}) = \frac{exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon^i_{kj}\right)}{1 + exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon^i_{kj}\right)}.$$

4. Normalize the set of $\{\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kl}) + \tau\varepsilon^i_{kl})\}_{l=1}^{J(k)}$ so that they sum to 1 by dividing each $\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj})$ by the sum

$$\sum_{l=1}^{J(k)} \frac{exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon^i_{kl}\right)}{1 + exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon^i_{kl}\right)}.$$

This belief-map does not satisfy *unbiasedness* (B4) exactly, but simulations (unreported) suggest that the bias is negligible for low τ, such as those estimated in the Melo et al. [2017] data in Section 6.4.

8.6 Details of games, data, and experiments

Study	Game	N	Notes		
			Rounds (R)	Subject pairs (I)	Feedback?
MPW 2000	A	1800	50	36	yes
	B	1200	50	24	yes
	C	1200	50	24	yes
	D	600	50	12	yes
SC 2008	1	9600	200	48	yes
	2	9600	200	48	yes
	3	9600	200	48	yes
	4	9600	200	48	yes
	5	9600	200	48	yes
	6	9600	200	48	yes
	7	4800	200	24	yes
	8	4800	200	24	yes
	9	4800	200	24	yes
	10	4800	200	24	yes
	11	4800	200	24	yes
	12	4800	200	24	yes
Ochs 1995	2	448	56	8	yes
	3	516	64	8	yes
GH 2001	AMP	25	1	25	no
	RA	25	1	25	no
NS 2002	NS	6720	60	112	yes
MPS 2018	2	1620	20,20,10	34,32,30	yes
	3	940	20,10	32,20	yes
	4	300	10	30	yes

Table 8: *Details of games, data, and experiments.* MPW 2000 = McKelvey et al. [2000], SC 2008 = Selten and Chmura [2008], Ochs 1995 = Ochs [1995], GH 2001 = Goeree and Holt [2001], NS 2002 = Nyarko and Schotter [2002], and MPS 2018 = Melo et al. [2017]

8.6.1 *MPW 2000* (McKelvey et al. [2000])

For any one game (see Figure 9), each subject takes an action 50 times against randomly matched opponents with feedback. Subjects play multiple games, and maintain their role as either player 1 or player 2 for the duration of the experiment. The exchange rate is \$0.10 (year 2000 Dollars).

8.6.2 *SC 2008* (Selten and Chmura [2008])

Each subject plays just one game, and takes an action 200 times against randomly matched opponents with feedback. Subjects maintain their role as either player 1 or player 2 for the duration of the experiment. The exchange rate is €0.016 (year 2008 Euros), and the experiment took place in Germany.

1

	L	R
U	10, 8	0, 18
D	9, 9	10, 8

2

	L	R
U	9, 4	0, 13
D	6, 7	8, 5

3

	L	R
U	8, 6	0, 14
D	7, 7	10, 4

4

	L	R
U	7, 4	0, 11
D	5, 6	9, 2

5

	L	R
U	7, 2	0, 9
D	4, 5	8, 1

6

	L	R
U	7, 1	1, 7
D	3, 5	8, 0

7

	L	R
U	10, 12	4, 22
D	9, 9	14, 8

8

	L	R
U	9, 7	3, 16
D	6, 7	11, 5

9

	L	R
U	8, 9	3, 17
D	7, 7	13, 4

10

	L	R
U	7, 6	2, 13
D	5, 6	11, 2

11

	L	R
U	7, 4	2, 11
D	4, 5	10, 1

12

	L	R
U	7, 3	3, 9
D	3, 5	10, 0

8.6.3 *Ochs 1995* (Ochs [1995])

Each subject plays only one of the two games. Those who play game 2 (game 3) take an action 56 (64) times against randomly matched opponents with feedback. Subjects maintain their role as either player 1 or player 2 for the duration of the experiment. Following McKelvey and Palfrey [1995], who note that "The subjects in the Ochs experiments were paid using a lottery procedure," we convert the payoffs described in Ochs [1995] to those in the matrices below before estimation. The exchange rate is \$0.01 (*expected* 1982 Dollars).

2

	L	R
U	1.1141, 0	0, 1.1141
D	0, 1.1141	0.1238, 0

3

	L	R
U	1.1141, 0	0, 1.1141
D	0, 1.1141	0.2785, 0

8.6.4 *GH 2001* (Goeree and Holt [2001])

Each subject takes just one action against an anonymous opponent, and the exchange rate is \$0.01 (year 2001 dollars). "AMP" refers to "asymmetric matching pennies" and "RA" refers to "reversed asymmetry."

AMP

	L	R
U	320, 40	40, 80
D	40, 80	80, 40

RA

	L	R
U	44, 40	40, 80
D	40, 80	80, 40

8.6.5 *NS 2002* (Nyarko and Schotter [2002])

Each subject plays one of four treatments, which differ in terms of details regarding a belief-elicitation procedure. Since there are no significant differences in empirical frequencies of actions across treatments, we pool all data together. Each subject takes an action 60 times against randomly matched opponents with feedback. Subjects maintain their role as either player 1 or player 2 for the duration of the experiment. The exchange rate is \$0.05 (year 2000 Dollars).

NS

	L	**R**
U	$6, 2$	$3, 5$
D	$3, 5$	$5, 3$

8.6.6 *MPS 2018* (Melo et al. [2017])

The data for games 2-4 (see Figure 12) were collected over three sessions. Game 2 was played with 20 rounds twice and 10 rounds once (34, 32, and 30 subject pairs respectively). Game 3 was played with 20 rounds and 10 rounds once (32 and 20 subject pairs respectively). Game 4 was played with 10 rounds with 30 subject pairs. In all cases, there was feedback and random rematching.

8.7 Bootstrap procedure

We estimate the distribution of $\Delta\hat{\mathcal{D}}$ (and hence $\mathcal{P}$) via a simple parametric bootstrap. For each game, each of 5,000 bootstrap samples is a number of draws (equal to the sample size N) from independent Bernoulli distributions (i.e. for each player) with parameters given by the aggregate empirical frequencies.

However, for all studies considered, except for Goeree and Holt [2001], each subject participated in multiple rounds. For a breakdown of how the total number of observations is broken into number of subject-pairs and rounds, see Table 8. Thus, if there is within-subject correlation in actions, then the assumption of independence would artificially lower sampling variation (relative to that of the population) and lead to overstated significance.

To address this concern, we use an alternate bootstrap procedure for robustness. At one extreme, within-subject correlation is perfect and each subject takes the same action in each of his rounds. At the other extreme, each subject's rounds are independent. In the former case, the *effective sample size* should only count one round from each subject. In the latter case, the effective sample size should count every round played by each subject, and the original bootstrap is fine. We thus proxy for within-subject correlation by "throwing away" a fraction $\chi \in [0, 1)$ of each subject's rounds. Specifically, for each subject who played a game R rounds, we re-run the bootstrap as if he only played $\lfloor (1 - \chi)R \rfloor$ rounds, where $\lfloor \cdot \rfloor$ is the "floor" function and larger χ represent more conservative tests (throwing away *at least* a fraction χ of each subject's data). We do not adjust the Goeree and Holt [2001] study in which each subject played only one round.

Tables 4, 5, 6, 7, and 13 assume $\chi = 0$, but we show in Tables 9 and 10 that inference is fairly robust to $\chi \in \{0.5, 0.75, 0.90\}$.

Study	QRE	$\mathcal{P}$			
		$\chi = 0$	$\chi = 0.5$	$\chi = 0.75$	$\chi = 0.90$
MPW 2000	Logit	0.000	0.000	0.012	0.093
SC 2008	Logit	1.000	1.000	1.000	0.985
MPS 2018	Logit	0.808	0.726	0.640	0.547
Survey	Logit	0.000	0.000	0.000	0.000
	Luce	0.048	0.080	0.132	0.231

Table 9: *Revisiting pooled fits.* The conclusions based on the initial bootstrap procedure ($\chi = 0$) remain largely unchanged.

$\chi = 0$

	A	B	C	D
A	0.000	0.019	0.325	0.261
B	0.000	0.996	0.000	0.002
C	0.016	0.389	0.000	0.019
D	0.174	0.097	0.245	0.000

$\chi = 0.5$

	A	B	C	D
A	0.000	0.068	0.366	0.308
B	0.000	0.954	0.002	0.014
C	0.064	0.430	0.002	0.072
D	0.272	0.227	0.341	0.000

$\chi = 0.75$

	A	B	C	D
A	0.000	0.134	0.379	0.327
B	0.002	0.882	0.035	0.049
C	0.149	0.477	0.032	0.167
D	0.326	0.398	0.415	0.001

$\chi = 0.90$

	A	B	C	D
A	0.000	0.215	0.392	0.368
B	0.032	0.744	0.115	0.129
C	0.236	0.539	0.096	0.253
D	0.375	0.516	0.484	0.027

Table 10: *Revisiting out-of-sample performance in* McKelvey et al. [2000]. This replicates Table 5 for different values of χ. That logit transform NBE outperforms logit QRE in out-of-sample tests is moderately-to-very robust, depending on the pair of games.

8.8 Risk aversion

Goeree et al. [2003] construct "game 4" in Figure 15 to "exaggerate the effects of possible risk aversion" by giving each player a "safe" option with payoffs of 200 and 160 and a "risky" option with payoffs of 370 and 10.[55] It is easy to show that under risk neutrality, the data from the game, $(\sigma_U, \sigma_L) = (0.53, 0.65)$, is inconsistent with any QRE, and hence NBE also by Theorem 3. With risk aversion, however, both models can rationalize the data.

Goeree et al. [2003] fit logit QRE to game 4 and games *A*-*D* from McKelvey et al. [2000] by jointly estimating λ and risk aversion parameter r, where the utility function takes the constant relative risk aversion (CRRA) form:

$$u_r(x) = \frac{x^{1-r} - 10^{1-r}}{370^{1-r} - 10^{1-r}}.$$

[55] Relative to how the matrix is given in Goeree et al. [2003], we have switched the rows so that the game has the form of Figure 1.

4

	L	**R**
U	370, 200	10, 370
D	200, 160	160, 10

Figure 15: *Matching Pennies with safe and risky decisions from* Goeree et al. [2003].

Note that utility is normalized so that $u_r(10) = 0$ and $u_r(370) = 1$. To make monetary payoffs comparable across game 4 and games A-D, the payoffs of A-D given in Figure 9 are first multiplied by 10 before the models are fit. Table 11 is essentially a replication of Table 3 of Goeree et al. [2003], but includes NBE for comparison (and minimizes squared distance instead of maximizing likelihood). We find that the fit of game 4 is statistically the same for both models. However, for games A-D, NBE's squared distance (0.0005) is 31% that of QRE (0.0016), and that NBE outperforms QRE is highly significant. Interestingly, the estimated risk aversion parameters are extremely stable, both across games (as Goeree et al. [2003] noted) as well as across models.

Study	Game	Data			NBE		QRE		NBE		QRE		NBE	QRE
		σ_U	σ_L	N	σ_U	σ_L	σ_U	σ_L	$\hat{\tau}$	$\hat{r}$	$\hat{\lambda}$	$\hat{r}$	$\hat{\mathcal{D}}$ ($\mathcal{P}$)	
GHP	4	0.53	0.65	340	0.53	0.67	0.53	0.67	1.091	0.45	6.71	0.45	9.6e7	2.5e6
2003													–	(0.555)
	A	0.64	0.24	1800	0.62	0.25	0.65	0.26	0.691	0.40	22.97	0.44	0.0005	0.0016
MPW	B	0.63	0.24	1200	0.62	0.25	0.57	0.25					–	(0.016)
2000	C	0.59	0.26	1200	0.62	0.25	0.58	0.24						
	D	0.55	0.33	600	0.57	0.33	0.59	0.35						

Table 11: *Parameter Estimates of Models with Risk Aversion.* MPW 2000 = McKelvey et al. [2000] and GHP 2003 = Goeree et al. [2003].

Finally, we show that for CRRA (but not for general utility functions), NBE predictions are invariant to scaling the monetary payoffs. We think this is potentially important as it provides a robustness argument for the prediction of scale invariance in the presence of risk aversion. For an arbitrary normal form game, we now interpret $u_i(a_{ij}, a_{-i})$ as the *monetary* payoff to player i of taking action a_{ij} given the opponents' play a_{-i}. The corresponding utility payoff is simply $u_r(u_i(a_{ij}, a_{-i}))$. After a β-scaling, the utility payoff becomes $u_r(\beta u_i(a_{ij}, a_{-i})) = \beta^{1-r}(u_i(a_{ij}, a_{-i}))$. Using this, it is clear that β drops out of the expression for the ij-response set for all i and j, from which the result is immediate:

$$R_{ij} = \{\sigma'_{-i} : \sum_{a_{-i}\in\Delta_{-i}} u_r(\beta u_i(a_{ij}, a_{-i}))\sigma'_{-i}(a_{-i}) \geq \sum_{a_{-i}\in\Delta_{-i}} u_r(\beta u_i(a_{ik}, a_{-i}))\sigma'_{-i}(a_{-i}) \ \forall k = 1, ..., J(i)\}$$
$$= \{\sigma'_{-i} : \sum_{a_{-i}\in\Delta_{-i}} u_r(u_i(a_{ij}, a_{-i}))\sigma'_{-i}(a_{-i}) \geq \sum_{a_{-i}\in\Delta_{-i}} u_r(u_i(a_{ik}, a_{-i}))\sigma'_{-i}(a_{-i}) \ \forall k = 1, ..., J(i)\}.$$

A consequence of this is that when fitting the logit transform model with risk aversion to games *A*-*D* pooled together in Table 11, we have that the predictions are the same in each of *A*-*C*, consistent with the fact that scale invariance cannot be rejected statistically (see Table 2).

8.9 Utility-money exchange rate conversions

8.9.1 Procedure

To make estimates of logit QRE-λ comparable across studies in the exercise in Section 6.5, we convert all utility-money exchange rates in the studies considered to be consistent with that of McKelvey et al. [2000], the arbitrarily chosen "base study" denominated in U.S. currency.

Given an exchange rate of 1 utile per C_{ft}, where C_{ft} is a numerical amount denominated in the currency of country f in year t, we calculate a "conversion factor" γ_{ft} using the formula

$$\gamma_{ft} = \left(\frac{C_{ft}}{E_{ft}}\right)\left(\frac{P_0}{P_t}\right)\frac{1}{C_0},$$

where E_{ft} is the year t PPP adjustment factor to U.S. dollars, P_0 and P_t are the U.S. CPI price indices in the base year and year t respectively, and C_0 is the dollar value of 1 utile in the base study. Before fitting QRE to a study with an exchange rate of 1 utile per C_{ft}, the payoff matrices are multiplied by γ_{ft}.

Study	Currency (C_{ft})	PPP (E_{ft})	CPI 2000 (P_0)	CPI (P_t)	C_0	γ_{ft}
MPW 2000	\$0.10 (2000 U.S.)	1	172.192	172.192	0.10	1
SC 2008	€0.016 (2008 Germany)	0.820401	172.192	215.254	0.10	0.1560
Ochs 1995	\$0.01 (1982 U.S.)	1	172.192	96.533	0.10	0.1784
GH 2001	\$0.01 (2001 U.S.)	1	172.192	177.042	0.10	0.0973
NS 2002	\$0.05 (2000 U.S.)	1	172.192	172.192	0.10	0.5

Table 12: *Utility-Money Exchange Rate Conversion Factors.*

PPP adjustment factors are from the World Bank[56] and CPI is from the St. Louis Federal Reserve Bank[57]. In all cases, we use annual statistics from the calendar year in which the studies were published (we acknowledge this may be a bit later than when the experiments took place). Table 12 gives the conversion factors as well as their components.

8.9.2 Robustness

After calculating conversion factors, we consider the effect of perturbing the factors on the overall performance of logit QRE in the exercise of Section 6.5. For each study $s \in \{1, 2, 3, 4\}$ (other than the base study), we calculate γ^s via the procedure in the previous section and then we consider all

$$\tilde{\gamma}^s \in \{.5\gamma^s, .6\gamma^s, ..., \gamma^s, ..., 1.4\gamma^s, 1.5\gamma^s\} \equiv \Pi^s,$$

i.e. perturbed factors that are off by as much as 50% in either direction in 10% increments. We calculate the pooled squared distance from fitting QRE to all 21 games across the 5 studies, just as in Section 6.5, for every possible combination of perturbed conversion factors:

$$(\tilde{\gamma}^1, ..., \tilde{\gamma}^4) \in \Pi^1 \times ... \times \Pi^4.$$

Figure 16 plots the histogram of QRE squared distances for all 14,641 factor combinations, as well as vertical lines representing the NBE squared distance which is invariant to the factors, the QRE squared distance presented in Section 6.5, and the minimum QRE squared distance across all factor combinations. Clearly, NBE outperforms QRE for all factors considered. At worst, the squared distance of NBE is no more than 68% that of QRE.

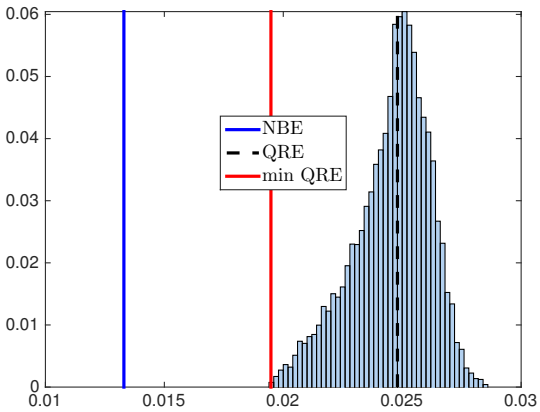

Figure 16: *Robustness to the Utility-Money Exchange Rate.*

[56]http://data.worldbank.org/indicator/PA.NUS.PPP?locations=DE
[57]https://fred.stlouisfed.org/series/CPIAUCSL

8.10 A survey of generalized matching pennies games

Study	Game	Data			NBE	QRE		NBE	QRE	
						Logit	Luce		Logit	Luce
		σ_U	σ_L	N	$\hat{\tau}$	$\hat{\lambda}$	$\hat{\mu}$		$\hat{\mathcal{D}}$	
MPW 2000	A	0.643	0.241	1800	1.411	6.459	0.908	0.0112	0.0174	0.0113
	B	0.630	0.244	1200	1.451	0.800	0.918	0.0142	0.0071	0.0143
	C	0.594	0.257	1200	0.436	2.513	0.254	0.0217	0.0234	0.0220
	D	0.550	0.328	600	1.239	9.579	0.002	0.0134	0.0153	0.0118
SC 2008[58]	1	0.079	0.690	9600	0.960	8.602	0.090	0.0004	0.0017	0.0020
	2	0.217	0.527	9600	0.998	4.012	0.288	0.0018	0.0019	0.0020
	3	0.163	0.793	9600	0.911	7.708	0.125	0.0002	0.0001	0.0003
	4	0.286	0.736	9600	0.689	7.096	0.172	0.0005	0.0001	0.0000
	5	0.327	0.664	9600	0.542	8.084	0.180	0.0003	0.0002	0.0000
	6	0.445	0.596	9600	0.436	8.641	0.002	0.0009	0.0009	0.0002
	7	0.141	0.564	4800	1.473	3.904	0.180	0.0043	0.0058	0.0058
	8	0.250	0.586	4800	0.687	6.143	0.143	0.0077	0.0092	0.0089
	9	0.254	0.827	4800	0.199	17.430	0.057	0.0055	0.0033	0.0030
	10	0.366	0.699	4800	0.689	7.064	0.173	0.0121	0.0094	0.0084
	11	0.331	0.652	4800	0.627	6.986	0.182	0.0007	0.0006	0.0002
	12	0.439	0.604	4800	0.406	9.407	0.167	0.0004	0.0004	0.0002
Ochs 1995	2	0.595	0.258	448	0.451	18.408	0.004	0.0219	0.0030	0.0090
	3	0.542	0.336	516	1.690	14.917	0.004	0.0155	0.0059	0.0062
GH 2001	AMP	0.960	0.160	25	1.280	0.455	0.002	0.0629	0.0228	0.0261
	RA	0.080	0.800	25	1.488	0.977	0.195	0.0266	0.0738	0.0635
NS 2002	NS	0.458	0.390	6720	0.791	3.348	0.158	0.0001	0.0001	0.0001
Pooled	–	–	–	–	1.057	6.233	0.181	**0.0133**	**0.0248**	**0.0148**
($\mathcal{P}$)	–	–	–	–	–	–	–	–	(0.000)	(0.048)

Table 13: *Parameter estimates from a survey of matching pennies games.* MPW 2000 = McKelvey et al. [2000], SC 2008 = Selten and Chmura [2008], Ochs 1995 = Ochs [1995], GH 2001 = Goeree and Holt [2001], and NS 2002 = Nyarko and Schotter [2002].

[58] Brunner et al. [2011] point out that the fit of logit QRE and the empirical frequencies in Game 3 were misreported in Selten and Chmura [2008]. We re-did the analysis using the raw data, which we thank Thorsten Chmura for providing.